Electronic Colloquium on Computational Complexity, Report No. 58 (2024)

Planted Clique Conjectures Are Equivalent

Shuichi Hirahara
National Institute of Informatics
s_hirahara@nii.ac.jp

Nobutaka Shimizu
Tokyo Institute of Technology
shimizu.n.ah@m.titech.ac.jp

March 29, 2024

Abstract

The planted clique conjecture states that no polynomial-time algorithm can find a hidden clique of size $k \ll \sqrt{n}$ in an n-vertex Erdős–Rényi random graph with a k-clique planted. In this paper, we prove the equivalence among many (in fact, *most*) variants of planted clique conjectures, such as search versions with a success probability exponentially close to 1 and with a non-negligible success probability, a worst-case version (the k-clique problem on incompressible graphs), decision versions with small and large success probabilities, and decision versions with adversarially chosen k and binomially distributed k. In particular, we establish the equivalence between the planted clique problem introduced by Jerrum and Kučera and its decision version suggested by Saks in the 1990s. Moreover, the equivalence among decision versions identifies the optimality of a simple edge counting algorithm: By counting the number of edges, one can efficiently distinguish an n-vertex random graph from a random graph with a k-clique planted with probability $\Theta(k^2/n)$ for any $k \leq \sqrt{n}$. We show that for *any* k, no polynomial-time algorithm can distinguish these two random graphs with probability $\gg k^2/n$ *if and only if* the planted clique conjecture holds. The equivalence among search versions identifies the first one-way function that admits a polynomial-time security-preserving self-reduction from exponentially weak to strong one-way functions. These results reveal a detection-recovery gap in success probabilities for the planted clique problem. We also present another equivalence between the existence of a refutation algorithm for the planted clique problem and an average-case polynomial-time algorithm for the k-clique problem with respect to the Erdős–Rényi random graph.

ISSN 1433-8092

Contents

1 Introduction

The *planted k-clique problem*, introduced by Jerrum [Jer92] and Kučera [Kuč95], asks to find a k-clique in an n-vertex random Erdős–Rényi graph with a k-clique planted on average. This is one of the most popular average-case problems, and its hardness assumptions and their variants have many applications in various areas, such as average-case complexity [ERSY22], cryptography [JP00; ABW10; ABIKN23], hardness of approximation [MRS21], game theory [HK11], property testing [AAKMRX07], mathematical finance [ABBG11], and high-dimensional statistics [BR13a; BBH18; BB20].

The planted k-clique problem can be formally stated as follows. Let $\mathcal{G}(n, 1/2)$ denote the distribution of the Erdős–Rényi random graph, i.e., a random graph on n vertices where every pair of vertices is connected by an edge independently with probability $1/2$. Let $\mathcal{G}(n, 1/2, k)$ denote the distribution of the random graph obtained by planting an additional clique of size k in $\mathcal{G}(n, 1/2)$; specifically, to sample a graph from $\mathcal{G}(n, 1/2, k)$, we choose a uniformly random size-k subset C of the n vertices, connect every pair of vertices in C by an edge, and connect every other pair of vertices independently with probability $1/2$. We say that a randomized algorithm A *finds a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability $\varepsilon(n)$* if

$$\Pr_{A,\, G\sim\mathcal{G}(n,1/2,k)}\big[A \text{ outputs a clique of size } k \text{ in } G \text{ on input } G\big] \geq \varepsilon(n),$$

where the probability is over a random graph G drawn from $\mathcal{G}(n, 1/2, k)$ and the internal randomness of A. The task of the planted k-clique problem is to find a k-clique in $\mathcal{G}(n, 1/2, k)$ with high probability $\varepsilon(n)$ (say, $\varepsilon(n) = 1 - o(1)$).

A state-of-the-art algorithm due to Alon, Krivelevich, and Sudakov [AKS98] solves the planted k-clique problem in polynomial time for any $k \geq \Omega(\sqrt{n})$. The maximum size of a clique of an Erdős–Rényi random graph is $(2 + o(1)) \cdot \log_2 n$ with high probability, and thus it is information-theoretically possible to solve the planted k-clique problem as long as $k \gg 2 \log_2 n$; i.e., it can be solved in quasi-polynomial time by using a brute-force search. Yet, no polynomial-time algorithm that solves the planted k-clique problem for $k = o(\sqrt{n})$ is known, despite that many algorithms that improve some aspect of [AKS98] have been developed in the literature [FK00; FR10; DM15; DGP14]. In fact, a large body of research suggests that for $k \ll \sqrt{n}$, the planted k-clique problem cannot be efficiently solved in restricted computational models, such as a fixed-temperature variant of simulated annealing [Jer92; CMZ23], constant depth circuits [Ros08], the statistical query model [FGRVX17], the Lovász–Schrijver semidefinite programming hierarchy [FK03], and the sum-of-squares hierarchy [BHKKMP19]. This leads us to the celebrated conjecture known as *the Planted Clique Conjecture*. There are several ways to formalize the Planted Clique conjecture, each of which leads to a different mathematical statement. The most standard version of the conjecture is as follows.

Conjecture 1.1 (The Planted Clique Conjecture with probability $\varepsilon(n)$)**.** *For any constant $\alpha \in (0, 1/2)$ and for any randomized polynomial-time algorithm A, for all large $n \in \mathbb{N}$ and for $k := n^{1/2-\alpha}$, A cannot find a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability $\varepsilon(n)$. By default, we assume $\varepsilon(n) := 1/2$.*

Conjecture 1.1 is often referred to as a search (or recovery) version because it postulates that no polynomial-time algorithm can recover the hidden k-clique. However, in the literature, it is more popular to assume a stronger variant of the conjecture: a *decision* version.

The decision version of the planted clique problem was suggested by Saks [AKS98; KV02]. The problem is defined as follows. Given as input a random graph G that is drawn from either

$\mathcal{G}(n, 1/2, k)$ or $\mathcal{G}(n, 1/2)$, the task of an algorithm is to decide which distributions the input graph G comes from. As in the case of the search version, it is conjectured that no polynomial-time algorithm can solve the decision version of the planted k-clique problem for $k \ll \sqrt{n}$. The conjecture can be formally stated as follows.

Conjecture 1.2 (A Decision Version of the Planted Clique Conjecture with advantage $\varepsilon(n, k)$)**.** *For any randomized polynomial-time algorithm A, for any constant $\alpha > 0$, for all large $n \in \mathbb{N}$ and for some $k \geq n^{1/2-\alpha}$,*

$$\left| \Pr_{A,\, G \sim \mathcal{G}(n,1/2,k)} \big[A(G) = 1\big] - \Pr_{A,\, G \sim \mathcal{G}(n,1/2)} \big[A(G) = 1\big] \right| < \varepsilon(n, k),$$

where $A(G)$ denotes the output of A on input G. In this case, we say that A cannot distinguish *$\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\varepsilon(n, k)$. By default, we assume $\varepsilon(n, k) := \frac{1}{3}$.*

Decision versions of the planted clique conjecture are extremely useful for determining the average-case hardness of many other problems. Under the planted clique conjectures, a growing body of work shows the average-case hardness of many problems, including sparse principal component detection [BR13a; BBH18], submatrix detection [MW15], the certification of the restricted isometry property [KZ14], community detection [HWX15], and robust sparse mean estimation [BB20].

We note that there are at least three ways to choose the size k of a planted clique. Conjecture 1.2 is formulated by using the *adversarial-k model* [AAKMRX07], in which an adversary can choose any $k \geq n^{1/2-\alpha}$ so that an algorithm fails to distinguish $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$. Another model is the *binomial-k model* [HWX15; BJ23], in which k is distributed according to a binomial distribution. In this model, the task is to distinguish $\widetilde{\mathcal{G}}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$, where $\widetilde{\mathcal{G}}(n, 1/2, k)$ is the distribution of the random graph obtained by planting in $\mathcal{G}(n, 1/2)$ a clique on a random subset of vertices each of which is included in the clique independently with probability $\frac{k}{n}$. In other words, $\widetilde{\mathcal{G}}(n, 1/2, k) \equiv \mathcal{G}(n, 1/2, \mathrm{Bin}(n, k/n))$, where $\mathrm{Bin}(n, p)$ denotes the random variable of the number of success out of n trials, each of which succeeds independently with probability p. Lastly, one may consider the *fixed-k model* [MRS21; ERSY22], in which k is fixed to be $n^{1/2-\alpha}$ for some constant $\alpha \in (0, 1/2)$.

Despite its usefulness and popularity of the planted clique conjectures, many fundamental and basic questions remain open. Could the different ways of choosing the size k of a planted clique result in different conjectures? Is Conjecture 1.1 equivalent to Conjecture 1.2? How much can the probability and the advantage in Conjectures 1.1 and 1.2 be increased or decreased? Identifying an equivalence class for the planted clique conjecture is recognized as a major open problem [BGP23].

We highlight the importance of understanding the optimal advantage of Conjecture 1.2. Although it is common to choose the advantage $\varepsilon(n, k)$ of Conjecture 1.2 to be $o(1)$ in the literature of high-dimensional statistics [BBH18], it is often useful to assume a smaller advantage in the literature of average-case complexity [ERSY22] and cryptography [ABIKN23]. One recent paper postulates that the advantage $\varepsilon(n, k)$ can be made as small as $1/n$ [ERSY22, Conjecture 1], under which the authors of [ERSY22] presented a pseudorandom self-reduction for NP-complete problems. In fact, it is a (not widely known)[1] folklore result that there exists a polynomial-time algorithm that distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\Theta(k^2/n)$ for every $k \leq \sqrt{n}$: The algorithm simply counts the number of the edges of a given graph and outputs 1 if and only if there are many edges; see Appendix B for the details. Thus, [ERSY22, Conjecture 1] is false, which invalidates the

[1] We became aware of this, thanks to Luca Trevisan's blog [Tre18].

results of the paper.[2] In order to avoid relying on such a false conjecture, it is important to base a strong conjecture on another weak and plausible conjecture, such as Conjecture 1.1.

1.1 Our Results

In this paper, we establish the equivalence among many variants of planted clique conjectures, including Conjectures 1.1 and 1.2. Our results are optimal in many cases. For example, we prove that the edge counting algorithm is an optimal distinguisher for $\widetilde{\mathcal{G}}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ in the sense that the existence of a polynomial-time algorithm that distinguishes these distributions with advantage $\gg k^2/n$ for some k falsifies Conjecture 1.1. Since there are many equivalent statements, we present the statements one by one while explaining the significance as well as the ideas of proofs.

Theorem 1.3. *The following (Items 1 to 11) are equivalent.*

1. *Conjecture 1.1 holds.*
2. *Conjecture 1.2 holds.*

That is, we present a search-to-decision reduction for the planted k-clique problem, where $k \approx \sqrt{n}$. Previous search-to-decision reductions either assumed a high success probability $(1 - 1/n^2)$ [AAKMRX07] or decreased the size k of planted cliques significantly [HS23], and thus could not show the equivalence between Conjectures 1.1 and 1.2. Our key idea is that the constant probabilities in Conjectures 1.1 and 1.2 can be significantly amplified to $1 - \exp(-n^{\Omega(1)})$. We formulate the following statements in the adversarial-k model in order to make them as weak as possible.

3. (An exponentially weak decision version) *For any constants $\alpha > 0$ and $\gamma > 0$, any randomized polynomial-time algorithm fails to distinguish $\mathcal{G}(n, 1/2, k)$ from $\mathcal{G}(n, 1/2)$ with advantage $1 - \exp(-n^{\gamma})$ for all large $n \in \mathbb{N}$ and for some $k \geq n^{1/2-\alpha}$.*

Alon, Andoni, Kaufman, Matulef, Rubinfeld, and Xie [AAKMRX07] presented a search-to-decision reduction in a low-error regime of the adversarial-k model, which shows that Item 3 is implied by the following search version.

4. (An exponentially weak search version) *For any constants $\alpha > 0$ and $\gamma > 0$, any randomized polynomial-time algorithm fails to find a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability $1 - \exp(-n^{\gamma})$ for all large $n \in \mathbb{N}$ and for some $k \geq n^{1/2-\alpha}$.*

We view these results as new evidence for the planted clique conjecture because the statements are "close" to the worst-case hardness of the maximum clique problem, which is NP-hard to approximate within a factor of $n^{1-\varepsilon}$ for any constant $\varepsilon > 0$ [Hås99; Zuc07].

In fact, using the notion of Kolmogorov complexity, we obtain an equivalent statement about some worst-case hardness of the maximum clique problem. Informally, the *Kolmogorov complexity* of a string $x \in \{0,1\}^*$ is defined as the size of a shortest program that prints x. The planted clique conjecture is equivalent to the *worst-case hardness* of finding a large clique on the instances whose Kolmogorov complexity is high.

[2] Our results indicate how to fix the bug: Their conjectures and results are valid if the advantage $1/n$ is changed to $1/n^{0.5-1.5\alpha-\gamma}$ for $k := n^{\alpha}$ and for any sufficiently small positive constants α and γ; see Item 9 of Theorem 1.3.

5. (A worst-case search version) *There exists a constant $\alpha > 0$ such that for all $\gamma > 0$, for any randomized polynomial-time algorithm A, for all large $n \in \mathbb{N}$ and for some $k \geq n^{1/2-\alpha}$, there exists an n-vertex graph G that contains a k-clique and has Kolmogorov complexity at least $\binom{n}{2} - \binom{k}{2} + \log_2 \binom{n}{k} - n^\gamma$ such that*

$$\Pr_A[A \text{ outputs a } k\text{-clique in } G \text{ on input } G] \leq \frac{1}{2}.$$

This should be compared with the NP-hardness of approximating the maximum clique problem [Hås99], which shows that, unless $\mathsf{NP} \subseteq \mathsf{BPP}$, for any constant $\varepsilon > 0$, for any randomized polynomial-time algorithm A, for $k := n^{1-\varepsilon}$, there exists an n-vertex graph G with a k-clique such that

$$\Pr_A[A \text{ outputs an } n^\varepsilon\text{-clique in } G \text{ on input } G] \leq \frac{1}{2}.$$

The only essential difference between $\mathsf{NP} \not\subseteq \mathsf{BPP}$ (which is widely believed) and the planted clique conjecture (Item 5, which is stronger and less believed) is that, in the latter, instances are promised to be incompressible, i.e., have high Kolmogorov complexity.

It should be noted that there are barrier results [FF93; BT06b], which show that any worst-case problem outside $\mathsf{NP/poly} \cap \mathsf{coNP/poly}$ cannot be reduced to an average-case analogue of NP via a nonadaptive reduction. Our proofs avoid this barrier by using a non-black-box reduction, as in the recent literature on meta-complexity [Hir18; San20; AFMV06; HN23; LP23] (see, e.g., [Hir22] for a survey).

Next, we explain how much the success probabilities can be decreased. We obtain optimal results in both search and decision versions. For the search version, the probability that a k-clique is found in polynomial time is negligible, i.e., smaller than the reciprocal of any polynomial.

6. (A strong search version) *For any constants $\alpha \in (0, 1/2)$ and $c > 0$, any randomized polynomial-time algorithm fails to find a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability n^{-c} for all large $n \in \mathbb{N}$ and $k := n^{1/2-\alpha}$.*

Item 6 identifies a *detection-recovery gap* in the success probabilities, which is actively studied in the literature [SW22; Mar21; KVWX23; BJ23]: Decision versions of the planted clique problem are "easier" than search versions in that the edge counting algorithm (Appendix B) distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with a non-negligible advantage, whereas the search version cannot be solved with a non-negligible success probability under the planted clique conjecture.

For decision versions, we obtain the optimal result in the binomial-k model, i.e., $\widetilde{\mathcal{G}}(n, 1/2, k) \equiv \mathcal{G}(n, 1/2, \mathrm{Bin}(n, k/n))$.

7. (A strong decision version in the binomial-k model) *For any constant $\gamma > 0$, for any randomized polynomial-time algorithm A, for all large $n \in \mathbb{N}$ and for all $k \in \mathbb{N}$, distributions $\widetilde{\mathcal{G}}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ cannot be distinguished by A with advantage $\frac{k^2}{n} \cdot n^\gamma$.*

By a standard concentration inequality, the size $\mathrm{Bin}(n, k/n)$ of a planted clique in the binomial-k model is concentrated around its mean k. Thus, the edge counting algorithm distinguishes $\widetilde{\mathcal{G}}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\Theta(k^2/n)$ for any $k \leq \sqrt{n}$. Item 7 shows that this advantage cannot be improved by a factor of n^γ for *any* $k \in \mathbb{N}$ if and only if the planted clique conjecture is true.[3] We emphasize that Item 7 determines an optimal advantage for *all* k.

Since the binomial-k model conjecture is stronger than the adversarial-k model conjecture, we also obtain an optimal result in the latter.

[3] If $k > \sqrt{n}$, the advantage $\frac{k^2}{n}$ is larger than 1, in which case Item 7 is vacuously true.

8. (A strong decision version in the adversarial-k model) *For any constant $\gamma > 0$, for any randomized polynomial-time algorithm A, for all large $n \in \mathbb{N}$ and for all $k \in \mathbb{N}$, there exists some $k' \geq k$ such that distributions $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ cannot be distinguished by A with advantage $\frac{k^2}{n} \cdot n^{\gamma}$.*

An immediate corollary is that the advantage $\frac{1}{3}$ of Conjecture 1.2 can be strengthened to $\varepsilon(n, k) = \frac{k^2}{n} \cdot n^{\gamma} = n^{-2\alpha+\gamma}$ for $k := n^{1/2-\alpha}$ and for any small positive constants α and γ.

In the case of the fixed-k model, we obtain an advantage smaller than $o(1)$ for any $k \ll n^{1/3}$; however, for a technical reason, obtaining a small advantage for $n^{1/3} \leq k \ll n^{1/2}$ is left open.

9. (A strong decision version in the fixed-k model) *For all sufficiently small $\gamma > 0$, for any randomized polynomial-time algorithm A, for all large $n \in \mathbb{N}$ and for all k, distributions $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ cannot be distinguished by A with advantage $\varepsilon(n, k)$, where*

$$\varepsilon(n, k) := \min\left\{\sqrt{\frac{k^3}{n}} \cdot n^{\gamma}, 1 - n^{-3}\right\}.$$

In [MRS21; ERSY22], a planted clique conjecture is formulated as follows: For *some* constant $\alpha \in \left(0, \frac{1}{2}\right)$ and for $k := n^{\alpha}$, distributions $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ cannot be distinguished by an efficient algorithm with advantage ε. This is implied by Item 9, which, moreover, upper-bounds the advantage $\varepsilon = \sqrt{\frac{k^3}{n}} \cdot n^{\gamma} = n^{-0.5+1.5\alpha+\gamma}$ for *any* small $\alpha > 0$.

We also obtain a small advantage for distinguishing $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$ for *all* $k \ll \sqrt{n}$, which is used to obtain Item 9.

10. (k vs. $k-1$ in the fixed-k model) *For any constant $\gamma > 0$, for any randomized polynomial-time algorithm A, for all large $n \in \mathbb{N}$ and for all k, distributions $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$ cannot be distinguished by A with advantage $\sqrt{\frac{k^2}{n}} \cdot n^{\gamma}$.*

Finally, we strengthen Item 6 to show that even a $(2+\beta) \cdot \log_2 n$-clique cannot be found in $\mathcal{G}(n, 1/2, k)$ for any constant $\beta > 0$, whose task is often referred to as *partial recovery* [BJ23]. This result is nearly optimal; see Remark 1.4.

11. (A strong search version; partial recovery) *For any positive constants α, β and c, any randomized polynomial-time algorithm fails to find a $(2+\beta) \cdot \log_2 n$-clique in $\mathcal{G}(n, 1/2, k)$ with probability n^{-c} for all large $n \in \mathbb{N}$ and $k := n^{1/2-\alpha}$.*

Remark 1.4. *Although we stated the equivalence statements for* randomized polynomial-time algorithms*, we may also obtain similar equivalent statements for* non-uniform *algorithms and algorithms running in time $n^{o(\log n)}$. Under the assumption that no $n^{o(\log n)}$-time algorithm can solve the planted clique problem [MRS21], the probability n^{-c} in Item 11 of Theorem 1.3 can be strengthened to $n^{-o(\log n)}$. This is asymptotically optimal because the simple algorithm that outputs a random subset of size $3\log_2 n$ finds a $3\log_2 n$-clique in $\mathcal{G}(n, 1/2, k)$ with probability $n^{-3\log_2 n}$.*

Refutation and Average-Case Polynomial Time

Next, we consider a stronger class of algorithms: a refutation algorithm. A refutation algorithm for the planted k-clique problem is an algorithm certifying that most graphs $G \sim \mathcal{G}(n, 1/2)$ do not have any k-clique. This type of an algorithm naturally arises when we consider proof systems,

such as the sum-of-squares hierarchy [BHKKMP19; Pan21], and was explicitly considered in, e.g., [KZ14]. We say that a randomized algorithm A *refutes a planted $k(n)$-clique with probability $\varepsilon(n)$* if $\Pr_A[A(G)=1] \geq \frac{2}{3}$ for every graph G that has a $k(n)$-clique and $\Pr_{G\sim\mathcal{G}(n,1/2)}\left[\Pr_A[A(G)=0] \geq \frac{2}{3}\right] \geq \varepsilon(n)$ for all large $n \in \mathbb{N}$. This definition is based on the adversarial-k model, in that the algorithm accepts every graph G with a k'-clique for some $k' \geq k$. The existence of an efficient refutation algorithm for the planted k-clique problem forms a different equivalence among statements stronger than the negation of each item of Theorem 1.3.

Theorem 1.5. *The following are equivalent.*

1. *There exists a randomized polynomial-time algorithm that refutes a planted $n^{1/2-\alpha}$-clique with probability n^{-c} for some constants $\alpha, c > 0$.*
2. *There exists a randomized polynomial-time algorithm that refutes a planted $n^{1/2-\alpha}$-clique with probability $1-\exp(-n^{\gamma})$ for some constants $\alpha, \gamma > 0$.*
3. *There exists an average-polynomial-time randomized algorithm for the $n^{1/2-\alpha}$-clique problem with respect to the Erdős–Rényi random graph $\mathcal{G}(n,1/2)$ for some constant $\alpha > 0$.*

Here, the $k(n)$-clique problem is the problem of deciding whether a given n-vertex graph has a clique of size $k(n)$. An *average-polynomial-time randomized algorithm M with respect to a distributional problem* $(L, \mathcal{G}(n,1/2))$ [Lev86] is an algorithm such that $\Pr_M[M(G)=L(G)] \geq \frac{2}{3}$ for every input G in the support of $\mathcal{G}(n,1/2)$ and there exists a constant $\varepsilon > 0$ such that $\mathbb{E}_{M,G\sim\mathcal{G}(n,1/2)}[t_M(G)^{\varepsilon}] \leq O(n)$ for all $n \in \mathbb{N}$, where $t_M(G)$ denotes the running time of M on input G; see the excellent survey of Bogdanov and Trevisan [BT06a] for background on average-case complexity.

1.2 Exponentially Weak to Strong One-Way Functions

We present the significance of our results from the viewpoint of cryptography. One of the most fundamental cryptographic primitives is a *one-way function*, which is a polynomial-time-computable function that cannot be inverted on average in polynomial time. There are two notions of one-way function — *strong* one-way and *weak* one-way functions. The former requires that the success probability of inversion by any polynomial-time algorithm is negligible, whereas the latter requires that the success probability to be at most $1-1/p(n)$ for some polynomial p. A strong one-way function is a building block for many important cryptographic primitives, such as a pseudorandom generator [HILL99]. Yao [Yao82] showed that a weak one-way function f can be transformed into another strong one-way function g. However, the transformation is not *security-preserving* [Gol11], i.e., the input of g is much larger than the input of f. Such a reduction is too inefficient to be used in practice. To construct an efficient cryptographic primitive, it is important to prove its security by a security-preserving reduction (see, e.g., [LTW05; BCKR21] and references therein for more background).

There is a natural construction of a one-way function $f^{k\text{-PC}}$ based on the planted clique problem [JP00]. For a parameter $k = k(n)$, consider a family of functions $f_n^{k\text{-PC}} \colon \{0,1\}^{\binom{n}{2}} \times \binom{[n]}{k} \to \{0,1\}^{\binom{n}{2}}$ defined as follows: $f_n^{k\text{-PC}}$ takes the adjacency matrix of a graph G over the vertex set $[n] := \{1,\cdots,n\}$ and a subset $C \subseteq [n]$ of size k, and outputs the graph obtained by adding all the edges inside C to G. Then, inverting $f^{k\text{-PC}}$ is equivalent to solving the planted k-clique problem. As an immediate corollary of Theorem 1.3, we obtain that $f^{k\text{-PC}}$ is an exponentially weak one-way function, i.e., cannot be inverted with probability $1-\exp(-n^{\Omega(1)})$ if and only if $f^{k\text{-PC}}$ is a strong one-way function.

Corollary 1.6. *The function $f^{k\text{-PC}}$ is an exponentially weak one-way function for some constant $\alpha \in (0, 1/2)$ and $k := n^{1/2-\alpha}$ if and only if the function $f^{k\text{-PC}}$ is a strong one-way function for some constant $\alpha \in (0, 1/2)$ and $k := n^{1/2-\alpha}$. Moreover, this equivalence is proved by a* strongly security-preserving reduction *[Gol11], i.e., it maps an instance for the weak one-way function to smaller instances for the strong one-way function; see Remark 2.5.*

To the best of our knowledge, this is the first one-way function that admits a polynomial-time security-preserving self-reduction from an exponentially weak one-way function to a strong one-way function. Note that a generic reduction, such as [Yao82; GILVZ90; CI99; HHR11], transforms an exponentially weak one-way function to an *exponential-time-computable* strong one-way function. Lin, Trevisan, and Wee [LTW05] showed that such an exponential time complexity is necessary for any fully black-box construction. Our results avoid this barrier by exploiting specific structures of the planted k-clique problem. Previously, Bogdanov and Rosen [BR13b] presented an *exponential-time* self-reduction for a one-way function with constant input locality.

Moreover, the security of the one-way function $f^{k\text{-PC}}$ is based on the *worst-case hardness* of the maximum k-clique problem on incompressible instances (Item 5 in Theorem 1.3). This gives an alternative to the lattice-based cryptography [Ajt96], whose security is based on the worst-case hardness of approximating the shortest vector problem. Compared to lattice-based one-way functions, an appealing feature of $f^{k\text{-PC}}$ is that (the decision version of) the maximum k-clique problem on incompressible instances is not known to be in coNP, whereas the approximation of the shortest vector problem is known to be in $\mathsf{NP} \cap \mathsf{coNP}$ [AR05]; the latter is inherent for any black-box reduction techniques [FF93; BT06b; AGGM06; BB15].[4]

We mention that the function $f^{k\text{-PC}}$ itself is not suitable for cryptographic purposes because it can be inverted in quasi-polynomial time; however, as noted by Juels and Peinado [JP00], one may consider a higher edge density $p \approx 1$ of the Erdős–Rényi random graph $\mathcal{G}(n, p)$, in which case the planted clique problem is conjectured to be exponentially hard.

2 Proof Overview

The proof of Theorem 1.3 is outlined in Figure 1. Our proofs are mainly based on the following reductions, each of which is fairly simple.

1. *A Shrinking Reduction.* This takes a large graph G of size N as input and queries a random induced subgraph $G[I]$ of n vertices for a uniformly random size-n subset I of the N vertices, where $n \ll N$. Here, $G[I]$ denotes the induced subgraph of G on the vertex set I. This reduction enables us to obtain the strong decision version in the binomial-k model (Item 7 of Theorem 1.3).

2. *An Embedding Reduction.* This takes a small graph G of size n as input and queries the large random graph obtained by planting G at a random position in $\mathcal{G}(N, 1/2)$, where $n \ll N$. This reduction enables us to obtain the strong search version (Item 6 of Theorem 1.3).

To prove the equivalence among the planted k-clique conjectures for $k \approx \sqrt{n}$, it is crucial to ensure that $n \leq N \leq n^{1+\alpha}$ for a small constant $\alpha > 0$. Our main technical contribution is to analyze the reductions almost optimally, by using a concentration inequality for the probability that a random induced subgraph satisfies a property S, where S is an arbitrary graph property. A

[4] We also mention that the existence of a one-way function can be characterized by some worst-case hardness [HN23; LP23]. This result is not proved by a security-preserving reduction because it uses a universal one-way function, which is quite inefficient.

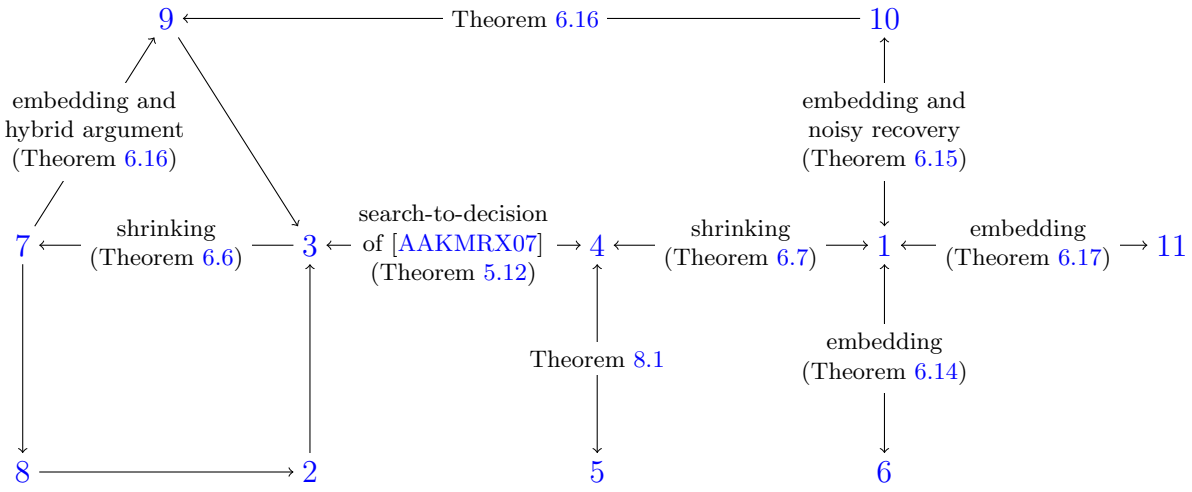

Figure 1: Outline of the proof of Theorem 1.3.

graph property is a subset of n-vertex graphs that is invariant under all the $n!$ permutations of the n vertices. At the core of our proofs is the following concentration inequality.

Theorem 2.1 (see also Theorem 4.4)**.** *Let n, N be positive integers, $k \in \mathbb{N}$, and S be a graph property over n-vertex graphs. For a graph G on the vertex set $[N] := \{1, \cdots, N\}$, define*

$$f(G) := \Pr_{I \sim \binom{[N]}{n}} \big[G[I] \in S\big],$$

where the probability is taken over a size-n subset I of $[N]$ chosen uniformly at random. Let $\mu := \mathbb{E}_{G \sim \mathcal{G}(N,1/2,k)}[f(G)]$. Then, for any $t \geq 0$,

$$\Pr_{G \sim \mathcal{G}(N,1/2,k)} \big[|f(G) - \mu| \geq t\big] \leq 2 \exp\left(-\frac{N(N-1)}{n(n-1)} \cdot 2t^2\right).$$

(The actual concentration inequality is stronger. For example, it holds for any $[0,1]$-valued graph property S; see Theorem 4.4.)

This concentration inequality can be proved by using the result of Gavinsky, Lovett, Saks, and Srinivasan [GLSS15], which shows a general concentration ineqaulity for read-κ families of functions.[5]

2.1 Shrinking Reduction

Using Theorem 2.1, we explain how to obtain the optimal advantage in the decision version of the planted clique problem.

Theorem 2.2. *In the following list, Item 1 $\Longrightarrow$ Item 2 $\Longrightarrow$ Item 3.*

1. (the negation of Item 7 of Theorem 1.3) *For some constant $\gamma > 0$, there exists a randomized polynomial-time algorithm that, for infinitely many $n \in \mathbb{N}$ and for some $k \in \mathbb{N}$, distinguishes $\widetilde{\mathcal{G}}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\frac{k^2}{n} \cdot n^{\gamma}$.*

[5] We note that popular concentration inequalities for the sum of dependent random variables, such as McDiarmid's inequality [McD89] and Janson's inequality [Jan04], show much weaker inequalities than Theorem 2.1 in the setting of Theorem 2.1.

2. (the negation of Item 3 of Theorem 1.3 in the fixed-k model) *For some constants $\alpha \in (0, 1/2)$ and $\gamma > 0$, there exists a randomized polynomial-time algorithm that, for infinitely many $N \in \mathbb{N}$ and for $K := N^{1/2-\alpha}$, distinguishes $\mathcal{G}(N, 1/2, K)$ and $\mathcal{G}(N, 1/2)$ with advantage $1 - \exp(-N^{\gamma})$.*

3. (the negation of Item 3 of Theorem 1.3) *For some constants $\alpha > 0$ and $\gamma > 0$, there exists a randomized polynomial-time algorithm that, for infinitely many $N \in \mathbb{N}$ and for any $K \geq N^{1/2-\alpha}$, distinguishes $\mathcal{G}(N, 1/2, K)$ and $\mathcal{G}(N, 1/2)$ with advantage $1 - \exp(-N^{\gamma})$.*

Proof Sketch. Item 1 $\Rightarrow$ 2: This can be proved by a shrinking reduction. Let A be an algorithm that distinguishes $\widetilde{\mathcal{G}}(n, 1/2, k)$ from $\mathcal{G}(n, 1/2)$ with advantage $\gg \frac{k^2}{n}$. For simplicity, we assume that A is a deterministic algorithm and a graph property, i.e., invariant under all the permutations of the n vertices.[6] The shrinking reduction $\mathcal{R}$ that, using A, distinguishes $\mathcal{G}(N, 1/2, K)$ and $\mathcal{G}(N, 1/2)$ with probability $1 - \exp(-N^{\Omega(1)})$ operates as follows: Given as input a graph G on N vertices, we estimate

$$f(G) := \Pr_{I \sim \binom{[N]}{n}}[A(G[I]) = 1],$$
$$\mu_0 := \Pr_{G_0 \sim \mathcal{G}(n,1/2)}[A(G_0) = 1]$$

by random sampling, and output 0 if and only if $f(G) \approx \mu_0$.

The correctness of the shrinking reduction $\mathcal{R}$ can be proved as follows. Let

$$\mu_1 := \Pr_{G_1 \sim \widetilde{\mathcal{G}}(n,1/2,k)}[A(G_1) = 1].$$

The assumption on A implies that $|\mu_1 - \mu_0| \gg \frac{k^2}{n}$. The concentration inequality (Theorem 2.1) shows that, for an appropriate choice of parameters, with probability $1 - \exp(-N^{\Omega(1)})$ over a choice of $G \sim \mathcal{G}(N, 1/2)$, it holds that $f(G) \approx \mu_0$, in which case $\mathcal{R}$ rejects G. Thus, $\mathcal{R}$ works correctly on most random graphs $\mathcal{G}(N, 1/2)$. It also works correctly on most random graphs $\mathcal{G}(N, 1/2, K)$: By the concentration inequality, with probability $1 - \exp(-N^{\Omega(1)})$ over a choice of $G \sim \mathcal{G}(N, 1/2, K)$, it holds that $f(G) \approx \mu_1'$, where we define

$$\mu_1' := \Pr_{\substack{G \sim \mathcal{G}(N,1/2,K) \\ I \sim \binom{[N]}{n}}}[A(G[I]) = 1].$$

We also have $\mu_1' \approx \mu_1$ by choosing K so that $k = K \cdot \frac{n}{N}$ because the distribution of $G[I]$ for $G \sim \mathcal{G}(N, 1/2, K)$ and $I \sim \binom{[N]}{n}$ is statistically close to $\widetilde{\mathcal{G}}(n, 1/2, k)$. Since $|\mu_1 - \mu_0| \gg \frac{k^2}{n}$, we obtain $f(G) \approx \mu_1' \approx \mu_1 \not\approx \mu_0$, in which case the reduction $\mathcal{R}$ accepts G.

The parameters can be chosen as follows. Let n and k be parameters such that A succeeds. We choose K so that $k = K \cdot \frac{n}{N}$, which ensures that the expected size of a planted clique in $\mathcal{G}(N, K)[I]$ over $I \sim \binom{[N]}{n}$ is k. We also choose N so that $K = N^{1/2-\alpha}$. For $\varepsilon = \frac{k^2}{n} \cdot n^{\gamma}$, Theorem 2.1 shows that with probability $1 - \delta$ over $G \sim \mathcal{G}(N, 1/2)$, where $\delta := 2\exp(-2(\varepsilon N/n)^2)$, it holds that $|f(G) - \mu_0| \leq \varepsilon$. The exponent of δ is proportional to $(\varepsilon N/n)^2 = (n^{\gamma} \cdot Nk^2/n^2)^2 = n^{2\gamma} \cdot (K^2/N)^2 = n^{2\gamma} \cdot N^{-4\alpha} \geq N^{\gamma - 4\alpha}$, where the last inequality holds because $n = kN/K = k \cdot N^{1/2+\alpha} \geq \sqrt{N}$. Thus, if we choose a constant $\alpha := \gamma/8$, the advantage of $\mathcal{R}$ is $\approx 1 - \exp(-N^{\gamma/2})$.

[6] This does not lose any generality by considering $S(G) := \mathbb{E}_{A,\pi}[A(\pi(G))]$, where the expectation is taken over the internal randomness of A and a uniformly random permutation π of the n vertices of G, using the fact that Theorem 2.1 holds for any $[0, 1]$-valued graph property S.

Item 2 $\Rightarrow$ 3: We show that the adversarial-k model and the fixed-k model are equivalent in a low-error regime. The reduction is fairly simple. Given a graph G of n vertices, let V_i denote the first i vertices of G and let G_i be the random graph obtained by planting $G[V_i]$ in $\mathcal{G}(n,1/2)$ at a uniformly random size-i subset of the n vertices. It is easy to observe that if $G \sim \mathcal{G}(n,1/2,k')$ and C is the planted k'-clique, then the marginal distribution of G_i is $\mathcal{G}(n,1/2,|C \cap V_i|)$. In particular, for every $k \le k'$, there exists i such that the marginal distribution of G_i is $\mathcal{G}(n,1/2,k)$, for which an algorithm in the fixed-k model works correctly. Details can be found in Section 5.3. □

We mention that a somewhat similar reduction is implicit in the work of Hazan and Krauthgamer [HK11, Lemma 2.2]. The reduction is quite simple and easier to analyze: Given a graph G of N vertices, it randomly partitions G into n-vertex graphs $G_1, \cdots, G_{N/n}$. This reduction can be used to prove the equivalence between Item 3 (in the binomial-k model) and Item 2 of Theorem 1.3, but is not sufficient for obtaining the optimal advantage of Item 7. See Appendix E for details.

2.2 Embedding Reduction

To obtain the optimal result for the search version (Item 6 of Theorem 1.3), the shrinking reduction is not sufficient; otherwise, we would not be able to show a detection-recovery gap. We use a different reduction to amplify the success probabilities of search algorithms from a non-negligible probability to a constant probability. From a constant probability, the shrinking reduction can amplify the probability to $1-\exp(-n^{\Omega(1)})$.

Theorem 2.3. *The first bullet implies the second in the following.*

- (the negation of Item 6 of Theorem 1.3) *There exist constants $\alpha \in (0,1/2)$ and $c>0$ and a randomized polynomial-time algorithm that, for infinitely many $N \in \mathbb{N}$ and for $k := N^{1/2-\alpha}$, finds a k-clique in $\mathcal{G}(N,1/2,k)$ with probability N^{-c}.*
- (the negation of Item 1 of Theorem 1.3) *There exist a constant $\alpha \in (0,1/2)$ and a randomized polynomial-time algorithm that, for infinitely many $n \in \mathbb{N}$ and for $k := n^{1/2-\alpha}$, finds a k-clique in $\mathcal{G}(n,1/2,k)$ with probability $\frac{1}{2}$.*

Hirahara and Shimizu [HS23] presented an *embedding reduction*, which embeds a given small graph of size n in a large random graph of size N. They used the reduction to show hardness self-amplification for the planted clique problem and proved that the planted clique conjectures remain equivalent even if the constant probabilities in Conjectures 1.1 and 1.2 are changed to arbitrary constants. They analyzed their reductions by using the coupling method of Markov chains and Chebyshev's inequality. However, Chebyshev's inequality does not show a strong concentration inequality, and their analysis is not optimal.

The key to obtaining the optimal search version is to *exponentially* improve their analysis. The embedding reduction $\mathcal{R}_{\mathrm{emb}}$ for parameters $n \le N$ operates as follows: Given as input a graph G of n vertices, it randomly permutes the vertices of G to obtain a graph $\pi(G)$, chooses a random graph $G' \sim \mathcal{G}(N,1/2)$ and a uniformly random size-n subset I of the N vertices, and outputs the N-vertex graph obtained by replacing $G'[I]$ with $\pi(G)$ in G'. We analyze this embedding reduction by using the following concentration inequality.

Theorem 2.4 (see also Lemma 6.13)**.** *Let $\delta, \varepsilon \in [0,1]$ be parameters such that $\varepsilon \ge 4\exp\left(-\frac{N\delta^2}{8n}\right)$. Let S be a function that maps an N-vertex graph to a real number in $[0,1]$ such that*

$$\mathop{\mathbb{E}}_{G\sim\mathcal{G}(N,1/2,k)}[S(G)] \ge \varepsilon.$$

Then, it holds that

$$\Pr_{G\sim\mathcal{G}(n,1/2,k)}\left[\mathop{\mathbb{E}}_{\mathcal{R}_{\mathrm{emb}}}[S(\mathcal{R}_{\mathrm{emb}}(G))] \geq \varepsilon/2\right] \geq 1-\delta,$$

where the expectation is taken over the internal randomness of $\mathcal{R}_{\mathrm{emb}}$.

This result shows that the random variable $\mathbb{E}_{\mathcal{R}_{\mathrm{emb}}}[S(\mathcal{R}_{\mathrm{emb}}(G))]$ is concentrated around its mean $\mathbb{E}_{G\sim\mathcal{G}(n,1/2,k)}[\mathbb{E}_{\mathcal{R}_{\mathrm{emb}}}[S(\mathcal{R}_{\mathrm{emb}}(G))]] = \mathbb{E}_{G\sim\mathcal{G}(N,1/2,k)}[S(G)] \geq \varepsilon$. The previous result of [HS23] showed the same concentration inequality under the stronger assumption that $\varepsilon \geq \frac{4n}{N\delta^2}$, which Theorem 2.4 improves to $\varepsilon \geq 4\exp\left(-\frac{N\delta^2}{8n}\right)$. This improvement is crucial to prove Theorem 2.3 because we set ε to be N^{-c} for a large constant c, in which case the previous work [HS23] does not show the concentration inequality.

Theorem 2.4 enables us to prove Theorem 2.3 as follows. Let A be a randomized polynomial-time algorithm that finds a k-clique in $\mathcal{G}(N,1/2,k)$ with probability N^{-c} for some constant $c>0$. Let S be the function that maps an N-vertex graph G' to the probability that A successfully finds a k-clique in G'. Theorem 2.4 shows that, with probability $\geq \frac{1}{2}$ over a random graph $G \sim \mathcal{G}(n,1/2,k)$, we have $\Pr_{A,\mathcal{R}_{\mathrm{emb}}}[A \text{ finds a } k\text{-clique in } \mathcal{R}_{\mathrm{emb}}(G)] \geq N^{-c}/2$ under the assumption that $N^{-c} \geq \exp(-\Omega(N/n))$, which is satisfied if we choose $n = N/(c'\log N)$ for a sufficiently large constant c'. To solve the planted k-clique problem on $G \sim \mathcal{G}(n,1/2,k)$, we consider the following algorithm: Given as input an n-vertex graph G, we repeat the following $O(N^c)$ times. We run A on $\mathcal{R}_{\mathrm{emb}}(G)$; if A outputs a k-clique in $\mathcal{R}_{\mathrm{emb}}(G)$, then we obtain a k-clique in the original graph G because the k-clique is unique in $\mathcal{R}_{\mathrm{emb}}(G)$ with high probability. This algorithm correctly finds a k-clique in any G such that $\Pr_{A,\mathcal{R}_{\mathrm{emb}}}[A \text{ finds a } k\text{-clique in } \mathcal{R}_{\mathrm{emb}}(G)] \geq N^{-c}/2$.

A brief outline of the proof of Theorem 2.4 is as follows. The main technical lemma is a concentration inequality analogous to Theorem 2.1 in which the subset I is promised to contain the planted location of $\mathcal{G}(n,1/2,k)$ (see Theorem 4.5 for the formal statement). This concentration inequality can be translated into Theorem 2.4 by using the "exchanging lemma" [IJKW10; IJK09; HS23] (Lemma 3.12). The concentration inequality (Theorem 4.5) can be proved as follows. We first prove an upper bound of the moment generating function when the planted location is fixed, and then use Jensen's inequality to take the average over all planted locations. The upper bound of the moment generating function can be proved by combining the proof ideas of [GLSS15] and the transportation method [BLM13, Chapter 8], which is a powerful method to bound the moment generating function using information-theoretic inequalities; see Appendix A for details.

Remark 2.5. *Although the embedding reduction is not security-preserving, the implication from Item 4 to Item 6 in Theorem 1.3 is proved by the composition of the shrinking reduction and the embedding reduction, which is security-preserving as mentioned in Corollary 1.6. The reason is that the shrinking reduction* decreases *the size n of an instance to $n^{1-\alpha}$ for a small constant $\alpha>0$, and the embedding reduction increases the size $n^{1-\alpha}$ of instances to $O(n^{1-\alpha}\log n^{1-\alpha}) \leq n^{1-\alpha/2}$, which is smaller than the size n of the original instance.*

2.3 Decision Versions in the Fixed-k Model

It is instructive to compare the shrinking and embedding reductions.

1. The shrinking reduction can amplify the success probability from $\omega(\frac{k^2}{n})$ to $1-\exp(-n^{\Omega(1)})$ in the binomial-k model. However, the size of a planted clique in the output of the reduction on $\mathcal{G}(n,1/2,k)$ is not necessarily fixed and cannot be used in the fixed-k model.

2. The embedding reduction can amplify the success probability from either $\omega(\frac{k^2}{n})^{1/2}$ in the decision version or $1/\mathsf{poly}(n)$ in the search version to $1-n^{-\gamma}$ for a sufficiently small constant $\gamma > 0$. The reduction does not change the value of k and can be used in the fixed-k model. However, the success probability $1-n^{-\gamma}$ is too small to take a union bound, which poses a technical challenge in the fixed-k model.

Now, we explain how to obtain Item 9 of Theorem 1.3, i.e., the strong decision version in the fixed-k model. Assume that there exists an efficient algorithm A that distinguishes $\mathcal{G}(n,1/2,k)$ and $\mathcal{G}(n,1/2)$ with advantage $\gg \sqrt{\frac{k^3}{n}}$ for some k. The main idea for obtaining Item 9 is to use A to prove the negation of Item 7, i.e., to distinguish $\widetilde{\mathcal{G}}(n,1/2,k)$ and $\mathcal{G}(n,1/2)$. Since $\mathrm{Bin}(n,k/n)$ is in $k \pm O(\sqrt{k})$ with high probability, the algorithm A also distinguishes $\widetilde{\mathcal{G}}(n,1/2,k) = \mathcal{G}(n,1/2,\mathrm{Bin}(n,k/n))$ from $\mathcal{G}(n,1/2)$ with probability $\gg \frac{k^2}{n}$ if

$$\Pr[A(\mathcal{G}(n,1/2,k+a))=1] \approx \Pr[A(\mathcal{G}(n,1/2,k))=1]$$

for every $a \in \mathbb{Z}$ such that $-O(\sqrt{k}) \le a \le O(\sqrt{k})$. To this end, we show Item 10, that is, that $\mathcal{G}(n,1/2,k-1)$ and $\mathcal{G}(n,1/2,k)$ cannot be distinguished with advantage $\gg \sqrt{\frac{k^2}{n}}$. By a hybrid argument, Item 10 enables us to show that for every $|a| \ll \sqrt{k}$,

$$\begin{aligned}
&|\Pr[A(\mathcal{G}(n,1/2,k+a))=1] - \Pr[A(\mathcal{G}(n,1/2))=1]| \\
\ge\ &|\Pr[A(\mathcal{G}(n,1/2,k))=1] - \Pr[A(\mathcal{G}(n,1/2))=1]| \\
-\ &|\Pr[A(\mathcal{G}(n,1/2,k))=1] - \Pr[A(\mathcal{G}(n,1/2,k+a))=1]| \\
\gg\ &\sqrt{\frac{k^3}{n}} - \sqrt{\frac{k^2}{n}} \cdot a \gg \frac{k^2}{n}.
\end{aligned}$$

In this way, we can obtain a contradiction to Item 7. See Theorem 6.16 for the details.

Thus, it remains to prove Item 10. The proof consists of two steps. These steps are based on beautiful ideas of Feige and Krauthgamer [FK00] and Dekel, Gurel-Gurevich, and Peres [DGP14], respectively.

Step 1. k versus $(k-1)$ to Noisy Recovery

Feige and Krauthgamer [FK00] presented a simple search-to-decision reduction from the planted k-clique problem to the decision problem of distinguishing $\mathcal{G}(n,1/2,k)$ and $\mathcal{G}(n,1/2,k-1)$ with advantage $1-n^{-2}$. Their reduction makes n queries $(G \setminus v_1, \cdots, G \setminus v_n)$ given a graph $G=(V,E)$ with $V=\{v_1,\cdots,v_n\}$; thus, to guarantee that all the queries are answered correctly, the advantage must be very close to 1.

In our case, however, the advantage cannot be assumed to be sufficiently close to 1. The embedding reduction can only increase the advantage to $1-n^{-\alpha}$ for a small constant $\alpha > 0$ without decreasing the size $k \approx \sqrt{n}$ of a planted clique too much. Because of the small advantage, we cannot simply use a union bound over all n queries. However, we do show that the reduction of [FK00] can be used to solve some non-trivial tasks, which we call *noisy recovery*.

The noisy recovery task for the planted k-clique problem is defined as follows. Given as input a graph $G \sim \mathcal{G}(n,1/2,k)$ with a planted k-clique C, the task is to output an induced subgraph $G[C']$ for some subset C' of the vertices of G such that $|C' \cap C| \ge (1-o(1)) \cdot k$ and $|C'| \le n^{1-\alpha}$ for a small constant $\alpha > 0$. That is, the task is to remove a $(1-n^{-\alpha})$-fraction of the vertices from G, while keeping the clique of size to be $\approx k$.

Step 2. Noisy Recovery to Exact Recovery

Then, we show that the noisy recovery algorithm can be used to solve the search version of the planted clique problem. Let A be the noisy recovery algorithm. A naïve attempt to construct a search algorithm using A would be as follows. Given a graph $G \sim \mathcal{G}(n, 1/2, k)$, we run A on G to obtain an induced subgraph G', and then we run existing algorithms, such as [AKS98], to recover a $(1 - o(1)) \cdot k$-clique in G'. Note that since the number of vertices of G' is at most $n^{1-\alpha}$, we can hope that the existing algorithm can recover a clique of size at least $\sqrt{n^{1-\alpha}} = n^{1/2-\alpha/2} \ll k$.

Unfortunately, this naïve approach does not work because of the subtlety of average-case algorithms. Intuitively, the output G' of the noisy recovery algorithm should be distributed according to $\mathcal{G}(n', 1/2, k')$ for some $n' \ll n$ and $k' \approx k$, in which case the existing algorithms for the planted clique problem work correctly. However, there is no guarantee that G' is distributed according to such a nice distribution, as the algorithm A is, in principle, adversarial. There are several algorithms, such as [FK00; BKS23], for the planted clique problem in *semi-random models*, in which an input graph is corrupted by an adversary; however, the existing algorithms do not seem to be robust with respect to the noisy recovery algorithm A.

Somewhat surprisingly, the issue can be resolved by a simple and elegant idea of Dekel, Gurel-Gurevich, and Peres [DGP14] (similar ideas were implicit in earlier works, e.g., [HK11; FR10]). The main issue is that the "adversary" A is too powerful in that it can see the whole input graph. We restrict the power of the adversary by randomly partitioning an input graph G into two subgraphs G_1 and G_2, and running A only on G_1. Let G_1' be the output of A on G_1. Then, the edges between G_1' and G_2 in the original graph G are not seen by A, so they can be seen as a purely random bipartite graph independent of A. This enables us to reduce the task to the following problem: Given as input a random bipartite graph on the vertex sets V_1' and V_2, in which a biclique of size $k \approx n^{1/2-\alpha'}$ is planted, where $|V_1'| \le n^{1-\alpha}$ and $|V_2| = \Theta(n)$, the task is to find a biclique in the bipartite graph for small constants $\alpha > 2\alpha' > 0$. This can be solved by the simple degree counting algorithm of Kučera [Kuč95] because $\sqrt{|V_1'|} \ll k$.

2.4 Organization

The rest of the paper is devoted to proving Theorems 1.3 and 1.5.

- In Section 3, we define notations and introduce several tools that are important in our proof.
- In Section 4, we prove the concentration inequality about random induced subgraphs.
- In Section 5, we present reductions for several settings of the planted clique problem. In particular, we prove the equivalence between Item 3 and Item 4 of Theorem 1.3.
- In Section 6, we use the shrinking reduction and the embedding reduction to prove hardness amplification results. In particular, we prove the equivalence of Items 2, 3, 7 and 8 (Theorem 6.6), Items 1 and 4 (Theorem 6.7), Items 1 and 6 (Theorem 6.14), Items 1 and 10 (Theorem 6.15), Items 7 and 9 (Theorem 6.16), and Items 1 and 11 (Theorem 6.17).
- In Section 7, we present hardness amplification for refuting planted cliques and prove Theorem 1.5.
- In Section 8, we present a problem whose worst-case hardness characterizes the Planted Clique conjecture by proving the equivalence between Items 4 and 5.
- In Section 9, we combine the equivalence results and prove Theorem 1.3.

3 Preliminaries

All logarithms in this paper are natural logarithms unless otherwise stated. For $n \in \mathbb{N}$, we write $[n] = \{1, \dots, n\}$. We use $x \sim \mathcal{D}$ to denote that x is drawn from a distribution $\mathcal{D}$. For a set S, we also use $x \sim S$ to denote that x is drawn uniformly at random from S.

A graph $G = (V, E)$ is a pair of a finite set V and $E \subseteq \binom{V}{2}$, where $\binom{V}{2}$ is the set of unordered pairs from V. For a vertex $v \in V$, let $\Gamma_G(v) = \{w \in V \mid \{v, w\} \in E\}$ denote the set of neighbors of v (which does not include v). If G is clear from the context, we simply write $\Gamma(v)$. For two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, let $G_1 \cup G_2 = (V_1 \cup V_2, E_1 \cup E_2)$. For $S, T \subseteq V$, let $E(S, T) = \left\{ \{s, t\} \in \binom{V}{2} \;\middle|\; s \in S, t \in T \right\}$ denote the set of edges such that one endpoint is in S and the other is in T.We usually assume that the vertex set of an n-vertex graph is $[n]$. For a graph $G = ([n], E)$ and $S = \{s_1, \dots, s_t\} \subseteq [n]$ with $s_1 < \cdots < s_t$, let $G[S] = ([t], E_S)$ denote the subgraph induced by S, where $E_S = \{\{a, b\} \mid \{s_a, s_b\} \in E\}$. For a graph $G = (V, E)$ and a permutation $\pi \colon V \to V$, let $\pi(G) = (V, E_\pi)$ be the graph obtained by shuffling G using π, i.e., $E_\pi = \{\{\pi(u), \pi(v)\} \mid \{u, v\} \in E\}$.

For $n \in \mathbb{N}$ and $p \in [0, 1]$, let $\mathrm{Bin}(n, p)$ denote the binomial distribution where n is the number of trials and p is the success probability; that is, $\Pr[\mathrm{Bin}(n, p) = k] = \binom{n}{k} p^k (1-p)^{n-k}$ for any $k \in \{0, \cdots, n\}$. For $n \in \mathbb{N}$, let $\mathcal{G}(n, 1/2)$ be the distribution of an n-vertex random graph, where each pair $\{u, v\}$ forms an edge independently with probability $1/2$. For a finite set C, let $K_C = (C, \binom{C}{2})$ be the clique on the vertex set C. Let $\mathcal{G}(n, 1/2, k)$ be the distribution of $G \cup K_C$ for $G \sim \mathcal{G}(n, 1/2)$ and $C \sim \binom{[n]}{k}$. We also consider the planted clique whose size is drawn from a binomial distribution. Let $\widetilde{\mathcal{G}}(n, 1/2, k)$ be the distribution of $G \cup K_C$, where $G \sim \mathcal{G}(n, 1/2)$ and $C \subseteq [n]$ is the random subset that contains each $i \in [n]$ with probability k/n independently. The set C is called a *planted location*. It is well known that $\mathcal{G}(n, 1/2, k)$ contains a unique k-clique with high probability (see, e.g., [HS23, Lemma 6.1]).

Lemma 3.1. *For any $n, k \in \mathbb{N}$, $G \sim \mathcal{G}(n, 1/2, k)$ contains a unique k-clique with probability at least $1 - 2kn2^{-k/2}$.*

For distributions $\mathcal{D}_1$ and $\mathcal{D}_2$, we say that a randomized algorithm A *distinguishes $\mathcal{D}_1$ and $\mathcal{D}_2$ with advantage ε* if

$$\left| \Pr_{A, x \sim \mathcal{D}_1}[A(x) = 1] - \Pr_{A, x \sim \mathcal{D}_2}[A(x) = 1] \right| \geq \varepsilon.$$

Lemma 3.2 (Jensen's Inequality)**.** *Let X be a $\mathbb{R}^n$-valued random variable and $f \colon \mathbb{R}^n \to \mathbb{R}$ be a convex function. Then, $\mathbb{E}[f(X)] \geq f(\mathbb{E}[X])$.*

Lemma 3.3 (The Chernoff Bound)**.** *Let $X_1, \dots, X_n \in [0, 1]$ be independent random variables and $\mu = \frac{1}{n} \sum_{i \in [n]} \mathbb{E}[X_i]$. Then, for any $0 \leq t \leq \mu$,*

$$\Pr\left[\frac{1}{n} \sum_{i \in [n]} X_i \geq \mu + t \right] \leq \exp\left(-\frac{nt^2}{3\mu} \right),$$
$$\Pr\left[\frac{1}{n} \sum_{i \in [n]} X_i \leq \mu - t \right] \leq \exp\left(-\frac{nt^2}{3\mu} \right).$$

Moreover, for any $t \geq 0$, the upper bounds can be replaced by $\exp(-2nt^2)$.

3.1 Information Theory

For a random variable X, let $\mathsf{supp}(X) = \{x\colon \Pr[X = x] > 0\}$ be the support of X. For random variables X_1 and X_2 over Ω, the *KL divergence* of X_1 and X_2 is defined as

$$\mathrm{KL}(X_1 \parallel X_2) = \sum_{x \in \mathsf{supp}(X_1)} \Pr[X_1 = x] \cdot \log \frac{\Pr[X_1 = x]}{\Pr[X_2 = x]},$$

where we define $\mathrm{KL}(X_1 \parallel X_2) = \infty$ if $\mathsf{supp}(X_1) \not\subseteq \mathsf{supp}(X_2)$. It is well known that the KL divergence is always non-negative. We sometimes identify a random variable with its distribution. For example, if μ, ν are distributions of random variables X, Y, then we write $\mathrm{KL}(\mu \parallel \nu)$ to mean $\mathrm{KL}(X \parallel Y)$.

For $p, q \in [0,1]$, let $\mathrm{KL}_{\mathrm{b}}(p \parallel q) = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$ denote the KL divergence of the two Bernoulli random variables with bias p and q. We invoke the following fact, which is a special case of Pinsker's inequality.

Fact 3.4. *For any $p, q \in [0,1]$, we have $\mathrm{KL}_{\mathrm{b}}(p \parallel q) \geq 2(p-q)^2$.*

Let (X, Y) be a random variable over a product space $\Omega_X \times \Omega_Y$. For $y \in \Omega_Y$, let $X|_{Y=y}$ denote the random variable X conditioned on $Y = y$.

Definition 3.5 (Conditional KL Divergence)**.**

$$\begin{aligned}\mathrm{KL}(X_1|Y_1 \parallel X_2|Y_2) &= \mathop{\mathbb{E}}_{y \sim Y_1}[\mathrm{KL}(X_1|_{Y_1=y} \parallel X_2|_{Y_2=y})] \\ &= \sum_{y \in \Omega_Y} \Pr[Y_1 = y] \sum_{x \in \Omega_X} \Pr[X_1 = x \mid Y_1 = y] \cdot \log \frac{\Pr[X_1 = x \mid Y_1 = y]}{\Pr[X_2 = x \mid Y_2 = y]}.\end{aligned}$$

Lemma 3.6 (Chain Rule of KL Divergence)**.** *Let $(X_1, Y_1), (X_2, Y_2)$ be random variables over a product space $\Omega_X \times \Omega_Y$. Then,*

$$\mathrm{KL}((X_1, Y_1) \parallel (X_2, Y_2)) = \mathrm{KL}(Y_1 \parallel Y_2) + \mathrm{KL}(X_1|Y_1 \parallel X_2|Y_2)$$

Lemma 3.7 (Data Processing Inequality)**.** *Let X, Y be random variables over Ω and f be a probabilistic function over Ω (i.e., $f(x)$ is a random variable for each $x \in \Omega$). Then, we have $\mathrm{KL}(f(X) \parallel f(Y)) \leq \mathrm{KL}(X \parallel Y)$, where $f(X)$ denotes the random variable that is sampled from $f(x)$ for $x \sim X$.*

Corollary 3.8. *Let X, Y be $[0,1]$-valued random variables. Then, $\mathrm{KL}(X \parallel Y) \geq \mathrm{KL}_{\mathrm{b}}(\mathbb{E}[X] \parallel \mathbb{E}[Y])$.*

Proof. For $x \in [0,1]$, let $B(x)$ be the probabilistic function that outputs 1 with probability x and 0 with probability $1-x$. Note that $B(X)$ is identical to the Bernoulli random variable $\mathrm{Ber}(\mathbb{E}[X])$ since

$$\Pr[B(X) = 1] = \sum_x \Pr[B(X) = 1 \mid X = x] \cdot \Pr[X = x] = \sum_x x \cdot \Pr[X = x] = \mathbb{E}[X].$$

It follows from the data processing inequality (Lemma 3.7) that

$$\mathrm{KL}(X \parallel Y) \geq \mathrm{KL}(B(X) \parallel B(Y)) = \mathrm{KL}_{\mathrm{b}}(\mathbb{E}[X] \parallel \mathbb{E}[Y]). \qquad \square$$

Lemma 3.9 (Conditioning Increases Divergence for Product Measure)**.** *Let $(X_1, Y_1), (X_2, Y_2)$ be random variables, and suppose that X_2 and Y_2 are independent. Then,*

$$\mathrm{KL}(X_1|Y_1 \parallel X_2|Y_2) \geq \mathrm{KL}(X_1 \parallel X_2).$$

Proof. By the convexity of KL divergence [CT06, Section 2.7], we have

$$\begin{aligned}\mathrm{KL}(X_1|Y_1 \parallel X_2|Y_2) &= \mathop{\mathbb{E}}_{y\sim Y_1}[\mathrm{KL}(X_1|_{Y_1=y} \parallel X_2|_{Y_2=y})] \\ &\geq \mathrm{KL}\left(\mathop{\mathbb{E}}_{y\sim Y_1}[X_1|_{Y_1=y}] \,\middle\|\, \mathop{\mathbb{E}}_{y\sim Y_1}[X_2|_{Y_2=y}]\right) \\ &= \mathrm{KL}(X_1 \parallel X_2),\end{aligned}$$

where $\mathbb{E}_{y\sim Y_1}[X_1|_{Y_1=y}]$ denotes the mixture of the distributions $X_1|_{Y_1=y}$ over $y \sim Y_1$, which is identical to the marginal distribution of X_1, and $\mathbb{E}_{y\sim Y_1}[X_2|_{Y_2=y}]$ is identical to X_2 because of the independence between X_2 and Y_2. □

3.2 Sampler

Following [IJK09], we introduce the notion of sampler. A (δ, ε)-sampler is a pair of random variables (X, Y) such that any function $S\colon \mathsf{supp}(Y) \to [0,1]$ cannot "distinguish" Y and $Y|_{X=x}$ with advantage ε for a $(1-\delta)$-fraction of $x \sim X$ in the following sense.

Definition 3.10. *We say that a pair of random variables (X, Y) is a* (δ, ε)-sampler *if, for any function $S\colon \mathsf{supp}(Y) \to [0,1]$, we have*

$$\Pr_{x\sim X}[|\mathbb{E}[S(Y) \mid X = x] - \mathbb{E}[S(Y)]| \geq \varepsilon] \leq \delta.$$

We say that a family of random variables $X = \{X_n\}_{n\in\mathbb{N}}$ is samplable in polynomial time if there exists a polynomial-time randomized algorithm M such that $\Pr_M[M(1^n) = x] = \Pr[X_n = x]$ for all $n \in \mathbb{N}$ and all x. We denote the family X by X_n when it is clear from the context. For example, $\mathcal{G}(n, 1/2)$ is samplable in polynomial time.

Lemma 3.11 (Hardness Amplification for Decision Problems)**.** *Let X_1, X_2 be independent random variables such that X_1 is samplable in polynomial time. Let $\mathcal{R}$ be a polynomial-time randomized algorithm such that, for $Y_i = \mathcal{R}(X_i)$, the pairs (X_1, Y_1) and (X_2, Y_2) are $\left(\frac{\delta}{4}, \frac{\varepsilon}{6}\right)$-samplers. Let n be the size of X_i and N be the size of Y_i. Suppose there exists a polynomial-time randomized algorithm A such that*

$$|\Pr[A(Y_1) = 1] - \Pr[A(Y_2) = 1]| \geq \varepsilon,$$

where the probability is taken over the internal randomness of A and Y_i. Then, there exists a randomized algorithm A' that runs in time $\mathsf{poly}(n, N, 1/\varepsilon, \log(1/\delta))$ and satisfies

$$\Pr[A'(X_1) = 1] - \Pr[A'(X_2) = 1] \geq 1 - \delta.$$

Proof. Let $p(x) = \Pr[A(\mathcal{R}(x)) = 1]$ and $\mu_i = \mathbb{E}[p(X_i)]$ for each $i \in \{1, 2\}$. Note that $|\mu_1 - \mu_2| \geq \varepsilon$ by the assumption of A. On input x, the algorithm A' runs as follows:

1. Let $T = c\varepsilon^{-2}\log(1/\delta)$ for a sufficiently large constant $c > 0$ and $y_1, \ldots, y_T$ be independent samples from $\mathcal{R}(x)$. Compute $\hat{p}(x) := \frac{1}{T}\sum_{i\in[T]} A(y_i)$.
2. Let $y'_1, \ldots, y'_T$ be independent samples from Y_1. Compute $\hat{\mu}_1 := \frac{1}{T}\sum_{i\in[T]} A(y'_i)$.
3. If $|\hat{p}(x) - \hat{\mu}_1| \leq \frac{\varepsilon}{3}$, then output 1. Otherwise, output 0.

In what follows, we prove $\Pr[A'(X_1)=1] \geq 1-\frac{\delta}{2}$ and $\Pr[A'(X_2)=1] \leq \frac{\delta}{2}$.

By the Chernoff bound (Lemma 3.3), with probability $1-4\exp(-\frac{T\varepsilon^2}{108}) \geq 1-\frac{\delta}{4}$ (over the randomness in Step 1 and 2), we have $|p(x)-\hat{p}(x)| \leq \frac{\varepsilon}{6}$ and $|\mu_1-\hat{\mu}_1| \leq \frac{\varepsilon}{6}$. By applying the sampler property of (X_i, Y_i) for $S(y) := \Pr_A[A(y)=1]$, we have $\Pr_{x\sim X_i}\left[|p(x)-\mu_i| > \frac{\varepsilon}{6}\right] \leq \frac{\delta}{4}$. Therefore, for a $(1-\delta/4)$-fraction of $x \sim X_i$, we have $|\hat{p}(x)-\hat{\mu}_i| \leq \frac{\varepsilon}{3}$ with probability $1-\frac{\delta}{4}$ over the internal randomness of A'. Thus, we obtain $\Pr[A'(X_1)=1] \geq \left(1-\frac{\delta}{4}\right)^2 \geq 1-\frac{\delta}{2}$.

Similarly, for a $(1-\delta/4)$-fraction of $x \sim X_2$, with probability $1-\frac{\delta}{4}$ over the internal randomness of A', we have

$$\begin{aligned}\varepsilon &\leq |\mu_1-\mu_2| \\ &\leq |\hat{\mu}_1-\hat{p}(x)| + |\mu_1-\hat{\mu}_1| + |\hat{p}(x)-p(x)| + |p(x)-\mu_2| \\ &\leq |\hat{\mu}_1-\hat{p}(x)| + \frac{\varepsilon}{2}.\end{aligned}$$

Thus, we have $\Pr[A'(X_2)=1] \leq 1-\left(1-\frac{\delta}{4}\right)^2 \leq \frac{\delta}{2}$. □

In the context of hardness amplification [IJKW10; IJK09; HS23], it is often easy to prove that (Y, X) is a sampler. For example, in the direct product theorem, for an input distribution $\mathcal{D}$, we consider the pair of inputs (X, Y) obtained by $X \sim \mathcal{D}$ and then $Y = (X_1, \ldots, X_{i-1}, X, X_{i+1}, \ldots, X_k)$ for $(X_1, \ldots, X_k) \sim \mathcal{D}^k$ and $i \sim [k]$. It is easy to see that (Y, X) is a sampler: Fix a function $S\colon \mathsf{supp}(X) \to [0,1]$. For any $y = (x_1, \ldots, x_k)$, we have $\mathbb{E}[S(X) \mid Y = y] = \frac{1}{k}\sum_{i\in[k]} S(x_i)$. If $y \sim \mathcal{D}^k$, then this quantity is the sum of independent random variables, which concentrates around its mean. In fact, the sampler property of (Y, X) is enough to ensure the sampler property of (X, Y). See Appendix C for the proof.

Lemma 3.12 (Exchange Lemma)**.** *If (Y, X) is a $\left(\frac{\varepsilon}{2}, \frac{\delta\varepsilon}{8}\right)$-sampler, then, (X, Y) is a (δ, ε)-sampler.*

4 Concentration Inequalities

4.1 Concentration of Random Shrinking

For an M-bit string $x = (x_1, \ldots, x_M) \in \{0,1\}^M$ and a subset of indices $I = \{i_1, \ldots, i_\ell\} \subseteq [M]$ with $i_1 < \cdots < i_\ell$, let $x_I = (x_{i_1}, \ldots, x_{i_\ell}) \in \{0,1\}^\ell$ be the substring induced by I. Each subset I is associated with a function $S_I\colon \{0,1\}^\ell \to [0,1]$. Gavinsky, Lovett, Saks, and Srinivasan [GLSS15] proved a concentration of the averaging function $f(X) = \mathbb{E}_I[S_I(X_I)]$, where X is a random M-bit string drawn from a product distribution and the expectation is taken over a random subset I. We state their result in the notation convenient for us.

Theorem 4.1 ([GLSS15])**.** *Let X be a random variable that is drawn from a product distribution $\mathcal{D}$ over $\{0,1\}^M$. Let $\mathcal{I}$ be a distribution over subsets of $[M]$, where each $I \in \mathsf{supp}(\mathcal{I})$ is associated with a function $S_I\colon \{0,1\}^{|I|} \to [0,1]$. Let $\rho := \max\{\Pr_{I\sim\mathcal{I}}[i \in I] \mid i \in [M]\}$. Define $f\colon \{0,1\}^M \to [0,1]$ by $f(x) = \mathbb{E}_{I\sim\mathcal{I}}[S_I(x_I)]$. Let $\mu = \mathbb{E}[f(X)]$. Then, for any $t \geq 0$, we have*

$$\begin{aligned}\Pr[f(X) \geq \mu + t] &\leq \exp\left(-\frac{1}{\rho}\cdot \mathrm{KL_b}(\mu+t \parallel \mu)\right), \\ \Pr[f(X) \leq \mu - t] &\leq \exp\left(-\frac{1}{\rho}\cdot \mathrm{KL_b}(\mu-t \parallel \mu)\right).\end{aligned}$$

For completeness, we include a proof of Theorem 4.1.

Lemma 4.2. *Under the same settings of Theorem 4.1, for every random variable Y such that* $\mathsf{supp}(Y) \subseteq \mathsf{supp}(X)$,

$$\mathop{\mathbb{E}}_{I\sim\mathcal{I}}[\mathrm{KL}(Y_I \parallel X_I)] \le \rho \cdot \mathrm{KL}(Y \parallel X).$$

Proof. For a string $x = (x_1, \dots, x_n) \in \{0,1\}^\ell$ and $i \in [\ell]$, let $x^{\le i} = (x_1, \dots, x_i) \in \{0,1\}^i$ denote the prefix of x of length i.

For every subset $I = \{i_1, \dots, i_\ell\} \subseteq [M]$ with $i_1 < \cdots < i_\ell$, we have

$$\begin{aligned}
\mathrm{KL}(Y_I \parallel X_I) &= \sum_{j\in[\ell]} \mathrm{KL}\Big(Y_I^{\le j} | Y_I^{\le j-1} \;\Big\|\; X_I^{\le j} | X_I^{\le j-1}\Big) && \because \text{chain rule (Lemma 3.6)} \\
&\le \sum_{j\in[\ell]} \mathrm{KL}\Big(Y_I^{\le j} | Y^{\le i_j-1} \;\Big\|\; X_I^{\le j} | X^{\le i_j-1}\Big) && \because \mathcal{D} \text{ is a product distribution (Lemma 3.9)} \\
&= \sum_{j\in[\ell]} \mathrm{KL}\big(Y^{\le i_j} | Y^{\le i_j-1} \;\big\|\; X^{\le i_j} | X^{\le i_j-1}\big) && \because \text{only the } i_j\text{-th bit is not conditioned} \\
&= \sum_{i\in I} \mathrm{KL}\big(Y^{\le i} | Y^{\le i-1} \;\big\|\; X^{\le i} | X^{\le i-1}\big).
\end{aligned}$$

By taking the average over $I \sim \mathcal{I}$, we obtain

$$\begin{aligned}
\mathop{\mathbb{E}}_{I\sim\mathcal{I}}[\mathrm{KL}(Y_I \parallel X_I)] &\le \mathop{\mathbb{E}}_{I\sim\mathcal{I}}\left[\sum_{i\in I} \mathrm{KL}\big(Y^{\le i} | Y^{\le i-1} \;\big\|\; X^{\le i} | X^{\le i-1}\big)\right] \\
&= \sum_{i\in[M]} \Pr_{I\sim\mathcal{I}}[i \in I] \cdot \mathrm{KL}\big(Y^{\le i} | Y^{\le i-1} \;\big\|\; X^{\le i} | X^{\le i-1}\big) \\
&\le \sum_{i\in[M]} \rho \cdot \mathrm{KL}\big(Y^{\le i} | Y^{\le i-1} \;\big\|\; X^{\le i} | X^{\le i-1}\big) \\
&= \rho \cdot \mathrm{KL}(Y \parallel X).
\end{aligned}$$

□

Proof of Theorem 4.1. Let E be the event that $f(X) \ge \mu + t$. We may assume $\Pr[E] > 0$. Let $Y = X|_E$ be X conditioned on E. Observe that

$$\begin{aligned}
\mathrm{KL}(Y \parallel X) &= \sum_{x\in E} \Pr[X = x \mid E] \cdot \log \frac{\Pr[X = x \mid E]}{\Pr[X = x]} \\
&= \sum_{x\in E} \Pr[X = x \mid E] \cdot \log \frac{\Pr[E \mid X = x]}{\Pr[E]} \\
&= \log \frac{1}{\Pr[E]} && \because \Pr[E \mid X = x] = 1 \text{ for any } x \in E.
\end{aligned}$$

Let Y' and X' be the random variables $S_I(Y_I)$ and $S_I(X_I)$ for $I \sim \mathcal{I}$, respectively. Note that

$\mu = \mathbb{E}_X[f(X)] = \mathbb{E}_X[\mathbb{E}_I[S_I(X_I)]] = \mathbb{E}_{X'}[X']$ and $\mathbb{E}_{Y'}[Y'] = \mathbb{E}_Y[f(Y)] \geq \mu + t$. Thus, we have

$$\begin{aligned}
\log \frac{1}{\Pr[E]} &= \mathrm{KL}(Y \parallel X) \\
&\geq \frac{1}{\rho} \mathop{\mathbb{E}}_{I\sim\mathcal{I}}[\mathrm{KL}(Y_I \parallel X_I)] && \text{by Lemma 4.2} \\
&\geq \frac{1}{\rho} \mathop{\mathbb{E}}_{I\sim\mathcal{I}}[\mathrm{KL}(S_I(Y_I) \parallel S_I(X_I))] && \text{data processing inequality (Lemma 3.7)} \\
&\geq \frac{1}{\rho}\mathrm{KL}\big(Y' \parallel X'\big) && \text{the convexity of KL} \\
&\geq \frac{1}{\rho}\mathrm{KL}_{\mathrm{b}}\left(\mathop{\mathbb{E}}_{Y'}\big[Y'\big] \,\middle\|\, \mathop{\mathbb{E}}_{X'}\big[X'\big]\right) && \text{Corollary 3.8} \\
&\geq \frac{1}{\rho}\mathrm{KL}_{\mathrm{b}}(\mu + t \parallel \mu).
\end{aligned}$$

This completes the proof of the upper tail. The lower tail can be proved in the same way. □

We will need an upper bound on the moment generating function, which can be proved by the transportation method. We defer the proof to Appendix A.

Theorem 4.3. *Under the same settings of Theorem 4.1, for any $\lambda \geq 0$, we have*

$$\mathbb{E}\Big[\mathrm{e}^{\lambda(f(X)-\mathbb{E}[f(X)])}\Big] \leq \exp\Big(\frac{\rho}{8}\lambda^2\Big).$$

4.2 Random Induced Subgraph of a Random Graph

By regarding an N-vertex graph as an $\binom{N}{2}$-bit string and applying Theorem 4.1, we obtain concentration inequalities regarding $G[I]$ for a random n-vertex set $I \subseteq [N]$.

Theorem 4.4. *Let S be a $[0,1]$-valued function over the set of all n-vertex graphs. For an N-vertex graph G, let $f(G) = \mathbb{E}[S(\pi(G[I]))]$, where the expectation is taken over uniformly random n-subset $I \sim \binom{[N]}{n}$ and permutation $\pi\colon [n] \to [n]$. Let $\mu = \mathbb{E}_{G\sim\mathcal{G}(N,1/2,k)}[f(G)]$. Then, for any $t \geq 0$,*

$$\begin{aligned}
\Pr_{G\sim\mathcal{G}(N,1/2,k)}[f(G) - \mu \geq t] &\leq \exp\left(-\frac{N^2}{n^2}\mathrm{KL}_{\mathrm{b}}(\mu + t \parallel \mu)\right), \\
\Pr_{G\sim\mathcal{G}(N,1/2,k)}[f(G) - \mu \leq -t] &\leq \exp\left(-\frac{N^2}{n^2}\mathrm{KL}_{\mathrm{b}}(\mu - t \parallel \mu)\right).
\end{aligned}$$

Proof. For each $C \in \binom{[N]}{k}$, let $\mathcal{G}|_C$ denote the distribution of $G_0 \cup K_C$ for $G_0 \sim \mathcal{G}(N, 1/2)$. Since $I \sim \binom{[N]}{n}$ is uniformly random and f applies a random permutation π, the expectation $\mathbb{E}_{G\sim\mathcal{G}|_C}[f(G)]$ does not depend on C. Therefore, we have $\mu = \mathbb{E}_{G\sim\mathcal{G}(N,1/2,k)}[f(G)] = \mathbb{E}_{C\sim\binom{[N]}{k}}\big[\mathbb{E}_{G\sim\mathcal{G}|_C}[f(G)]\big] = \mathbb{E}_{G\sim\mathcal{G}|_C}[f(G)]$ for any $C \in \binom{[N]}{k}$.

Note that $\mathcal{G}|_C$ is a product distribution. By regarding an N-vertex graph as an $\binom{N}{2}$-bit string and applying Theorem 4.1 for $\mathcal{I}$ being the uniform distribution over $\binom{[N]}{n}$, we have

$$\rho = \max\left\{\Pr_{I\sim\mathcal{I}}[e \subseteq I] \,\middle|\, e \in \binom{[N]}{2}\right\} = \frac{\binom{N-2}{n-2}}{\binom{N}{n}} = \frac{n(n-1)}{N(N-1)} \leq \frac{n^2}{N^2}$$

and thus

$$\Pr_{G\sim\mathcal{G}|_C}[f(G)\ge\mu+t]\le\exp\biggl(-\frac{N^2}{n^2}\mathrm{KL_b}(\mu+t\parallel\mu)\biggr),$$

from which the upper tail follows because

$$\Pr_{G\sim\mathcal{G}(N,1/2,k)}[f(G)\ge\mu+t]=\mathop{\mathbb{E}}_{C\sim\binom{[N]}{k}}\biggl[\Pr_{G\sim\mathcal{G}|_C}[f(G)\ge\mu+t]\biggr]\le\exp\biggl(-\frac{N^2}{n^2}\mathrm{KL_b}(\mu+t\parallel\mu)\biggr).$$

The lower tail can be proved in the same way. □

To analyze the embedding reduction, we will need the following variant of the concentration inequality.

Theorem 4.5. *Let S be a $[0,1]$-valued function over the set of n-vertex graphs. For an N-vertex graph $G\in\mathsf{supp}(\mathcal{G}(N,1/2,k))$, let $C_k(G)$ be the set of k-cliques in G and let $f(G)=\mathbb{E}[S(\pi(G[I]))]$, where the expectation is taken over $C\sim C_k(G)$, a uniformly random $I\in\binom{[N]}{n}$ such that $I\supseteq C$, and a uniformly random permutation π over $[n]$. Let $\mu=\mathbb{E}_{G\sim\mathcal{G}(N,1/2,k)}[f(G)]$. Then, for any $t\ge0$,*

$$\Pr_{G\sim\mathcal{G}(N,1/2,k)}[f(G)+\mu\ge t]\le\exp\biggl(-\frac{2N}{n}t^2\biggr),$$
$$\Pr_{G\sim\mathcal{G}(N,1/2,k)}[f(G)-\mu\le -t]\le\exp\biggl(-\frac{2N}{n}t^2\biggr).$$

Proof. For simplicity, write $\mathcal{G}=\mathcal{G}(N,1/2,k)$. For fixed $C\in\binom{[N]}{k}$, let $\mathcal{G}|_C$ be the distribution of $G_0\cup K_C$ for $G_0\sim\mathcal{G}(N,1/2)$. For $G\in\mathsf{supp}(\mathcal{G}|_C)$, define $f_C(G)$ by

$$f_C(G)=\mathop{\mathbb{E}}_{I\sim\binom{[N]}{n},\pi}[S(\pi(G[I]))\mid I\supseteq C]=\mathop{\mathbb{E}}_{I'\sim\binom{[N]\setminus C}{n-|C|},\pi}\bigl[S(\pi(G[I'\cup C]))\bigr].$$

Note that $f(G)=\mathbb{E}_{C\sim C_k(G)}[f_C(G)]$. The marginal distribution of $\pi(G[I])$ (over the choices of $G\sim\mathcal{G}|_C,\pi$, and I) is identical to $\mathcal{G}(n,1/2,k)$ (since the planted location is uniformly random due to the random shuffle π and edges outside the planted clique are independent). Therefore, for every $C\in\binom{[N]}{k}$, we have $\mathbb{E}_{G\sim\mathcal{G}|_C}[f_C(G)]=\mathbb{E}_{G\sim\mathcal{G}(n,1/2,k)}[S(G)]$ and thus $\mu=\mathbb{E}_{G\sim\mathcal{G}}[f(G)]=\mathbb{E}_{G\sim\mathcal{G}}\bigl[\mathbb{E}_{C\sim C_k(G)}[f_C(G)]\bigr]=\mathbb{E}_{C\sim\binom{[N]}{k}}\bigl[\mathbb{E}_{G\sim\mathcal{G}|_C}[f_C(G)]\bigr]=\mathbb{E}_{G\sim\mathcal{G}(n,1/2,k)}[S(G)]$.

A graph $G\in\mathsf{supp}(\mathcal{G}|_C)$ can be identified with a $\binom{N}{2}-\binom{k}{2}$-bit string that specifies edges outside C. With this identification, we view $\mathcal{G}|_C$ as a product distribution over $\{0,1\}^{\binom{N}{2}-\binom{k}{2}}$. Applying Theorem 4.3 for the uniform distribution $\mathcal{I}$ over $\binom{[N]\setminus C}{n-k}$, we obtain

$$\mathop{\mathbb{E}}_{G\sim\mathcal{G}_C}\Bigl[\mathrm{e}^{\lambda(f_C(G)-\mu)}\Bigr]\le\exp\biggl(\frac{n}{8N}\lambda^2\biggr)$$

for every $\lambda\ge0$. Here, note that $\rho=\frac{\binom{N-k-1}{n-k-1}}{\binom{N-k}{n-k}}=\frac{n-k}{N-k}\le\frac{n}{N}$ since any edge lying between C and

$[N] \setminus C$ appears in $G[I' \cup C]$ for $\binom{N-k-1}{n-k-1}$ times. Therefore, we have

$$\begin{aligned}
\mathop{\mathbb{E}}_{G\sim\mathcal{G}}\left[\mathrm{e}^{\lambda(f(G)-\mu)}\right] &= \mathop{\mathbb{E}}_{G\sim\mathcal{G}}\left[\mathrm{e}^{\lambda\,\mathbb{E}_{C\sim C_k(G)}[f_C(G)-\mu]}\right] \\
&\le \mathop{\mathbb{E}}_{G\sim\mathcal{G},C\sim C_k(G)}\left[\mathrm{e}^{\lambda(f_C(G)-\mu)}\right] && \text{Jensen's inequality for } x\mapsto \mathrm{e}^{\lambda x} \\
&= \mathop{\mathbb{E}}_{C\sim\binom{[N]}{k}}\left[\mathop{\mathbb{E}}_{G\sim\mathcal{G}|_C}\left[\mathrm{e}^{\lambda(f_C(G)-\mu)}\right]\right] \\
&\le \exp\Bigl(\frac{n}{8N}\lambda^2\Bigr).
\end{aligned}$$

By the standard argument (e.g., [BLM13, Section 2.3]), we obtain the upper tail:

$$\begin{aligned}
\Pr_{G\sim\mathcal{G}}[f(G)\ge\mu+t] &\le \inf_{\lambda>0}\left\{\Pr_{G\sim\mathcal{G}}\left[\mathrm{e}^{\lambda(f(G)-\mu)}\ge\mathrm{e}^{\lambda t}\right]\right\} \\
&\le \inf_{\lambda>0}\left\{\mathrm{e}^{-\lambda t}\mathop{\mathbb{E}}_{G\sim\mathcal{G}}\left[\mathrm{e}^{\lambda(f(G)-\mu)}\right]\right\} \\
&\le \inf_{\lambda>0}\left\{\exp\Bigl(\frac{n}{8N}\lambda^2-\lambda t\Bigr)\right\} \\
&\le \exp\left(-\frac{2N}{n}t^2\right),
\end{aligned}$$

where the last inequality holds by choosing $\lambda = \frac{4N}{n}t$. For the lower tail, apply the upper tail for the function $f' := 1 - f$.

□

5 Search to Decision Reductions

5.1 Auxiliary Results

We present useful structural properties of $\mathcal{G}(n, 1/2, k)$. First, we show how to find the whole clique given $G \sim \mathcal{G}(n, 1/2, k)$ and a large subset of the planted location. Such a result is already known in the literature (cf. [DGP14, Lemma 3.4]). For completeness, we present a proof.

Lemma 5.1. *Let $n \ge k \ge \ell > 0$ be any integers. With probability $1 - n\exp\left(-\frac{2\ell^2}{k}\right)$, the random graph $G \sim \mathcal{G}(n, 1/2, k)$ with planted location C satisfies the following: For any $(k/2+\ell)$-clique $C_0 \subseteq C$, any maximal clique containing C_0 is equal to C.*

Proof. Let $\mathcal{E}_{\mathrm{good}}$ be the event on $G \sim \mathcal{G}(n, 1/2, k)$ that any $v \in [n] \setminus C$ satisfies $|\Gamma(v) \cap C| < k/2 + \ell$. By the Chernoff bound (Lemma 3.3), we have $\Pr[\mathcal{E}_{\mathrm{good}}] \ge 1 - n\exp\left(-\frac{2\ell^2}{k}\right)$.

Let C_0 be any subset of C such that $|C_0| \ge k/2+\ell$. Any vertex $v \in C \setminus C_0$ satisfies $|\Gamma(v)\cap C_0| = |C_0| \ge k/2+\ell$. On the other hand, conditioned on $\mathcal{E}_{\mathrm{good}}$, any $v \notin C$ satisfies $|\Gamma(v)\cap C_0| \le |\Gamma(v)\cap C| < k/2 + \ell$. If $C_0 \cup \{v\}$ forms a clique, then v must be adjacent to all vertices in C_0 and we have $|\Gamma(v) \cap C_0| \ge k/2 + \ell$. Thus $v \in C \setminus C_0$. This proves the claim. □

In the second auxiliary result, we show that any large clique in $\mathcal{G}(n, 1/2, k)$ has a overlap of size at least $(1+\beta)\log_2 n$ with the planted location.

Lemma 5.2. *Let $0 < \beta < 1/2$ be any constant and $k \in \mathbb{N}$ be a parameter satisfying $(2+\beta)\log_2 n \le k \le n^{1/2-\beta}$. Then, with probability $1 - n^{-\Omega(\beta \log n)}$, the random graph $G \sim \mathcal{G}(n, 1/2, k)$ with planted location $C \subseteq \binom{[n]}{k}$ satisfies the following: For any $(2+\beta)\log_2 n$-clique $C' \subseteq [n]$, we have $|C \cap C'| \ge (1+\beta)\log_2 n$.*

Proof. Let C' be a $(2+\beta)\log_2 n$-clique such that $|C \cap C'| = t$. Such a clique appears in G with probability $2^{-\binom{(2+\beta)\log_2 n}{2}+\binom{t}{2}}$. Since there are $\binom{k}{t} \cdot \binom{n-k}{(2+\beta)\log_2 n - t}$ ways to choose such cliques, by the union bound over C' and t, we have

$$
\begin{aligned}
&\Pr_{G,C}\left[\exists (2+\beta)\log_2 n\text{-clique } C' \text{ such that } \left|C \cap C'\right| < (1+\beta)\log_2 n\right] \\
&\le \sum_{t=0}^{(1+\beta)\log_2 n} \binom{k}{t}\binom{n-k}{(2+\beta)\log_2 n - t} 2^{-\binom{(2+\beta)\log_2 n}{2}+\binom{t}{2}} \\
&\le n^{(2+\beta)\log_2 n} \cdot 2^{-\binom{(2+\beta)\log_2 n}{2}} \cdot \sum_{t=0}^{(1+\beta)\log_2 n} \binom{k}{t} n^{-t} 2^{\binom{t}{2}} \\
&\le n^{-\Omega(\beta \log n)} \cdot \sum_{t=0}^{k} \binom{k}{t} n^{-\frac{1-\beta}{2}t} \\
&= n^{-\Omega(\beta \log n)} \cdot \left(1 + n^{-\frac{1-\beta}{2}}\right)^k \\
&\le n^{-\Omega(\beta \log n)}.
\end{aligned}
$$

In the third inequality, note that

$$
\begin{aligned}
n^{-t}2^{\binom{t}{2}} &\le 2^{-t\log_2 n + \frac{t^2}{2}} \\
&\le 2^{-t\log_2 n + \frac{1+\beta}{2}t\log_2 n} && (\because t \le (1+\beta)\log_2 n) \\
&= n^{-\frac{1-\beta}{2}}.
\end{aligned}
$$

□

5.2 Search to Partial Recovery Reductions

First, we show that, if we can compute a list of $(1+\beta)\log_2 n$-cliques that contains a subset of the planted location, then we can find a $3k$-clique in $\mathcal{G}(3n, 1/2, 3k)$ with a slight loss in the success probability.

Lemma 5.3. *Let $\beta, \varepsilon > 0$ be any constants and $k \ge 5\log_2 n$ be a parameter. Suppose there exists a randomized polynomial-time algorithm A that, given $G \sim \mathcal{G}(n, 1/2, k)$ as input, with probability ε, outputs a list of $(1+\beta)\log_2 n$-cliques $\mathcal{F} = \{C_1, \dots, C_m\}$ such that for some $i \in [m]$, the clique C_i is a subset of the planted location of G. Then, there exists a randomized polynomial-time algorithm A' that finds a $3k$-clique in $\mathcal{G}(3n, 1/2, 3k)$ with probability $\varepsilon - o(1)$.*

Proof. The algorithm A', given $G \sim \mathcal{G}(3n, 1/2, 3k)$ as input, runs as follows:

1. Let (L, R) be a partition of $[3n]$ chosen uniformly at random from those partitions satisfying $|L| = n$ and $|R| = 2n$. Let $G_L = G[L]$ be the induced subgraph.
2. Run A on input G_L and let $\mathcal{F} = \{C_1, \dots, C_m\}$ be the output of A.

3. For each $C_i \in \mathcal{F}$, do the following:
   (a) Let $\widetilde{C}_R \subseteq R$ be the set of vertices adjacent to all vertices in C_i. That is, $\widetilde{C}_R = \{v \in R \mid \Gamma_G(v) \supseteq C_i\}$.
   (b) If $\widetilde{C}_R$ is a clique of G, let $\widetilde{C}$ be any maximal clique on G containing $\widetilde{C}_R$.
4. Repeat Steps 1–3 for n times and output the largest clique among those $\widetilde{C}$ found in Step 3(b).

We prove the correctness of A'. Let $G \sim \mathcal{G}(3n, 1/2, 3k)$ be the input with planted location C. For the random partition (L, R) of Step 1, let $C_L = C \cap L$ and $C_R = C \cap R$. Fix one iteration of Steps 1–3 and consider the following "good" events:

- Let $\mathcal{E}_1$ be the event on G, L, R that the random partition (L, R) of Step 1 satisfies $|L \cap C| = n$. Note that $\Pr[\mathcal{E}_1] = \frac{\binom{3k}{k}\binom{3n-3k}{n-k}}{\binom{3n}{n}} = \Omega(n^{-1/2})$.
- Let $\mathcal{E}_2$ be the event on G, L, R, A that the output $\mathcal{F}$ of A at Step 2 contains a $(1+\beta)\log_2 n$-clique C_i that is a subset of C_L. Conditioned on $\mathcal{E}_1$, the marginal distribution of G_L is $\mathcal{G}(n, 1/2, k)$; thus $\Pr[\mathcal{E}_2 \mid \mathcal{E}_1] \geq \varepsilon$.
- Let $\mathcal{E}_{3\mathrm{a}}$ be the event that $\widetilde{C}_R = C_R$ for some $C_i \in \mathcal{F}$ at Step 3(a). Conditioned on $\mathcal{E}_1$ and $\mathcal{E}_2$, for some $C_i \in \mathcal{F}$, we have $C_i \subseteq C_L$. Fix this C_i. Since any vertex $v \in C_R$ is adjacent to all vertices in $C_i \subseteq C_L$, we have $C_R \subseteq \widetilde{C}_R$. On the other hand, for any vertex $v \in R \setminus C_R$, edges lying between v and L appear with probability $1/2$ and are independent to A and any other edges of G conditioned on L and R. Therefore, by the union bound over v, we have $\Pr[\mathcal{E}_{3\mathrm{a}} \mid \mathcal{E}_1, \mathcal{E}_2] \geq 1 - 2n \cdot 2^{-(1+\beta)\log_2 n} \geq 1 - 2n^{-\beta}$.
- Finally, let $\mathcal{E}_{3\mathrm{b}}$ be the event on $G \sim \mathcal{G}(3n, 1/2, 3k)$ that G satisfies the property of Lemma 5.1 for $\ell = k/2$. That is, for any $2k$-clique $C' \subseteq [3n]$, any maximal clique of G containing C' is C. From Lemma 5.1, we have $\Pr[\mathcal{E}_{3\mathrm{b}}] \geq 1 - 3n\exp\left(-\frac{2(k/2)^2}{3k}\right) > 1 - 3n^{-0.1}$ since $k \geq 5\log_2 n$.

We bound the success probability of A'. Suppose the events $\mathcal{E}_{3\mathrm{a}}$ and $\mathcal{E}_{3\mathrm{b}}$ occur. Then, $\widetilde{C}_R = C_R$ is a $2k$-clique of G and any maximal clique of G containing $\widetilde{C}_R \subseteq C$ is C; thus, A' find the planted location C at Step 3. Therefore, we have

$$\Pr[A' \text{ outputs } C] \geq \Pr\left[\mathcal{E}_{3\mathrm{b}} \cap \bigcup_{i \in [n]} \{\mathcal{E}_{3\mathrm{a}} \text{ occurs at the } i\text{-th iteration of Step 4}\}\right].$$

Note that $\mathcal{E}_{3\mathrm{b}}$ is the event on $G \sim \mathcal{G}(3n, 1/2, 3k)$ and $\Pr[\mathcal{E}_{3\mathrm{b}}] \geq 1 - o(1)$. Conditioned on the event $\mathcal{E}_1$, we have

$$\begin{aligned}
\Pr[\mathcal{E}_{3\mathrm{a}} \mid \mathcal{E}_1] &\geq \frac{\Pr[\mathcal{E}_{3\mathrm{a}} \cap \mathcal{E}_1 \cap \mathcal{E}_2]}{\Pr[\mathcal{E}_1 \cap \mathcal{E}_2]} \cdot \frac{\Pr[\mathcal{E}_1 \cap \mathcal{E}_2]}{\Pr[\mathcal{E}_1]} \\
&= \Pr[\mathcal{E}_{3\mathrm{a}} \mid \mathcal{E}_1, \mathcal{E}_2] \cdot \Pr[\mathcal{E}_2 \mid \mathcal{E}_1] \\
&\geq \varepsilon - o(1).
\end{aligned}$$

Since A' repeats Steps 1–3 for n times, the event $\mathcal{E}_1$ occurs at least once during the repetition with probability $1 - (1 - O(n^{-1/2}))^n = 1 - o(1)$. If $\mathcal{E}_1$ occurs at the i-th iteration, $\mathcal{E}_{3\mathrm{a}}$ occurs at this iteration with probability $\Pr[\mathcal{E}_{3\mathrm{a}} \mid \mathcal{E}_1] \geq \varepsilon - o(1)$. Therefore, we have

$$\Pr\big[A' \text{ outputs } C\big] \geq \varepsilon - o(1).$$

□

Now we prove the search-to-partial-recovery reduction. Specifically, we show that, if we can find a $(2+\beta)\log_2 n$-clique in $\mathcal{G}(n, 1/2, k)$ with $k \geq 5\log_2 n$ (not necessarily a subset of the planted clique), then we can recover the whole planted clique in a slightly larger random graph with a slight loss in the success probability.

Lemma 5.4. *Let $\beta, \varepsilon > 0$ be constants and $k \geq 5\log_2 n$ be a parameter. If there exists a randomized polynomial-time algorithm A that finds a $(2+\beta)\log_2 n$-clique in $\mathcal{G}(n, 1/2, k)$ with probability ε, then there exists a randomized polynomial-time algorithm A' that finds a $3k$-clique in $\mathcal{G}(3n, 1/2, 3k)$ with probability $\varepsilon - o(1)$.*

Proof. Let B be the auxiliary algorithm that, given $G \sim \mathcal{G}(n, 1/2, k)$ as input, runs as follows:

1. Run A on G and let C_0 be the output of A.
2. If C_0 is a $(2+\beta)\log_2 n$-clique of G, output $\mathcal{F} = \{C' \subseteq C_0 \mid |C'| = (1+\beta)\log_2 n\}$.

Let $G \sim \mathcal{G}(n, 1/2, k)$ be the input with planted location C. We claim that, with probability $\varepsilon - o(1)$, the output $\mathcal{F}$ of B contains a $(1+\beta)\log_2 n$-clique that is a subset of C. Let $\mathcal{E}_1$ be the event that A outputs a $(2+\beta)\log_2 n$-clique at Step 1. By the assumption on A, we have $\Pr[\mathcal{E}_1] \geq \varepsilon$. Let $\mathcal{E}_2$ be the event that the input $G \sim \mathcal{G}(n, 1/2, k)$ satisfies the property of Lemma 5.2. From Lemma 5.2, we have $\Pr[\mathcal{E}_2] \geq 1 - o(1)$. If $\mathcal{E}_1 \cap \mathcal{E}_2$ occurs, the output $\mathcal{F}$ of B satisfies the desired property. This occurs with probability $\Pr[\mathcal{E}_1 \cap \mathcal{E}_2] \geq \varepsilon - o(1)$.

Since B satisfies the condition of Lemma 5.3, from Lemma 5.3, there exists a randomized polynomial-time algorithm that finds a $3k$-clique in $\mathcal{G}(3n, 1/2, 3k)$ with probability $\varepsilon - o(1)$. □

Remark 5.5. *From the proof of Lemmas 5.3 and 5.4, it is not hard to see that if the algorithm A of Lemma 5.4 is oblivious to the clique size k (i.e., the execution of A does not depend on value of k), then so does A'. Therefore, if A finds a $(2+\beta)\log_2 n$-clique in $\mathcal{G}(n, 1/2, k')$ for* all *$k' \geq 5\log_2 n$, then A' finds a $3k$-clique in $\mathcal{G}(3n, 1/2, 3k')$ for* all *$k' \geq 5\log_2 n$.*

5.3 Planted Clique of Adversarial Size

Here, we present the equivalence between the adversarial-k model and the fixed-k model. At the core of the proof, we consider the *resampling procedure* defined as follows.

Definition 5.6. *Let* Resample(G, F) *be the randomized algorithm that takes a graph $G = (V, E)$ and $F \subseteq \binom{V}{2}$ as input and flips the existence of every vertex pair $f \in F$ independently with probability* $1/2$.

We start with proving the equivalence for the decision versions.

Lemma 5.7. *If there exists a polynomial-time randomized algorithm that distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $1 - \frac{\delta}{2n}$, then, there exists a polynomial-time randomized algorithm that, for all $k' \geq k$, distinguishes $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ with advantage $1 - \delta$.*

Proof. Let A be the algorithm that distinguish $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$. Without loss of generality, we may drop the absolute value and assume $\mathbb{E}[A(\mathcal{G}(n, 1/2, k))] - \mathbb{E}[A(\mathcal{G}(n, 1/2))] \geq 1 - \frac{\delta}{2n}$, where the expectations are taken over the random graph and the internal randomness of A. Since $A(\mathcal{G}(n, 1/2, k)) \leq 1$, we have $\mathbb{E}[A(\mathcal{G}(n, 1/2))] \leq \frac{\delta}{2n}$ and similarly $\mathbb{E}[A(\mathcal{G}(n, 1/2, k))] \geq 1 - \frac{\delta}{2n}$.

Let A' be the algorithm that, given $G = ([n], E)$ as input, runs as follows:

1. For each $i = 1, \ldots, n$, do the following:

(a) Let $\pi_i \colon [n] \to [n]$ be a uniformly random permutation.

(b) Let $G_i \leftarrow \pi_i(\mathrm{Resample}(G, E([i], [n])))$.

(c) If $A(G_i) = 1$, then output 1 and terminate.

2. Output 0.

We prove the correctness of A'. If an input G is drawn from $\mathcal{G}(n, 1/2)$, then the marginal distribution of each G_i is $\mathcal{G}(n, 1/2)$. Thus, by the union bound over $i \in [n]$, we have that $\Pr[A'(\mathcal{G}(n, 1/2)) = 0] \geq 1 - \frac{\delta}{2}$.

Suppose $G \sim \mathcal{G}(n, 1/2, k')$ and let $C \sim \binom{V}{k'}$ be the k'-clique planted in G and $C_i = C \setminus [i]$ for every $i \in \{0, \cdots, n\}$. We claim that the marginal distribution of each G_i conditioned on the size $|C_i|$ is $\mathcal{G}(n, 1/2, |C_i|)$. Indeed, each C_i forms a clique in G_i and the marginal distribution of C_i is uniform due to the random permutation π_i. Moreover, each pair $\{u, v\}$ with $u \notin C_i$ forms an edge of G_i with probability $1/2$ independent of any other pairs. Therefore, the marginal distribution of each G_i is $\mathcal{G}(n, 1/2, |C_i|)$. Let $i \in \{0, \ldots, n\}$ be the first index such that $|C_i| = k$ (i is a random variable). Note that such i must exist since $|C_0| = |C| = k' \geq k$ and $|C_n| = 0$. Then, we obtain $\Pr[A'(\mathcal{G}(n, 1/2, k')) = 1] = \Pr[\bigcup_{i=0}^{n} \{A(G_i) = 1\}] \geq \Pr[A(G_i) = 1] \geq 1 - \frac{\delta}{2n}$. In the last inequality, note that the marginal distribution of G_i is $\mathcal{G}(n, 1/2, k)$.

It follows that $\mathbb{E}[A'(\mathcal{G}(n, 1/2, k))] - \mathbb{E}[A'(\mathcal{G}(n, 1/2))] \geq 1 - \delta$. □

Next, we present the case of the search versions. First, we show how to find a $(k + \ell)$-clique in $\mathcal{G}(n, 1/2, k + \ell)$ for any constant ℓ given an algorithm finding a k-clique in $\mathcal{G}(n, 1/2, k)$.

Lemma 5.8. *Let $\ell \in \mathbb{N}$, $\beta > 0$ be any constant and $k \geq (2 + \beta) \log_2 n$. Suppose there exists a randomized polynomial-time algorithm A that runs in time T and finds a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability ε. Then, there exists a randomized polynomial-time algorithm A' that, for all $k \leq k' \leq k + \ell$, runs in time $T \cdot O(n^\ell)$ and finds a k'-clique in $\mathcal{G}(n, 1/2, k')$ with probability ε.*

Proof. The algorithm A' on input $G \sim \mathcal{G}(n, 1/2, k')$ enumerates all vertex subset $S \subseteq [n]$ of $|S| \leq \ell$ and runs A on input $G_S := \mathrm{Resample}(G, E(S, [n]))$. If A outputs a clique C_S and $S \cup C_S$ forms a clique in G, let $\widetilde{C}_S = S \cup C_S$. Then, A' outputs a largest clique among all $\widetilde{C}_S$.

Let $G \sim \mathcal{G}(n, 1/2, k')$ be the input and C be the planted location. Since $G \sim \mathcal{G}(n, 1/2, k')$ and $k \leq k' \leq k + \ell$, for some $S \in \binom{[n]}{\leq \ell}$, the marginal distribution of G_S is $\mathcal{G}(n, 1/2, k)$ whose planted location is $C_S = C \setminus S$ (here, we used the condition $k \geq (2 + \beta) \log_2 n$ to ensure that G_S contains a unique k-clique that is a subset of C). For such S, A outputs $C_S = C \setminus S$ with probability ε and thus we have $C = C_S \cup S$. □

Remark 5.9. *Since A' of Lemma 5.8 is oblivious to the clique size k, if A finds planted cliques in the adversarial-k model, then so does A' (see Remark 5.5).*

We now prove the equivalence of the fixed-k and adversarial-k models for the search versions.

Lemma 5.10. *Let $k \geq 5 \log_2 n$ and $\varepsilon > 0$ be a constant. Suppose there exists a randomized polynomial-time algorithm A that finds a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability ε. Then, there exists a randomized polynomial-time algorithm A' that, for all $k' \geq 3k$, finds a k'-clique in $\mathcal{G}(3n, 1/2, k')$ with probability $\varepsilon - o(1)$.*

Proof. Let $k' \geq k$ and $\mathrm{Resample}(G, F)$ be the algorithm of Definition 5.6. Consider the following auxiliary algorithm B: On input $G \sim \mathcal{G}(n, 1/2, k')$,

1. For each $i = 1, \ldots, n$, do the following:

(a) Let π_i be a uniformly random permutation over $[n]$.

(b) Let $G_i \leftarrow \pi_i(\mathrm{Resample}(G, E([i], [n]))$.

(c) If A outputs a k-clique S of G_i on input G_i, output $\pi_i^{-1}(S)$ and terminate.

2. Output $\perp$.

We claim that, for all $k' \geq k$, B finds a k-clique in $G \sim \mathcal{G}(n, 1/2, k')$ with probability $\varepsilon - kn2^{-k/2}$. Let $G \sim \mathcal{G}(n, 1/2, k')$ be the input and C be the planted location. Let $i \in \{0, \dots, n\}$ be the first index such that $|C \setminus [i]| = k$. Then, the marginal distribution of G_i is $\mathcal{G}(n, 1/2, k)$. Thus, with probability ε, we obtain a k-clique C' in G_i at Step 1-(c). From Lemma 3.1, G_i contains a unique k-clique with probability $1 - kn2^{-k/2}$. Thus, A outputs $\pi_i(C \setminus [i])$ at Step 1-(c) with probability $\varepsilon - kn2^{-k/2}$, where π_i is the random permutation in Step 1-(a). Note that this output of B is a k-clique of $G \sim \mathcal{G}(n, 1/2, k')$. Therefore, for all $k' \geq k$, B finds a k-clique in $G \sim \mathcal{G}(n, 1/2, k')$ with probability $\varepsilon - o(1)$.

Since $k \geq 5 \log_2 n$, from Lemma 5.4, there exists a randomized polynomial-time algorithm B' that, for all $k' \geq k$, finds a $3k'$-clique in $\mathcal{G}(3n, 1/2, 3k')$ with probability $\varepsilon - o(1)$ (see also Remark 5.5). Then, from Lemma 5.8, there exists a randomized polynomial-time algorithm A' that, for all $k' \geq k$ and all $i \in \{0, 1, 2\}$, finds a $(3k'+i)$-clique in $\mathcal{G}(3n, 1/2, 3k'+i)$ with probability $\varepsilon - o(1)$. □

5.4 Search to Decision Reduction by Alon et al.

We present the search to decision reduction by Alon, Andoni, Kaufman, Matulef, Rubinfeld, and Xie [AAKMRX07]. They proved that, if one can distinguish $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ for *all* $k' \geq k/3$ with advantage $1 - 1/n$, then one can find a k-clique in $\mathcal{G}(n, 1/2, k)$. For completeness, we present a proof of the result in Appendix D.1.

Lemma 5.11 ([AAKMRX07])**.** *Let $k \geq 18 \log n$ for a sufficiently large constant $c > 0$. Suppose there exists a randomized polynomial-time algorithm A that, for all $k' \geq k/3$, distinguishes $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ with advantage $1 - \frac{\delta}{n}$. Then, there exists a randomized polynomial-time algorithm A' that, for every $k' \geq k$, finds a k'-clique in $\mathcal{G}(n, 1/2, k')$ with probability $1 - 2\delta - n\mathrm{e}^{-k/18}$.*

Using this search-to-decision reduction, we show the equivalence between decision and search versions of planted clique conjectures in a low-error regime.

Theorem 5.12. *Items 3 and 4 of Theorem 1.3 are equivalent. That is, the following are equivalent.*

- ¬3. *For some constants $\alpha > 0$ and $c > 0$, there exists a randomized polynomial-time algorithm A such that, for infinitely many n and for any $k \geq n^{1/2-\alpha}$, A distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $1 - \mathrm{e}^{-n^c}$.*
- ¬4. *For some constants $\alpha, \gamma > 0$, there exists a randomized polynomial-time algorithm B such that, for infinitely many n and any $k \geq n^{1/2-\alpha}$, B finds a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability $1 - \mathrm{e}^{-n^\gamma}$.*

Proof. The direction ¬4⇒¬3 is straightforward: Let B be the algorithm of ¬4. Let A be the following algorithm: On input G (an n-vertex graph), if B outputs a clique of size at least $n^{1/2-\alpha}$, A outputs 1. Otherwise, A outputs 0. For any $k' \geq n^{1/2-\alpha}$ and $G \sim \mathcal{G}(n, 1/2, k')$, we have $A(G) = 1$ with probability $1 - \mathrm{e}^{-n^\gamma}$. For $G \sim \mathcal{G}(n, 1/2)$, we have $A(G) = 1$ with probability at most $\Pr[\mathcal{G}(n, 1/2) \text{ contains an } n^{1/2-\alpha}\text{-clique}] \leq 2^{-\Omega(n^{1/2-\alpha})}$. Therefore, A has advantage $1 - \mathrm{e}^{-n^{\gamma'}}$, where $\gamma' < \min\{1/2 - \alpha, \gamma\}$ is any constant.

The opposite direction follows from Lemma 5.11. Since we can distinguish $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ for any $k \geq n^{1/2-\alpha}$ with advantage $1 - \mathrm{e}^{-n^{\gamma}}$, we can find a $3n^{1/2-\alpha}$-clique with probability $1 - \mathrm{e}^{-\Omega(n^{\gamma})} - \mathrm{e}^{-\Omega(n^{1/2-\alpha})}$.

□

5.5 Distinguishing k- and $(k-1)$-Clique

In this section, we show how to find the planted clique using a distinguisher for $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$.

Lemma 5.13. *Let k, δ be such that $k \geq \max\{5\sqrt{6\delta n \log n}, 5\log_2 n\}$. Suppose there exists a randomized polynomial-time algorithm A that distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$ with advantage $1 - \delta$. Then, there exists a randomized polynomial-time algorithm A' that finds a $3k$-clique in $\mathcal{G}(3n, 1/2, 3k)$ with probability $2/3 - o(1)$.*

The proof of Lemma 5.13 consists of the following two steps:

Step 1. Noisy Recovery to Decision Reduction. For a graph $G \sim \mathcal{G}(n, 1/2, k)$ with a planted location C, a vertex set $S \subseteq [n]$ is said to be *γ-noisy recovery* if $|S \cap C| \geq (1-\gamma)k$ and $|S \setminus C| \leq \gamma n$. Note that outputting a uniformly random vertex subset attains $\gamma \approx \frac{1}{2}$; thus we always assume $\gamma \leq \frac{1}{2}$. We show that, if we can distinguish $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$ with a high advantage, then we can compute a noisy recovery. The idea is based on [FK00].

Lemma 5.14. *Suppose there exists a randomized polynomial-time algorithm A that distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$ with advantage $1 - \delta$. Then, there exists a randomized polynomial-time algorithm A' that, on an input $G \sim \mathcal{G}(n, 1/2, k)$, outputs a 6δ-noisy recovery with probability $2/3$.*

Proof. Let A be the algorithm that distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$. For simplicity, suppose that $\mathbb{E}_{\mathcal{G}(n,1/2,k)}[A] - \mathbb{E}_{\mathcal{G}(n,1/2,k-1)}[A] \geq 1 - \delta$. Note that $\Pr_{A, G \sim \mathcal{G}(n,1/2,k)}[A(G) = 0] \leq \delta$ and $\Pr_{A, G \sim \mathcal{G}(n,1/2,k-1)}[A(G) = 0] \geq 1 - \delta$. Let A' be the algorithm that, on input $G \sim \mathcal{G}(n, 1/2, k)$,

1. Initialize $S \leftarrow \emptyset$.
2. For each $u \in [n]$, let G_u be the graph obtained from G by resampling edges incident to u. Specifically, for each $v \in [n] \setminus \{u\}$, flip the existence of an edge $\{u, v\}$ in G independently with probability $1/2$.
3. For each $u \in [n]$, if $A(G_u) = 0$, let $S \leftarrow S \cup \{u\}$.
4. Output S.

Let $C \subseteq [n]$ be the location of the planted clique of G. Conditioned on $u \in C$, the marginal distribution of G_u is $\mathcal{G}(n, 1/2, k-1)$ and thus $\Pr[u \in S \mid u \in C] = \Pr_{A, G \sim \mathcal{G}(n,1/2,k-1)}[A(G) = 0] \geq 1 - \delta$; in particular, $\mathbb{E}[|C \setminus S|] \leq \delta n$. By the Markov inequality, we have $\Pr[|C \cap S| \leq (1 - 6\delta)k] = \Pr[|C \setminus S| \geq 6\delta k] \leq \frac{\mathbb{E}[|C \setminus S|]}{6\delta n} \leq \frac{1}{6}$.

Similarly, conditioned on $u \notin C$, the marginal distribution of G_u is $\mathcal{G}(n, 1/2, k)$ and thus $\Pr[u \in S \mid u \notin C] = \Pr_{A, G \sim \mathcal{G}(n,1/2,k)}[A(G) = 0] \leq \delta$. By the Markov inequality, we have $\Pr[|S \setminus C| \geq 6\delta n] \leq \frac{\mathbb{E}[|S \setminus C|]}{6\delta n} \leq \frac{1}{6}$.

By the union bound over these two cases, with probability $2/3$, we have $|C \cap S| \geq (1 - 6\delta)k$ and $|S \setminus C| \leq 6\delta n$. □

Step 2. Partial Recovery to Noisy Recovery Reduction. We show that, if we can compute a noisy recovery, then we can compute a large subset of the planted location. The algorithm is similar to the recovery algorithm of Lemma 5.4: It randomly divides the vertex set $[3n]$ into two sets L and R of sizes n and $2n$, respectively, runs the noisy recovery algorithm on the induced subgraph $G[L]$, computes a $2k$-clique C_R of $G[R]$ via degree counting, and finally outputs the maximal clique containing C_R.

Lemma 5.15. *Let $k \in \mathbb{N}, \gamma \in (0, 1/2]$ be parameters satisfying $k \geq \max\{5\sqrt{\gamma n \log n}, 5\log_2 n\}$. Suppose there exists a randomized polynomial-time algorithm A that, on input size n, computes a γ-noisy recovery of $G \sim \mathcal{G}(n, 1/2, k)$ with probability $2/3$. Then, there exists a randomized polynomial-time algorithm A' that finds a $3k$-clique in $\mathcal{G}(3n, 1/2, 3k)$ with probability $2/3 - o(1)$.*

Proof. The algorithm A', given $G \sim \mathcal{G}(3n, 1/2, 3k)$ as input, runs as follows:

1. Let (L, R) be a partition of $[3n]$ chosen uniformly at random from those partitions satisfying $|L| = n$ and $|R| = 2n$. Let $G_L = G[L]$ be the induced subgraph.
2. Run A on input G_L and let $L' \subseteq L$ be the output.
3. Let $\widetilde{C}_R \subseteq R$ be the set of vertices of R having at least $\frac{|L'|}{2} + \sqrt{|L'| \log n}$ neighbors in L'.
4. If $\widetilde{C}_R \subseteq [3n]$ is a $2k$-clique of G, output any maximal clique of G containing $\widetilde{C}_R$ and terminate.
5. Repeat Step 1–4 for n times. If A' did not terminate, output $\perp$.

We prove the correctness of A'. Let $G \sim \mathcal{G}(3n, 1/2, 3k)$ be the input with planted location C. For the random partition (L, R) of Step 1, let $C_L = C \cap L$ and $C_R = C \cap R$. Fix one iteration of Steps 1–4 and consider the following "good" events:

- Let $\mathcal{E}_1$ be the event on G, L, R that the random partition (L, R) of Step 1 satisfies $|L \cap C| = n$. Note that $\Pr[\mathcal{E}_1] = \frac{\binom{3k}{k}\binom{3n-3k}{n-k}}{\binom{3n}{n}} = \Omega(n^{-1/2})$.
- Let $\mathcal{E}_2$ be the event on G, L, R, A that L' is a γ-noisy recovery of G_L at Step 2. Conditioned on $\mathcal{E}_1$, the marginal distribution of G_L is $\mathcal{G}(n, 1/2, k)$; thus $\Pr[\mathcal{E}_2 \mid \mathcal{E}_1] \geq 2/3$.
- Let $\mathcal{E}_3$ be the event on G, L, R, A that $\widetilde{C}_R = C_R$. Let $v \in C_R$. Conditioned on (L, R), the edges of G lying between L' and R occurs independently to G_L and G_R (note that the execution of A on input G_L is independent of these edges). Therefore, for any $v \in R \setminus C_R$, the Chernoff bound (Lemma 3.3) yields $\Pr\left[|\Gamma_G(v) \cap L'| \geq \frac{|L'|}{2} + \sqrt{|L'| \log n}\right] \leq n^{-2}$. On the other hand, for $v \in C_R$, with probability $1 - n^{-2}$ conditioned on $\mathcal{E}_1$ and $\mathcal{E}_2$, we have

$$\begin{aligned}
\left|\Gamma_G(v) \cap L'\right| &= \left|L' \cap C_L\right| + \left|\Gamma_G(v) \cap (L' \setminus C_L)\right| \\
&= \left|L' \cap C_L\right| + \frac{|L' \setminus C_L|}{2} - \sqrt{|L' \setminus C_L| \log n} && \text{by the Chernoff bound} \\
&= \frac{|L'|}{2} + \frac{|L' \cap C_L|}{2} - \sqrt{|L' \setminus C_L| \log n} \\
&\geq \frac{|L'|}{2} + (1-\gamma)k - \sqrt{\gamma n \log n} && \because L' \text{ is a } \gamma\text{-noisy recovery} \\
&\geq \frac{|L'|}{2} + \sqrt{(\gamma n + k) \log n} && \because \gamma \leq 1/2 \text{ and } k \geq 5\sqrt{\gamma n \log n} \\
&\geq \frac{|L'|}{2} + \sqrt{|L'| \log n} && \because |L'| \leq \gamma n + k.
\end{aligned}$$

Therefore, by the union bound over $v \in R$, we have $\Pr[\mathcal{E}_3 \mid \mathcal{E}_1, \mathcal{E}_2] \geq 1 - O(n^{-1})$.

- Let $\mathcal{E}_4$ be the event on G that $G \sim \mathcal{G}(3n, 1/2, 3k)$ satisfies the property of Lemma 5.1 for $\ell = k/2$. That is, for any $2k$-clique $C_0 \subseteq C$, we can obtain C by taking any maximal clique of G containing C_0. From Lemma 5.1, we have $\Pr[\mathcal{E}_4] \geq 1 - 3n \exp\left(-\frac{2\cdot(k/2)^2}{3k}\right) = 1 - o(1)$ (here, we used $k \geq 5 \log_2 n$).

We bound the success probability of A'. If $\mathcal{E}_3$ and $\mathcal{E}_4$ occur, then $\widetilde{C}_R \subseteq C$ of Step 3 is a $2k$-clique of G and thus, A' outputs C at Step 4. Therefore, it suffices to show that $\mathcal{E}_3 \cap \mathcal{E}_4$ occurs with probability $2/3 - o(1)$ at least once during the repetition of Step 5. Note that $\Pr[\mathcal{E}_4] = 1 - o(1)$. Fix an iteration. Conditioned on $\mathcal{E}_1$,

$$\begin{aligned}
\Pr[\mathcal{E}_3 \mid \mathcal{E}_1] &\geq \frac{\Pr[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3]}{\Pr[\mathcal{E}_1]} \\
&= \frac{\Pr[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3]}{\Pr[\mathcal{E}_1 \cap \mathcal{E}_2]} \cdot \frac{\Pr[\mathcal{E}_1 \cap \mathcal{E}_2]}{\Pr[\mathcal{E}_1]} \\
&= \Pr[\mathcal{E}_3 \mid \mathcal{E}_1, \mathcal{E}_2] \cdot \Pr[\mathcal{E}_2 \mid \mathcal{E}_1] \\
&= 1 - o(1)
\end{aligned}$$

and we obtain

$$\begin{aligned}
\Pr\left[A' \text{ outputs } C\right] &\geq \Pr\left[\mathcal{E}_4 \text{ and } \bigcup_{i\in[n]} \{\mathcal{E}_3 \text{ at } i\text{-th iteration of Step 5}\}\right] \\
&\geq \Pr\left[\bigcup_{i\in[n]} \{\mathcal{E}_3 \text{ at } i\text{-th iteration}\}\right] - \Pr[\neg\mathcal{E}_4] \\
&\geq \Pr[\mathcal{E}_3 \mid \mathcal{E}_1] \cdot \Pr\left[\bigcup_{i\in[n]} \{\mathcal{E}_1 \text{ at } i\text{-th iteration}\}\right] - o(1) \\
&\geq \frac{2}{3} - o(1).
\end{aligned}$$

In the last inequality, since $\Pr[\mathcal{E}_1] \geq \Omega(n^{-1/2})$ and the choice of (L, R) is independently random at every iteration, we have $\Pr\left[\bigcup_{i\in[n]}\{\mathcal{E}_1 \text{ at } i\text{-th iteration}\}\right] \geq 1 - (1 - O(n^{-1/2}))^n \geq 1 - o(1)$. This proves the claim. □

Proof of Lemma 5.13. Suppose A distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$ with advantage $1 - \delta$. From Lemma 5.14, there exists a randomized polynomial-time algorithm A_1 that computes a 6δ-noisy recovery of $G \sim \mathcal{G}(n, 1/2, k)$ with probability $2/3$. From Lemma 5.15, there exists a randomized polynomial-time algorithm A_2 that finds a $3k$-clique in $\mathcal{G}(3n, 1/2, 3k)$ with probability $2/3 - o(1)$. □

6 Hardness Amplification

6.1 Shrinking Reduction

We start with a formal definition of the shrinking reduction.

Definition 6.1. *Define $\mathcal{R}_{\mathrm{shr}}$ as a randomized algorithm that outputs $\pi(G[I])$ given an N-vertex graph G as input, where $I \sim \binom{[N]}{n}$ is a uniformly random n-subset and $\pi\colon [n] \to [n]$ is a uniformly random permutation.*

We prove that $\mathcal{R}_{\mathrm{shr}}$ yields samplers and then prove hardness amplification results of planted clique conjectures.

Theorem 6.2. *Let $N, n, \varepsilon, \delta$ be parameters satisfying $\delta \geq 8\exp\left(-\frac{N^2\varepsilon^2}{18n^2}\right)$. Suppose there exists a randomized polynomial-time algorithm A that distinguishes $\mathcal{G}(n, 1/2)$ and $\widetilde{\mathcal{G}}(n, 1/2, k)$ with advantage ε. Then, there exists a randomized algorithm $A'(G; n)$ that, given an N-vertex graph G as input and n as a nonuniform advice, distinguishes $\mathcal{G}(N, 1/2)$ and $\mathcal{G}(N, 1/2, k')$ with advantage $1-\delta$ in time $\mathsf{poly}(N, 1/\varepsilon, \log(1/\delta))$, where $k' = \frac{N}{n}k$.*

Remark 6.3. *The nonuniform advice n of A can be eliminated by using the fact that (the search version of) the planted clique problem admits a selector [Hir15]. Specifically, consider the following nonuniform variant of ¬1 of Theorem 1.3:*

¬1' There exist a constant $\alpha \in (0, 1/2)$ and a randomized polynomial-time algorithm A that, for infinitely many n, given some advice string $\alpha_n \in \{0,1\}^{O(\log n)}$, finds a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability $1/2$, where $k := n^{1/2-\alpha}$.

It is easy to observe the equivalence between ¬1 and ¬1'. Given a nonuniform *randomized algorithm A that takes $O(\log n)$ bits of advice, we may define a* uniform *algorithm S^A that takes a graph G of n vertices as input and, for all advice strings $\alpha \in \{0,1\}^{O(\log n)}$, computes $A(G; \alpha)$ and outputs the first k-clique of G that is found by $A'(G; \alpha)$. The oracle algorithm $S^{(\cdot)}$ is called a* selector*, and is known to characterize the property that $O(\log n)$ bits of advice can be eliminated [Hir15].*

For two computational problems L_1 and L_2 such that L_1 is reducible to L_2 and vice versa, if L_1 admits a selector, then L_2 also admits a selector [Hir15]. Thus, by our reductions that show the equivalence among many variants of planted clique problems, all the variants admit selectors, which imply that $O(\log n)$ bits of advice can be eliminated.

We first show that $\mathcal{R}_{\mathrm{shr}}$ yields samplers.

Lemma 6.4. *Let $X \sim \mathcal{G}(N, 1/2, k)$ and $Y = \mathcal{R}_{\mathrm{shr}}(X)$ be random variables. Then, for any $\varepsilon > 0$ and $\delta \geq 2\exp\left(-\frac{2N^2}{n^2}\varepsilon^2\right)$, (X, Y) is a (δ, ε)-sampler.*

Proof. Let $S\colon \mathsf{supp}(Y) \to [0,1]$ be any function. For $G \in \mathsf{supp}(X)$, let $f(G) = \mathbb{E}[S(Y) \mid X = G] = \mathbb{E}_{\mathcal{R}_{\mathrm{shr}}}[S(\mathcal{R}_{\mathrm{shr}}(G))] = \mathbb{E}_{\pi,I}[\pi(G[I])]$. From Theorem 4.4 and Fact 3.4, we have

$$\begin{aligned}\Pr_{G\sim\mathcal{G}(N,1/2,k)}[|f(G) - \mathbb{E}[S(Y) \mid X = G]| \geq \varepsilon] &= \Pr_{G\sim\mathcal{G}(N,1/2,k)}[|f(G) - \mathbb{E}[f(G)]| \geq \varepsilon] \\ &\leq 2\exp\left(-\frac{2N^2}{n^2}\varepsilon^2\right) \\ &\leq \delta\end{aligned}$$

and obtain the claim. □

Let $k' = \frac{N}{n}k$. Let $\mathcal{G}(N, 1/2, k')$ is $\mathcal{G}(n, 1/2, \mathrm{HG}(N, k', n))$ be the distribution of the n-vertex graph obtained by $\ell \sim \mathrm{HG}(N, k', n)$ and then outputting $\mathcal{G}(n, 1/2, \ell)$, where $\mathrm{HG}(N, k', n)$ denotes the hypergeometric distribution, i.e., $\Pr_{\ell\sim\mathrm{HG}(N,k',n)}[\ell = i] = \frac{\binom{k'}{i}\binom{N-k'}{n-i}}{\binom{N}{n}}$. Note that the distribution

of $Y = \mathcal{R}_{\mathrm{shr}}(X)$ for $X \sim \mathcal{G}(N, 1/2, k')$ is $\mathcal{G}(n, 1/2, \mathrm{HG}(N, k', n))$, which receives a k-clique in expectation.

To prove Theorem 6.2, we recall the following known result about the statistical distance between the distribution and the binomial distribution.

Lemma 6.5 (Theorem (4) of [DF80])**.** *For any N, k, n, $d_{\mathrm{TV}}(\mathrm{HG}(N, k, n), \mathrm{Bin}(n, k/N)) \le \frac{n}{N}$.*

Proof of Theorem 6.2. Let $X_1 \sim \mathcal{G}(N, 1/2), Y_1 = \mathcal{R}_{\mathrm{shr}}(X_1), X_2 \sim \mathcal{G}(N, 1/2, k')$, and $Y_2 = \mathcal{R}_{\mathrm{shr}}(X_2)$ be random variables. For simplicity, write $\mathcal{G}_{\mathrm{bin}} = \widetilde{\mathcal{G}}(n, 1/2, k)$ and $\mathcal{G}_{\mathrm{hg}} = \mathcal{G}(n, 1/2, \mathrm{HG}(N, k', n))$. Note that the marginal distributions of Y_1 and Y_2 are $\mathcal{G}(n, 1/2)$ and $\mathcal{G}_{\mathrm{hg}}$, respectively.

Since the only difference between $\mathcal{G}_{\mathrm{bin}}$ and $\mathcal{G}_{\mathrm{hg}}$ is the distribution of the size of the planted clique, from Lemma 6.5, we have $d_{\mathrm{TV}}(\mathcal{G}_{\mathrm{bin}}, \mathcal{G}_{\mathrm{hg}}) \le \frac{n}{N}$. Therefore, if an algorithm A distinguishes $\mathcal{G}(n, 1/2)$ and $\mathcal{G}_{\mathrm{bin}}$ with advantage ε, then A also distinguishes $\mathcal{G}(n, 1/2)$ and $\mathcal{G}_{\mathrm{hg}}$ with advantage $\varepsilon - \frac{n}{N} \ge \frac{\varepsilon}{2}$. Here, note that we may assume $\varepsilon \ge \frac{2n}{N}$; otherwise, we would have $\delta \ge 8 \exp\left(-\frac{N^2\varepsilon^2}{18n^2}\right) > 8\mathrm{e}^{-2/9} > 1$.

From Lemma 6.4, both (X_1, Y_1) and (X_2, Y_2) are $\left(\frac{\delta}{4}, \frac{\varepsilon}{6}\right)$-samplers (note that Lemma 6.4 holds even if $k = 0$). Therefore, from Lemma 3.11, we can distinguish X_1 and X_2 with advantage $1 - \delta$. □

Now we prove the equivalence results of planted clique conjectures.

Theorem 6.6. *Items 2, 3, 7 and 8 of Theorem 1.3 are equivalent. That is, the following are equivalent.*

- ¬2. *For some constant $\alpha > 0$ there exists a randomized polynomial-time algorithm A that, for infinitely many N and for any $K \ge N^{1/2-\alpha}$, distinguishes $\mathcal{G}(N, 1/2, K)$ and $\mathcal{G}(N, 1/2)$ with advantage $1/3$.*
- ¬3. *For some constants $\alpha > 0$ and $c > 0$, there exists a randomized polynomial-time algorithm A that, for infinitely many N and for any $K \ge N^{1/2-\alpha}$, distinguishes $\mathcal{G}(N, 1/2, K)$ and $\mathcal{G}(N, 1/2)$ with advantage $1 - \mathrm{e}^{-N^c}$.*
- ¬7. *For some constant $\gamma > 0$, There exists a randomized polynomial-time algorithm A that, for infinitely many n and for some $k \in \mathbb{N}$, distinguishes $\widetilde{\mathcal{G}}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\frac{k^2}{n} \cdot n^{\gamma}$.*
- ¬8. *For some constant $\gamma > 0$, there exists a randomized polynomial-time algorithm A that, for infinitely many n and for some k, distinguishes $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ with advantage $\frac{k^2}{n} \cdot n^{\gamma}$ for any $k' \ge k$.*

Proof. We prove ¬3 ⇒ ¬2 ⇒ ¬8 ⇒ ¬7 ⇒ ¬3.

Proof of ¬3⇒¬2. Assume ¬3 and let $\alpha, c > 0$ be the constants and A be the algorithm of ¬3. Then, the algorithm A also satisfies the condition of ¬2.

Proof of ¬2⇒¬8. Assume ¬2 and let $\alpha > 0$ be the constant and A be the algorithm of ¬2. Set $\gamma = \alpha$ and $k = n^{1/2-\alpha}$. For infinitely many n, A distinguishes $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ with advantage $1/3 \ge \frac{k^2}{n} \cdot n^{\gamma} = n^{-\alpha}$ for every $k' \ge k$. This proves ¬8.

Proof of ¬8⇒¬7. Assume ¬8 and let $\gamma > 0$ be the constant and A be the algorithm of ¬8. Let n, k be such that A distinguishes $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ with advantage $\frac{k^2}{n} \cdot n^\gamma$ for any $k' \geq k$. Let $k^* = \max\{k, (\log n)^2\}$. By the Chernoff bound (Lemma 3.3), $\widetilde{\mathcal{G}}(n, 1/2, 2k^*)$ contains a clique of size at least k^* with probability $1 - 2^{-\Omega(k^*)} = 1 - n^{-\omega(1)}$. Since A distinguishes $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ for any $k' \geq 2k^*$, A distinguishes $\widetilde{\mathcal{G}}(n, 1/2, 2k^*)$ and $\mathcal{G}(n, 1/2)$ with advantage $\frac{k^2}{n} \cdot n^\gamma - n^{-\omega(1)} = \frac{k^2}{n} \cdot n^{\gamma - o(1)}$.

Proof of ¬7⇒¬3. Assume ¬7 and let $\gamma > 0$ be the constant and A be the algorithm of ¬7. Let n, k be such that A distinguishes $\widetilde{\mathcal{G}}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\frac{k^2}{n} \cdot n^\gamma$. Let $\alpha, c > 0$ be constants that will be specified later. Take $K, N \in \mathbb{N}$ such that $K = \frac{N}{n}k = N^{1/2-\alpha}$. In particular, we have $N = \left(\frac{n}{k}\right)^{\frac{2}{1+2\alpha}}$ and $K = \left(\frac{n}{k}\right)^{\frac{1-2\alpha}{1+2\alpha}}$. We set α and c such that $\delta \geq 8\exp\left(-\frac{N^2\varepsilon^2}{18n^2}\right)$ holds for $\delta = \mathrm{e}^{-N^c}$ and $\varepsilon = \frac{k^2}{n} \cdot n^\gamma$. We can choose as $\alpha = \frac{\gamma}{5}$ and $c = \frac{\gamma}{6}$ since

$$\begin{aligned}
\frac{N^2\varepsilon^2}{n^2} &= N^2 \cdot \frac{k^4}{n^4} \cdot n^{2\gamma} \\
&= N^{-4\alpha} \cdot n^{2\gamma} \\
&\geq N^{\gamma - 4\alpha} && \text{since } N = \left(\frac{n}{k}\right)^{\frac{2}{1+2\alpha}} \leq n^2 \\
&\geq N^{c + \gamma/30} \geq N^{c + \Omega(1)}.
\end{aligned}$$

From Theorem 6.2 and Lemma 5.7, there exists a randomized polynomial-time algorithm A' that, on input size N, distinguishes $\mathcal{G}(N, 1/2, K')$ and $\mathcal{G}(N, 1/2)$ with advantage $1 - O(N\delta) \geq 1 - \mathrm{e}^{-N^{\Omega(1)}}$ for any $K' \geq K$. Note that we can remove the nonuniform advice n (Remark 6.3). □

Theorem 6.7. *Items 1 and 4 of Theorem 1.3 are equivalent. That is, the following are equivalent.*

- ¬1. *There exist a constant $\alpha \in (0, 1/2)$ and a randomized polynomial-time algorithm that, for infinitely many n, finds an $n^{1/2-\alpha}$-clique in $\mathcal{G}(n, 1/2, n^{1/2-\alpha})$ with probability $1/2$.*
- ¬4. *There exist constants $\alpha, \gamma > 0$ and a randomized polynomial-time algorithm that, for infinitely many $n \in \mathbb{N}$ and for any $k \geq n^{1/2-\alpha}$, finds a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability $1 - \mathrm{e}^{-n^\gamma}$.*

Proof. Note that ¬4⇒¬1 is trivial. In what follows, we prove the opposite direction.

Let $\alpha > 0$ be the constant and A be the algorithm of ¬1. From Lemma 5.10, we may assume that A finds a k'-clique in $\mathcal{G}(n, 1/2, k')$ with probability $1/2 - o(1)$ for infinitely many n and for every $k' \geq k_1 := n^{1/2-\alpha}$. Using A, we can distinguish $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ for all $k' \geq k_1$ with advantage $1/2 - o(1)$. To see this, consider the algorithm A_1 that, on an input G, outputs 1 if and only if $A(G)$ finds a clique of G of size at least k_1. Then, for any $k' \geq k_1$, we have

$$\begin{aligned}
\mathop{\mathbb{E}}_{A_1, G \sim \mathcal{G}(n,1/2,k)}\left[A_1'(G)\right] &\geq \frac{1}{2} - o(1), \\
\mathop{\mathbb{E}}_{A_1, G \sim \mathcal{G}(n,1/2)}\left[A_1'(G)\right] &\leq \Pr_{G \sim \mathcal{G}(n,1/2)}[G \text{ contains an } k_1\text{-clique}] = o(1).
\end{aligned}$$

Therefore, A_1 distinguishes $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ with advantage $1/2 - o(1)$ for all $k' \geq k_1$.

Let $k_1' = 2k_1$. We claim that A_1 indeed distinguishes $\widetilde{\mathcal{G}}(n, 1/2, k_1')$ (planted clique in the binomial-k model) and $\mathcal{G}(n, 1/2)$ with advantage $1/2 - o(1)$. Let C be the planted location of $G \sim \widetilde{\mathcal{G}}(n, 1/2, k_1')$. By the Chernoff bound (Lemma 3.3), we have $\Pr[|C| < k_1] = \Pr[\mathrm{Bin}(n, 2k_1/n) < k_1] =$

$o(1)$. Conditioned on $|C| \geq k_1$, A_1 distinguishes $\mathcal{G}(n, 1/2, |C|)$ and $\mathcal{G}(n, 1/2)$ with advantage $1/2 - o(1)$. Therefore, A_1 distinguishes $\widetilde{\mathcal{G}}(n, 1/2, k_1')$ and $\mathcal{G}(n, 1/2)$ with advantage $1/2 - o(1)$.

Let $\beta > 0$ be a parameter that will be specified later. From Theorem 6.2 (for $\delta = \mathrm{e}^{-N^\beta}$ and $\varepsilon = 1/2$) and Lemma 5.7, there exist $N = O\Big(\frac{n}{\varepsilon}\sqrt{\log(1/\delta)}\Big) = O\big(n^{1+\beta/2}\big)$ and a randomized algorithm A_2 that runs in time $\mathsf{poly}(N)$ and distinguishes $\mathcal{G}(N, 1/2, k')$ and $\mathcal{G}(N, 1/2)$ with advantage $1 - O(N\mathrm{e}^{-N^\beta})$ for all $k' \geq k_2 := \frac{N}{n}k_1'$. Note that we can remove the nonuniform advice n (Remark 6.3).

Set $\beta = \alpha$ and let $k_3 := 3k_2 = \frac{2N}{n} \cdot n^{1/2-\alpha} = O\big(n^{1/2-\alpha/2}\big)$. From Lemma 5.11, there exists a randomized polynomial-time algorithm A_3 that, for infinitely many N and every $k' \geq k_3$, finds a k'-clique in $\mathcal{G}(N, 1/2, k')$ with probability $1 - \mathrm{e}^{-\Omega(N^\beta)} - \mathrm{e}^{-\Omega(k_3)}$. This implies $\neg 4$. □

6.2 Embedding Reduction

In this part, we consider a reduction that, given a graph G, randomly embeds G into a large Erdős–Rényi graph. Formally, we consider an algorithm $\mathcal{R}_{\mathrm{emb}}$ that, given an n-vertex graph G and a parameter $N \in \mathbb{N}$, runs as follows:

1. Sample $I = \{v_1, \dots, v_n\} \sim \binom{[N]}{n}$ (suppose $v_1 < \cdots < v_n$), uniformly random permutation π over $[n]$, and $G' \sim \mathcal{G}(N, 1/2)$.
2. Replace $G'[I]$ with $\pi(G)$. Specifically, for every $1 \leq i < j \leq n$, let $G'[v_i, v_j] \leftarrow G[\pi(i), \pi(j)]$ (here, we identify a graph with its adjacency matrix).
3. Output G'.

We prove two hardness amplification results for planted clique problems using $\mathcal{R}_{\mathrm{emb}}$.

Theorem 6.8. *Let $n, N, \delta, \varepsilon$ be parameters such that $\varepsilon \geq 4\exp\Big(-\frac{N\delta^2}{8n}\Big)$. Suppose there exists a randomized polynomial-time algorithm A that finds a k-clique in $\mathcal{G}(N, 1/2, k)$ with probability ε. Then, there exists a randomized algorithm A' that, given an n-vertex graph as input and N as nonuniform advice, runs in time $\mathsf{poly}(n, N, 1/\varepsilon, \log(1/\delta))$, and satisfies*

$$\Pr_{A', G \sim \mathcal{G}(n,1/2,k)}\big[A'(G, N) \text{ is a } k\text{-clique in } G\big] \geq 1 - 2\delta.$$

Theorem 6.9. *Let $n, N, \delta, \varepsilon$ be parameters such that $\varepsilon \geq 24\exp\Big(-\frac{\delta^2\varepsilon^2 N}{18432n}\Big)$. Let $k_1, k_2 \geq 0$. Suppose there exists a randomized polynomial-time algorithm that, on input size N, distinguishes $\mathcal{G}(N, 1/2, k_1)$ and $\mathcal{G}(N, 1/2, k_2)$ with advantage ε. Then, there exists a randomized algorithm A' that is given an n-vertex graph as input and N as nonuniform advice, runs in time $\mathsf{poly}(n, N, 1/\varepsilon, \log(1/\delta))$, and distinguishes $\mathcal{G}(n, 1/2, k_1)$ and $\mathcal{G}(n, 1/2, k_2)$ with advantage $1 - \delta$.*

Remark 6.10. *We can remove the nonuniform advice N of Theorems 6.8 and 6.9. See Remark 6.3 for details.*

To prove the hardness amplification results, we first show that $\mathcal{R}_{\mathrm{emb}}$ exhibits a certain kind of sampler property. To state it more formally, we introduce the notion of *one-sided multiplicative sampler* as follows.

Definition 6.11 (Definition 3.4 of [HS23])**.** *A pair of random variables (X, Y) is a* one-sided multiplicative (δ, c)-sampler for density ε *if, for any function $S\colon \mathsf{supp}(Y) \to [0, 1]$ such that $\mathbb{E}[S(Y)] \geq \varepsilon$, we have*

$$\Pr_{x \sim X}[\mathbb{E}[S(Y) \mid X = x] \leq (1 - c)\,\mathbb{E}[S(Y)]] \leq \delta.$$

As well as Lemma 3.12, an exchange lemma for one-sided multiplicative samplers is known [HS23, Lemma 3.7]. For completeness, we prove it in Appendix C.

Lemma 6.12. *If (Y, X) is a one-sided multiplicative $\left(\frac{c\varepsilon}{2}, \frac{c}{2}\right)$-sampler for density δ, then (X, Y) is a one-sided multiplicative (δ, c)-sampler for density ε.*

The sampler property of $\mathcal{R}_{\mathrm{emb}}$ can be stated as follows.

Lemma 6.13. *Let (X, Y) be a pair of random variables obtained by $X \sim \mathcal{G}(n, 1/2, k)$ and then $Y = \mathcal{R}_{\mathrm{emb}}(X)$.*

1. *For any $\delta > 0$ and $\varepsilon \geq 4\exp\left(-\frac{N\delta^2}{8n}\right)$, (X, Y) is a one-sided multiplicative $\left(\delta, \frac{1}{2}\right)$-sampler for density ε.*
2. *For any $\delta > 0$ and $\varepsilon \geq 4\exp\left(-\frac{N\delta^2\varepsilon^2}{32n}\right)$, (X, Y) is a (δ, ε)-sampler.*

Proof. For any $[0,1]$-valued function S over $\mathsf{supp}(X)$, let $f\colon \mathsf{supp}(Y) \to [0,1]$ be the function defined by $f(y) = \mathbb{E}[S(X) \mid Y = y]$. For a graph G and $k \in \mathbb{N}$, let $C_k(G)$ be the set of k-cliques in G.

Consider the distribution of X conditioned on $Y = y$. We claim that $X|_{Y=y}$ can be obtained by choosing a k-clique $C \in C_k(y)$ and an n-vertex subset $I \subseteq [N]$ uniformly at random conditioned on $C \subseteq I$ and then outputting the induced subgraph $y[I]$. Note that $\mathcal{G}(n, 1/2, k)(G) = \frac{|C_k(G)|}{\binom{n}{k}} \cdot 2^{-\binom{n}{2}+\binom{k}{2}}$ (see, e.g., [JP00, Lemma 1]). Since $\Pr[Y = y \mid X = x] = \Pr[\mathcal{R}_{\mathrm{emb}}(x) = y] = 2^{-\binom{n}{2}+\binom{k}{2}} \Pr[\pi(y[I]) = x]$ (the probability is taken over $I \sim \binom{[N]}{n}$ and uniformly random permutation $\pi\colon [n] \to [n]$), we have

$$\begin{aligned}
\Pr[X = x \mid Y = y] &= \frac{\mathcal{G}(n, 1/2, k)(x)}{\mathcal{G}(N, 1/2, k)(y)} \cdot \Pr[Y = y \mid X = x] \\
&= \frac{|C_k(x)|}{\binom{n}{k}} \cdot \frac{\binom{N}{k}}{|C_k(y)|} \cdot \Pr_{I,\pi}[\pi(y[I]) = x] \\
&= \frac{\Pr_{I, C \sim \binom{I}{k}, \pi}[C \in C_k(y) \text{ and } \pi(y[I]) = x]}{\Pr_{C \sim \binom{[N]}{k}}[C \in C_k(y)]} \\
&= \frac{\Pr_{C \sim \binom{[N]}{k}, I, \pi}[C \in C_k(y) \text{ and } \pi(y[I]) = x \mid I \supseteq C]}{\Pr_{C \sim \binom{[N]}{k}}[C \in C_k(y)]} \\
&= \Pr_{C \sim C_k(G), I, \pi}[\pi(y[I]) = x \mid I \supseteq C].
\end{aligned}$$

Therefore, we have $f(y) = \mathbb{E}_{C \sim C_k(y), I, \pi}[S(\pi(y[I])) \mid I \supseteq C]$.

Proof of Item 1. Suppose $\mathbb{E}[S(X)] = \mathbb{E}[f(Y)] \geq \delta$. Then, from Theorem 4.5 and Fact 3.4, we obtain

$$\begin{aligned}
\Pr\left[f(Y) \leq \frac{3}{4}\mathbb{E}[f(Y)]\right] &\leq \exp\left(-\frac{2N}{n} \cdot \frac{\mathbb{E}[f(Y)]^2}{16}\right) \\
&\leq \exp\left(-\frac{N\delta^2}{8n}\right) \leq \frac{\varepsilon}{4}.
\end{aligned}$$

In other words, (Y, X) is a one-sided multiplicative $\left(\frac{\varepsilon}{4}, \frac{1}{4}\right)$-sampler for density δ. From Lemma 6.12, (X, Y) is a $\left(\delta, \frac{1}{2}\right)$-sampler for density ε.

Proof of Item 2. From Theorem 4.5 and Fact 3.4, we obtain

$$\Pr\left[|f(Y) - \mathbb{E}[f(Y)]| \geq \frac{\delta\varepsilon}{8}\right] \leq 2\exp\left(-\frac{N\delta^2\varepsilon^2}{32n}\right) \leq \frac{\varepsilon}{2}.$$

In other words, (Y, X) is a $\left(\frac{\varepsilon}{2}, \frac{\delta\varepsilon}{8}\right)$-sampler. From Lemma 3.12, (X, Y) is a (δ, ε)-sampler. □

Proof of Theorem 6.8. Let (X, Y) be the pair of random variables obtained by $X \sim \mathcal{G}(n, 1/2, k)$ and then $Y = \mathcal{R}_{\mathrm{emb}}(X)$. From Lemma 6.13 (item 1), (X, Y) is a one-sided multiplicative $\left(\delta, \frac{1}{2}\right)$-sampler for density ε. Let $S\colon \mathsf{supp}(Y) \to [0, 1]$ be the function defined by

$$S(y) = \Pr_A[A(y) \text{ is a } k\text{-clique of } y].$$

By assumption on A, we have $\mathbb{E}[S(Y)] \geq \varepsilon$. The sampler property of (X, Y) implies

$$\Pr_{x\sim\mathcal{G}(n,1/2,k)}\left[\mathbb{E}[S(Y) \mid X = x] \leq \frac{\varepsilon}{2}\right] = \Pr_{x\sim\mathcal{G}(n,1/2,k)}\left[\Pr_{A, y=\mathcal{R}_{\mathrm{emb}}(x)}[A(y) \text{ is a } k\text{-clique of } y] \leq \frac{\varepsilon}{2}\right] \leq \delta.$$

Let A' be the following algorithm: On input $x \sim \mathcal{G}(n, 1/2, k)$ and N, let $y = \mathcal{R}_{\mathrm{emb}}(x)$. If A finds a k-clique of y, recover a k-clique of x by inverting the random shuffle π and the embedding function $i \mapsto v_i$ for $I = \{v_1, \ldots, v_n\}$ (Step 1 and 2 of $\mathcal{R}_{\mathrm{emb}}$). Repeat this for $\frac{2\log(1/\delta)}{\varepsilon}$ times and if A' does not output any k-clique during this iteration, output $\perp$ and terminate.

We prove the correctness. For a $(1-\delta)$-fraction of x, $A(y)$ outputs a k-clique with probability at least $\varepsilon/2$ (randomness is over A and $\mathcal{R}_{\mathrm{emb}}$). Therefore, $\Pr_{A',x}[A'(x) \text{ is a } k\text{-clique of } x] \geq (1-\delta)^2 \geq 1 - 2\delta$. □

Proof of Theorem 6.9. Let (X_i, Y_i) be the pair of random variables obtained by $X_i \sim \mathcal{G}(n, 1/2, k_i)$ and then $Y_i = \mathcal{R}_{\mathrm{emb}}(X_i)$ $(i = 1, 2)$. From Lemma 6.13, both (X_1, Y_1) and (X_2, Y_2) are $\left(\frac{\delta}{4}, \frac{\varepsilon}{6}\right)$-samplers. Note that the marginal distributions of Y_i is $\mathcal{G}(N, 1/2, k_i)$. From Lemma 3.11, we obtain the claim. □

Theorem 6.14. *Items 1 and 6 of Theorem 1.3 are equivalent. That is, the following are equivalent.*

- ¬*1. There exist a constant $\alpha \in (0, 1/2)$ and a randomized polynomial-time algorithm that finds an $n^{1/2-\alpha}$-clique in $\mathcal{G}(n, 1/2, n^{1/2-\alpha})$ with probability $1/2$ for infinitely many n.*
- ¬*6. For some constants $\alpha \in (0, 1/2), c > 0$, there exists a randomized polynomial-time algorithm that, for infinitely many N, finds an $N^{1/2-\alpha}$-clique in $\mathcal{G}(N, 1/2, N^{1/2-\alpha})$ with probability N^{-c}.*

Proof. Note that ¬1⇒¬6 is trivial. We prove the opposite direction. Let $\alpha \in (0, 1/2), c > 0$ be the constant and A be the algorithm of ¬6. Suppose A finds an $N^{1/2-\alpha}$-clique in $\mathcal{G}(N, 1/2, N^{1/2-\alpha})$ with probability N^{-c}. Let $\varepsilon = N^{-c}$ and $\delta = 1/3$. Let $n \in \mathbb{N}$ be such that $\varepsilon \geq 4\exp\left(-\frac{N\delta^2}{8n}\right)$. We can take such n to be $n = \Omega(N/\log N)$. From Theorem 6.8 (and by removing the nonuniform advice), we can find a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability $2/3$, where $k = N^{1/2-\alpha} = n^{1/2-\alpha+o(1)}$. This proves ¬6⇒¬1. □

Theorem 6.15. *Items 1 and 10 of Theorem 1.3 are equivalent. That is, the following are equivalent.*

- ¬*1 For some constant $\alpha \in (0, 1/2)$, there exists a randomized polynomial-time algorithm A that, for infinitely many n and for $k := n^{1/2-\alpha}$, finds a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability $1/2$.*

¬10 *For some constant $\gamma > 0$, there exists a randomized polynomial-time algorithm B that, for infinitely many N, and for some $k \in \mathbb{N}$, distinguishes $\mathcal{G}(N, 1/2, k)$ and $\mathcal{G}(N, 1/2, k-1)$ with advantage $\frac{k^2}{N} \cdot N^\gamma$.*

Proof. We prove ¬1⇒¬10. Let $\alpha > 0$ be the constant and A be the algorithm of ¬1. Let $n \in \mathbb{N}$ and $k = n^{1/2-\alpha}$ be such that A finds a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability 1/2. We can distinguish $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$ by the following algorithm B: On an input G, B outputs 1 if and only if A finds a k-clique of G. To see this, note that

$$\Pr_{G \sim \mathcal{G}(n,1/2,k), B}[B(G) = 1] = \Pr_{G \sim \mathcal{G}(n,1/2,k), A}[A(G) \text{ is a } k\text{-clique of } G] \geq \frac{1}{2},$$
$$\Pr_{G \sim \mathcal{G}(n,1/2,k-1), A'}[A'(G) = 1] \leq \Pr_{G \sim \mathcal{G}(n,1/2,k-1)}[G \text{ contains a } k\text{-clique}] \leq o(1).$$

In the last inequality, we used Lemma 3.1 (if $\mathcal{G}(n, 1/2, k-1)$ contains a k-clique, then the planted $(k-1)$-clique is not unique). Therefore, B satisfies the condition of ¬10.

Now we prove ¬10⇒¬1. Let γ be the constant and B be the algorithm of ¬10. Let $N, k \in \mathbb{N}$ be such that B distinguishes $\mathcal{G}(N, 1/2, k)$ and $\mathcal{G}(N, 1/2)$ with advantage $\sqrt{\frac{k^2}{N}} \cdot N^\gamma$. Let $\varepsilon = \sqrt{\frac{k^2}{N}} \cdot N^\gamma$ and $\delta = n^{-\gamma}$. Let $\alpha = \gamma/4$, $n = \left\lfloor \sqrt{N} \right\rfloor$, and $k = n^{1/2-\alpha}$. Then, we have $\frac{\delta^2 \varepsilon^2 N}{n} = \left(\frac{N}{n^{\alpha+\gamma}}\right)^2 = N^{\frac{3}{8}\gamma}$ and thus $\varepsilon = N^{-O(1)} \geq 24 \exp\left(-\frac{\delta^2 \varepsilon^2 N}{18432 n}\right) = \mathrm{e}^{-N^{\Omega(1)}}$ and $k = n^{1/2-\alpha} \geq 5\sqrt{6\delta n \log n}$ for sufficiently large N. From Theorem 6.9, we can distinguish $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$ with advantage $1 - \delta$ in randomized polynomial time, given N as a nonuniform advice. We can eliminate the nonuniform advice N by, given n, enumerating all possible N such that $n = \left\lfloor \sqrt{N} \right\rfloor$ (there are at most $\mathsf{poly}(n)$ candidates for such N) and then approximate the advantage within an additive error $o(\delta)$. In other words, there exist a constant $\alpha > 0$ and a randomized polynomial-time algorithm B' that, for infinitely many n, B' distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2, k-1)$ with advantage $1 - \delta$. From Lemma 5.13, we can find a $3k$-clique in $\mathcal{G}(n, 1/2, 3k)$ with probability $2/3 - o(1)$; this implies ¬1. □

Theorem 6.16. *Items 7 and 9 of Theorem 1.3 are equivalent. That is, the following are equivalent:*

¬7 *For some constant $\gamma > 0$, there exists a randomized polynomial-time algorithm A that, for infinitely many n and for some k, distinguishes $\widetilde{\mathcal{G}}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\frac{k^2}{n} \cdot n^\gamma$.*

¬9 *For some constant $\gamma > 0$, there exists a randomized polynomial-time algorithm B that, for infinitely many n and for some k, distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\varepsilon(n, k)$, where*

$$\varepsilon(n, k) := \min\left\{\sqrt{\frac{k^3}{n}} \cdot n^\gamma, 1 - n^{-3}\right\}.$$

Proof. We prove ¬7⇒¬9. From Theorem 6.6, Items 3 and 7 are equivalent. Therefore, there exists a randomized polynomial-time algorithm that satisfies the condition of ¬3. This algorithm also satisfies the condition of ¬9.

We prove ¬9⇒¬7. Let $\gamma > 0$ be the constant and B be the algorithm of ¬9. Suppose that, for infinitely many n, there exists $k > n^{\frac{1}{3}(1-2\gamma)}$ such that B distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\varepsilon(n, k) = 1 - n^{-3}$. From Lemma 5.7, we can distinguish $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ with advantage $1 - O(n^{-2})$ for every $k' \geq k$. Then, from Lemma 5.11, we can find a $3k$-clique

in $\mathcal{G}(n, 1/2, 3k)$ with probability $1 - O(1/n)$. This implies ¬1, which is equivalent to ¬7 from Theorems 5.12, 6.6 and 6.7.

Therefore, we may assume that, for infinitely many n, there exists $k \le n^{\frac{1}{3}(1-2\gamma)}$ such that B distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\sqrt{\frac{k^3}{n}} \cdot n^{\gamma}$. We consider two cases: Whether Item 10 holds or not. Suppose Item 10 does not hold. Then, from Theorem 6.15, we obtain ¬1, which is equivalent to ¬7 from Theorems 5.12, 6.6 and 6.7.

Suppose that Item 10 holds. We claim that the algorithm B satisfies the condition of ¬7. Let n, k be such that $k \le n^{\frac{1}{3}(1-2\gamma)}$ and B distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\sqrt{\frac{k^3}{n}} \cdot n^{\gamma}$. For $\ell \ge 0$, let $p_\ell := \mathbb{E}[B(\mathcal{G}(n, 1/2, \ell))]$. By assumption on B, we have $|p_k - p_0| \ge \sqrt{\frac{k^3}{n}} \cdot n^{\gamma}$. Let $\sigma := \sqrt{3k \log n}$. By the Chernoff bound (Lemma 3.3), we have $\Pr[|\mathrm{Bin}(n, k/n) - k| > \sigma] \le \frac{2}{n}$. By Item 10, for any $k - \sigma \le \ell \le k + \sigma$ and any constant $\gamma' > 0$, we have

$$|p_\ell - p_k| \le \sqrt{\frac{\max\{\ell, k\}^2}{n}} \cdot n^{\gamma'} \cdot |\ell - k| \le \sqrt{\frac{3k(k+\sigma)^2 \log n}{n}} \cdot n^{\gamma'} \le \sqrt{\frac{k^3}{n}} \cdot n^{\gamma' + o(1)}.$$

Then, we have

$$\begin{aligned}
\left| \mathop{\mathbb{E}}_{\ell \sim \mathrm{Bin}(n,k/n)} [p_\ell - p_0] \right| &\ge |p_k - p_0| - \left| \mathop{\mathbb{E}}_{\ell \sim \mathrm{Bin}(n,k/n)} [p_\ell - p_k] \right| \\
&\ge \sqrt{\frac{k^3}{n}} \cdot n^{\gamma} - \sqrt{\frac{k^3}{n}} \cdot n^{\gamma' + o(1)} - \frac{2}{n} \\
&\ge (1 - o(1)) \cdot \sqrt{\frac{k^3}{n}} \cdot n^{\gamma} \qquad \text{(for any fixed } \gamma' < \gamma).
\end{aligned}$$

This shows that B satisfies the condition of ¬7.

Therefore, we obtain ¬9⇒(¬10 or (10 and ¬9))⇒ ¬7, which proves Theorem 6.16. □

We also show that the embedding reduction works well for partial recovery.

Theorem 6.17. *Items 1 and 11 of Theorem 1.3 are equivalent. That is, the following are equivalent.*

- ¬1 *There exist a constant $\alpha \in (0, 1/2)$ and a randomized polynomial-time algorithm that, for infinitely many n, finds an $n^{1/2-\alpha}$-clique in $\mathcal{G}(n, 1/2, n^{1/2-\alpha})$ with probability $1/2$.*
- ¬11 *There exist positive constants α, β and c, and a randomized polynomial-time algorithm that, for infinitely many N, finds a $(2+\beta)\log_2 N$-clique in $\mathcal{G}(N, 1/2, N^{1/2-\alpha})$ with probability N^{-c}.*

Proof. The direction ¬1 ⇒ ¬11 is straightforward. We prove the opposite direction 11 ⇒ 1 by boosting the success probability from N^{-c} to $2/3$ and then apply Lemma 5.4.

Let α, β and c be positive the constants and A be the algorithm of ¬11. Let $N \in \mathbb{N}$ be input size on which A finds a $(2+\beta)\log_2 N$-clique in $\mathcal{G}(N, 1/2, N^{1/2-\alpha})$ with probability N^{-c}. Let $\delta = 1/3, \varepsilon = N^{-c}$ and $n \in \mathbb{N}$ be such that $\varepsilon \ge 4\exp\left(-\frac{N\delta^2}{8n}\right)$ holds. We can take such n to be $n = \Omega(N/\log N)$ so that the right hand side of the inequality becomes $\exp(-n^{\Omega(1)}) \ll n^{-O(1)} = \varepsilon$.

Let B be the nonuniform algorithm that, given G (an n-vertex graph) as input and N as nonuniform advice, runs as follows:

1. Run A on input $G' = \mathcal{R}_{\mathrm{emb}}(G)$ and let $I = \{v_1, \dots, v_n\} \subseteq [N]$ and $\pi\colon [n] \to [n]$ be the subset and permutation used in $\mathcal{R}_{\mathrm{emb}}$. Let $\varphi\colon [n] \to I$ be the mapping $\varphi\colon i \mapsto v_{\pi(i)}$.

2. If A outputs a $(2+\beta)\log_2 N$-clique $C_0 \subseteq [N]$ of G', let $\mathcal{F}'$ be the list of $(1+\beta)\log_2 N$-cliques of G' that is contained in $C_0 \cap I$. That is, $\mathcal{F}' = \left\{\widetilde{C} \subseteq C_0 \cap I \;\middle|\; \left|\widetilde{C}\right| = (1+\beta)\log_2 N\right\}$. Then, output $\mathcal{F} := \left\{\phi^{-1}\left(\widetilde{C}\right) \;\middle|\; \widetilde{C} \in \mathcal{F}'\right\}$.
3. Repeat Step 1 and 2 for $\lceil \log n/\varepsilon \rceil$ times.

We say that B *succeeds* if B outputs a list of $(1+\beta)\log_2 N$-clique that contains a subset of the planted location C of G. We bound the probability that B succeeds on input $G \sim \mathcal{G}(n, 1/2, k)$. Let $G \sim \mathcal{G}(n, 1/2, k)$ be the input of B with planted location C. Fix an iteration of Step 3 and let $\mathcal{E}_1$ be the event on A, G' that A outputs a $(2+\beta)\log_2 N$-clique on input G' and $\mathcal{E}_2$ be the event on G' that G' satisfies the property of Lemma 5.2. That is, any $(2+\beta)\log_2 N$-clique C_0 of G' satisfies $|C' \cap C_0| \geq (1+\beta)\log_2 N$, where C' is the planted location of G'. Note that the planted location of G' at Step 2 is $\phi(C)$. If $\mathcal{E}_1 \cap \mathcal{E}_2$ occurs, then the $(2+\beta)\log_2 N$-clique C_0 obtained by A at Step 2 satisfies $|C_0 \cap \phi(C)| \geq (1+\beta)\log_2 N$. Then, $\mathcal{F}'$ contains a $(1+\beta)\log_2 N$-clique $\widetilde{C}$ of G' that is a subset of the planted location $\phi(C)$. Therefore, we have

$$\Pr[B \text{ succeeds}] \geq \Pr[\mathcal{E}_1 \cap \mathcal{E}_2 \text{ occurs at some iteration of Step 3}].$$

We bound the probability that $\mathcal{E}_1 \cap \mathcal{E}_2$ occurs at some iteration of Step 3. From Lemma 5.2 (note that $(2+\min\{\alpha,\beta\})\log_2 n \leq k \leq N^{1/2-\min\{\alpha,\beta\}}$) and the union bound over the repetition of Step 3, we have $\Pr[\mathcal{E}_2 \text{ occurs at every iteration of Step 3}] \geq 1 - N^{-\omega(1)}$. Let (X, Y) be the pair of random variables obtained by $X \sim \mathcal{G}(n, 1/2, k)$ and then $Y = \mathcal{R}_{\mathrm{emb}}(X)$. From Lemma 6.13, the pair (X, Y) is a one-sided multiplicative $\left(\frac{1}{3}, \frac{1}{2}\right)$-sampler for density ε. Therefore, for a $\frac{2}{3}$-fraction of $G \sim \mathcal{G}(n, 1/2, k)$, we have $\Pr[\mathcal{E}_1] \geq \varepsilon/2$. Call such G *good*. Note that $\Pr_G[G \text{ is good}] \geq \frac{2}{3}$. For any fixed good G, we have $\Pr_{A,G'}[\mathcal{E}_1 \mid G] \geq \varepsilon/2$. Because of the repetition of Step 3, we have $\Pr[\mathcal{E}_1 \text{ occurs at some iteration of Step 3} \mid G] \geq 1 - (1-\varepsilon/2)^{\log n/\varepsilon} \geq 1 - o(1)$. Therefore, we obtain

$$\begin{aligned}
&\Pr[\mathcal{E}_1 \cap \mathcal{E}_2 \text{ occurs at some iteration of Step 3}] \\
&\geq \Pr[\mathcal{E}_1 \text{ occurs at some iteration of Step 3} \mid G \text{ is good}] \Pr[G \text{ is good}] - N^{-\omega(1)} \\
&\geq \frac{2}{3} - o(1)
\end{aligned}$$

and thus B succeeds with probability $2/3 - o(1)$. Since we can eliminate the nonuniform advice of B (see Remark 6.3), from Lemma 5.3 (note that $\log_2 N = (1+o(1))\log_2 n$), we can find a k-clique in $\mathcal{G}(n, 1/2, k)$ for $k = N^{1/2-\alpha} = n^{1/2-\alpha+o(1)}$ with probability $2/3 - o(1)$. This proves ¬1. □

7 Refutation and Average-Case Polynomial Time

In this section, we prove Theorem 1.5. We first show how to boost the success probability of refutation algorithms.

Lemma 7.1. *Items 1 and 2 of Theorem 1.5 are equivalent. That is, the following are equivalent.*

1. *There exists a randomized polynomial-time algorithm that refutes a planted $n^{1/2-\alpha}$-clique with probability n^{-c} for some constants $\alpha, c > 0$.*
2. *There exists a randomized polynomial-time algorithm that refutes a planted $n^{1/2-\alpha}$-clique with probability $1 - \exp(-n^{\gamma})$ for some constants $\alpha, \gamma > 0$.*

Proof. Note that 2⇒1 is trivial. In what follows, we prove the opposite direction. The proof consists of two steps. In the first step, we boost the success probability from n^{-c} to $1/2$. In the second step, we further boost the success probability from $1/2$ to $1 - \mathrm{e}^{-n^c}$.

Step 1. Let A be the algorithm that refutes a planted $n^{1/2-\alpha}$-clique with probability n^{-c}. By repetition, we may assume without loss of generality that $\Pr_A[A(G)=1] \geq 1-2^{-n}$ for any n-vertex graph that contains a k-clique.

Let $N = N(n)$ be a function that satisfies $N^{-c} \geq 4\exp\left(-\frac{N}{32n}\right)$ for all large n. We can choose as $N = O(n(\log n)^2)$ since then the right hand side of the inequality becomes $n^{-\Omega(\log n)} \ll N^{-c}$. Let (X,Y) be the pair of random variables obtained by $X \sim \mathcal{G}(n,1/2)$ and then $Y = \mathcal{R}_{\mathrm{emb}}(X)$. From Lemma 6.13, (X,Y) is a one-sided multiplicative $\left(\frac{1}{2},\frac{1}{2}\right)$-sampler for density N^{-c}. Since $\Pr[A(Y)=0] \geq N^{-c}$, the sampler property of (X,Y) implies

$$\Pr_{x\sim\mathcal{G}(n,1/2)}\left[\Pr_{A,y=\mathcal{R}_{\mathrm{emb}}(x)}[A(y)=0] \leq \frac{N^{-c}}{2}\right] \leq \frac{1}{2}.$$

Let A_1 be the following algorithm: On an input x (an n-vertex graph),

1. Compute any N that satisfies $N^{-c} \geq 4\exp\left(-\frac{N}{32n}\right)$ (since we can take $N = \mathsf{poly}(n)$, this can be done in time $\mathsf{poly}(n)$).
2. For $i = 1,\dots,10N^c$, do the following:
   (a) Let $y = \mathcal{R}_{\mathrm{emb}}(x)$
   (b) If $A(y) = 0$, output 0.
3. Output 1.

Let $k = N^{1/2-\alpha} = n^{1/2-\alpha+o(1)}$. If x contains a k-clique, then so does y and thus we have $\Pr[A(y)=1] \geq 1-2^{-N}$. By the union bound over $i = 1,\dots,10N^c$, we have $\Pr_{A_1}[A_1(x)=1] = 1-o(1)$. On the other hand, for a 1/2-fraction of $x \sim \mathcal{G}(n,1/2)$, in Step 1(b), we have $\Pr[A(y)=0] \geq N^{-c}/2$. For such x, we have $\Pr_{A_1}[A_1(x)=1] \leq (1-N^{-c}/2)^{10N^c} \leq 1/3$. Therefore, A_1 refutes planted $n^{1/2-\alpha+o(1)}$-clique with probability 1/2 over $\mathcal{G}(n,1/2)$.

Step 2. Let $\gamma > 0$ be a parameter that will be chosen later. Let A_2 be the algorithm that, given an N-vertex graph x as input, runs as follows:

1. Set $n = N^{1-\frac{\gamma}{2}}/c'$, where c' is a sufficiently large constant.
2. For $i = 1,\dots,100$, do the following:
   (a) Let $y = \mathcal{R}_{\mathrm{shr}}(x)$.
   (b) If $A_1(y) = 0$, output 0 and terminate.
3. If A' does not terminate, output 1.

By the choice of n at Step 1, the random variables (X,Y) for $X \sim \mathcal{G}(N,1/2)$ and $Y = \mathcal{R}_{\mathrm{shr}}(X)$ is an $\left(\mathrm{e}^{-N^\gamma},1/4\right)$-sampler. Thus, for a $\left(1-\mathrm{e}^{-N^\gamma}\right)$-fraction of $x \sim \mathcal{G}(N,1/2)$, we have $\Pr_{y=\mathcal{R}_{\mathrm{shr}}(x)}[A(y)=0] \geq 1/4$. For such x, A' outputs 0 with probability $1-(3/4)^{100} \geq 2/3$.

Let $k = n^{1/2-\alpha+o(1)}$ so that A_1 refutes planted k-clique with probability 1/2 over $\mathcal{G}(n,1/2)$. Suppose x contains a k'-clique $K \subseteq [N]$, where $k' = \frac{2N}{n}k = N^{1/2-\left(\alpha-\frac{\gamma}{4}-\frac{\alpha\gamma}{2}\right)+o(1)}$. Then, each y of Step 1(a) contains a clique of size $\mathrm{HG}(N,k',n)$ (the hypergeometric distribution). From Lemma 6.5 and the Chernoff bound (Lemma 3.3), we have

$$\Pr\left[\mathrm{HG}(N,k',n) < k\right] \leq \Pr[\mathrm{Bin}(n,k'/N) < k] + \frac{n}{N} \leq \exp(-k/6) + \frac{n}{N} = o(1).$$

Conditioned on y contains a k-clique, we have $A_1(y) = 0$ with probability $o(1)$ (here, we boost the probability of A_1 from $2/3$ to $1 - o(1)$ by repetition). Therefore, for any x that contains a k'-clique, we have $\Pr[A_2(x) = 0] = 100 \cdot o(1) = o(1)$.

From Step 1 and 2 with setting $\gamma = \alpha$, there exists a randomized polynomial-time algorithm that refutes planted k''-clique with probability $1 - \mathrm{e}^{-n^\gamma}$, where $k'' = n^{1/2-\alpha'}$ and $\alpha' < \alpha/4$ is any constant. This proves ¬2. □

We now prove Theorem 1.5.

Proof of Theorem 1.5. We prove 2⇒3⇒1. Note that Items 1 and 2 are equivalent from Lemma 7.1.

Proof of 2⇒3. Let A be the algorithm of Item 2. Let $\alpha, \gamma > 0$ be the constants such that A refutes a planted $n^{1/2-\alpha}$-clique with probability $1 - \mathrm{e}^{-n^\gamma}$. Let n^d be an upper bound of the running time of A on n-vertex graphs. Let A' be the algorithm that, given an n-vertex graph G as input, runs as follows:

1. If $A(G) = 0$, output 0 and terminate.
2. Let $\gamma_0 := \gamma/2$.
3. For $i = 1, 2, \dots,$ do the following:
   (a) Let $\gamma_i := \min\{2\gamma_{i-1} - \gamma/3, 1/2 - \alpha\}$.
   (b) Check if G contains an n^{γ_i}-clique by enumerating all n^{γ_i}-subsets of $[n]$.
   (c) If not, output 0 and terminate.
   (d) If $\gamma_i = 1/2 - \alpha$, break the loop.
4. Output 1.

For simplicity, we first assume that A is a deterministic algorithm. Let $T(G)$ be the running time of A' on input G. Since $\gamma > 0$, the iteration of Step 3 ends in $O(\log(1/\gamma))$ rounds; thus $T(G) < \infty$ for any G. Moreover, A' always outputs the correct answer: If G contains a $n^{1/2-\alpha}$-clique, A' outputs 1. Otherwise, A' outputs 0.

Consider the expected running time of A'. For a $(1 - \mathrm{e}^{-n^\gamma})$-fraction of $G \sim \mathcal{G}(n, 1/2)$, the algorithm A' terminates at Step 1 and thus $T(G) = O(n^d)$. For the remaining G, A' proceeds to Step 3. The first iteration of Step 3 runs in time $\binom{n}{n^{\gamma_1}} = \exp\big(O(n^{2\gamma/3} \log n)\big)$. Suppose A' proceeds to the i-th loop of Step 3 for $i \geq 2$, which takes time $\exp(O(n^{\gamma_i} \log n))$. Then, G must contain an $n^{\gamma_{i-1}}$-clique. This occurs with probability $\binom{n}{n^{\gamma_{i-1}}} 2^{-\binom{n^{\gamma_{i-1}}}{2}} = \exp(-\Omega(n^{2\gamma_{i-1}})) = \exp\big(-\Omega(n^{\gamma_i + \gamma/3})\big)$.

Let $\varepsilon = 1/d$. Then, we have

$$\begin{aligned}
\mathop{\mathbb{E}}_{G \sim \mathcal{G}(n,1/2)}[T(G)^\varepsilon] &\leq (1 - \mathrm{e}^{-n^\gamma}) \cdot O(n^{d\varepsilon}) \\
&\quad + \mathrm{e}^{-n^\gamma} \cdot \exp\Big(\varepsilon \cdot O(n^{2\gamma/3} \log n)\Big) \\
&\quad + \sum_{i=1}^{O(\log(1/\gamma))} \exp\Big(\varepsilon \cdot O(n^{\gamma_i} \log n - \Omega(n^{\gamma_i + \gamma/3})\Big) \\
&= O(n).
\end{aligned}$$

That is, A' runs in average-case polynomial time.

If A is a randomized algorithm, by repetition, we may assume that A outputs the correct answer with probability $1 - 2^{-n^c}$ for sufficiently large constant c. Since $T(G) \leq 2^{O(n)}$ for any n-vertex graph G and for any internal randomness of A', we have $\mathbb{E}_{G,A}[T(G)^\varepsilon] \leq O(n) + 2^{-n^c} \cdot 2^{O(n)} = O(n)$.

Proof of 3⇒1. Let A be an average-polynomial-time randomized algorithm for the clique problem and $T(G)$ be the running time of A on input G. Let $\varepsilon > 0$ be the constant such that $\mathbb{E}_{G\sim\mathcal{G}(n,1/2),A}[T(G)^\varepsilon] \le O(n)$. Let $d = O(1/\varepsilon)$ be a sufficiently large constant and A' be the following algorithm: On an input G (n-vertex graph), emulate A for n^d time on input G. If A outputs $b \in \{0,1\}$, output b. Otherwise, output 1. Note that $\Pr_{G,A}\left[T(G) > n^d\right] \le n^{-d\varepsilon}\,\mathbb{E}_{G,A}[T(G)^\varepsilon] = o(1)$.

Let $k = n^{1/2-\alpha}$. If G contains a k-clique, then we have $\Pr_{A'}[A'(G) = 0] \le \Pr_A[A(G) = 0] \le 1/3$. For $G \sim \mathcal{G}(n, 1/2)$, we have

$$\Pr_{G\sim\mathcal{G}(n,1/2),A'}[A'(G) = 0] \ge 1 - \Pr_{A,G}[T(G) > n^d] - \Pr_G[G \text{ contains a } k\text{-clique}] \ge 1 - o(1).$$

Therefore, A' refutes a planted k-clique with probability $1 - o(1)$. □

8 A Worst-Case Version of the Planted Clique Problem

In this section, we present a problem whose worst-case hardness characterizes the Planted Clique conjecture. The Kolmogorov complexity $\mathrm{K}(x)$ of a string x is defined as the minimum length of a string $d \in \{0,1\}^*$ such that $U(d) = x$ for a universal Turing machine; see [LV19] for more background.

Theorem 8.1. *Item 4 of Theorem 1.3 is equivalent to Item 5 of Theorem 1.3. That is, the following are equivalent.*

- ¬4. *For some constants $\alpha > 0$ and $\gamma > 0$, there exists a randomized polynomial-time algorithm A that finds a k-clique in $\mathcal{G}(n, 1/2, k)$ with probability $1 - 2^{-n^\gamma}$ for infinitely many $n \in \mathbb{N}$ and for any $k \ge n^{1/2-\alpha}$.*
- ¬5. *For some constants $\alpha > 0$ and $\gamma > 0$, there exists a randomized polynomial-time algorithm A such that for infinitely many $n \in \mathbb{N}$ and for any $k \ge n^{1/2-\alpha}$, for any n-vertex graph G that contains a k-clique and has Kolmogorov complexity at least $\binom{n}{2} - \binom{k}{2} + \log_2 \binom{n}{k} - n^\gamma$, it holds that*

$$\Pr_A[A \text{ outputs a } k\text{-clique in } G \text{ on input } G] > \frac{1}{2}.$$

The main idea for the proof is that the set of inputs on which an algorithm errs is small. Thus, by enumerating all the inputs in the error set, one can obtain a short description of any input in the set. Taking the contrapositive, the algorithm must be successful on any incompressible input. We use only basic properties of Kolmogorov complexity, one of which is the coding theorem.

Lemma 8.2 (Coding Theorem; see, e.g., [LV19])**.** *For any computable family $\mathcal{D} = \{\mathcal{D}_n\}_{n\in\mathbb{N}}$ of distributions, for any n and any $x \in \mathsf{supp}(\mathcal{D}_n)$,*

$$\mathrm{K}(x) \le -\log_2 \mathcal{D}_n(x) + O(\log n).$$

Proof of Theorem 8.1. We first prove ¬4 ⇒ ¬5. Assume ¬4, and fix n and $k \in \mathbb{N}$ such that

$$\Pr_{A,\,G\sim\mathcal{G}(n,1/2,k)}[A(G) \text{ is not a } k\text{-clique of } G] \le 2^{-n^\gamma}.$$

Let $m := \langle n, k\rangle$, where $\langle\cdot,\cdot\rangle$ is the bijection defined as $\langle n, k\rangle := \sum_{i=0}^{n+k} i + n$. Let E_m be the set of all the graphs $G \in \mathsf{supp}(\mathcal{G}(n, 1/2, k))$ such that

$$\Pr_A[A(G) \text{ is not a } k\text{-clique of } G] \le \frac{1}{2}.$$

By Markov's inequality, we have $\Pr_{G\sim\mathcal{G}(n,1/2,k)}[G \in E_m] \le 2 \cdot 2^{-n^\gamma}$. We claim that any $G \in E_m$ satisfies $\mathrm{K}(G) \le \theta$ for some threshold θ to be chosen. The statement is obvious if $E_m = \emptyset$, and thus we assume $E_m \neq \emptyset$. Let $\mathcal{D}_m$ be the distribution of the random variable G conditioned on $G \in E_m$ for $G \sim \mathcal{G}(n, 1/2, k)$. For a fixed graph $G \in E_m$, the probability that G is sampled according to $\mathcal{D}_m$ is

$$\mathcal{D}_m(G) = \frac{\Pr[G = G']}{\Pr[G' \in E_m]} \ge 2^{-\binom{n}{2}+\binom{k}{2}} \cdot \frac{1}{\binom{n}{k}} \cdot 2^{n^\gamma - 1}.$$

Thus, by Lemma 8.2,

$$\begin{aligned}
\mathrm{K}(G) &\le -\log_2 \mathcal{D}_m(G) + O(\log m) \\
&\le \binom{n}{2} - \binom{k}{2} + \log_2 \binom{n}{k} - n^\gamma + O(\log m) \\
&\le \binom{n}{2} - \binom{k}{2} + \log_2 \binom{n}{k} - n^{\gamma/2} =: \theta,
\end{aligned}$$

where the last inequality holds for all large $n \in \mathbb{N}$. This completes the proof that any $G \in E_m$ satisfies $\mathrm{K}(G) \le \theta$. By the contrapositive of this statement, for any G such that G contains a k-clique and $\mathrm{K}(G) > \theta$, it holds that $G \notin E_m$, i.e.,

$$\Pr_A[A(G) \text{ is a } k\text{-clique of } G] > \frac{1}{2}.$$

This proves the negation of Item 5 for the constant $\gamma/2$.

Next, we prove the converse, i.e., $\neg$ 5 $\Rightarrow$ $\neg$ 4. Assume the negation of Item 5. We may assume without loss of generality that $\alpha < 1/2$. By repeating A, one can amplify the success probability $\frac{1}{2}$; thus, without loss of generality, we may assume that the success probability of A is 2^{-n^γ}. It suffices to show that with probability $1 - 2^{-n^{\gamma'}}$ over $G \sim \mathcal{G}(n, 1/2, k)$, it holds that $\mathrm{K}(G) \ge \theta$ for some constant $\gamma' > 0$. From Lemma 3.1, the probability that $\mathcal{G}(n, 1/2, k)$ does not contain a unique k-clique is at most $2kn2^{-k/2}$. Let E be the set of graphs G such that G contains a unique k-clique. The probability that $\mathrm{K}(G) < \theta$ and $G \in E$ is at most $2^\theta \cdot 2^{-\binom{n}{2}+\binom{k}{2}}/\binom{n}{k} \le 2^{-n^\gamma}$ by taking a union bound over all programs of length $< \theta$. Thus, the probability that $\mathrm{K}(G) < \theta$ is at most $2^{-n^\gamma} + 2kn2^{-k/2} \le 2^{-n^{\gamma'}}$, where $\gamma' := \min\{\gamma, 1/2 - \alpha\}/2$. □

Remark 8.3. *Since decision reduces to search, one can obtain an equivalent statement that* $\Pi \notin$ i.o.BPP *based on the promise problem* $\Pi = (\Pi_{\mathrm{YES}}, \Pi_{\mathrm{NO}})$ *defined as follows.* Π_{YES} *consists of* (G, k) *such that* G *is an* n*-vertex graph that contains a clique of size* $k \ge n^{1/2-\alpha}$ *and* $\mathrm{K}(G) \ge \binom{n}{2} - \binom{k}{2} + \log_2 \binom{n}{k} - n^\gamma$. Π_{NO} *consists of* (G, k) *such that* G *does not contain a clique of size* k.

9 Putting It All Together

In this section, we prove Theorem 1.3.

Proof of Theorem 1.3. We already proved the following equivalence results.

- From Theorem 5.12, Items 3 and 4 are equivalent.
- From Theorem 6.6, Items 2, 3, 7 and 8 are equivalent.
- From Theorem 6.7, Items 1 and 4 are equivalent.

- From Theorem 6.14, Items 1 and 6 are equivalent.
- From Theorem 8.1, Items 4 and 5 are equivalent.
- From Theorem 6.15, Items 1 and 10 are equivalent.
- From Theorem 6.16, Items 7 and 9 are equivalent.
- From Theorem 6.17, Items 1 and 11 are equivalent.

Therefore, Items 1 to 11 are equivalent. □

Acknowledgement

We thank anonymous reviewers for helpful comments. Shuichi Hirahara was supported by JST, PRESTO Grant Number JPMJPR2024, Japan. Nobutaka Shimizu is supported by JSPS KAKENHI Grant Number 23K16837, Japan.

References

[AAKMRX07] Noga Alon, Alexandr Andoni, Tali Kaufman, Kevin Matulef, Ronitt Rubinfeld, and Ning Xie. "Testing k-wise and almost k-wise independence". In: *Proceedings of the Symposium on Theory of Computing (STOC)*. 2007, pp. 496–505. DOI: 10.1145/1250790.1250863.

[ABBG11] Sanjeev Arora, Boaz Barak, Markus Brunnermeier, and Rong Ge. "Computational complexity and information asymmetry in financial products". In: *Commun. ACM* 54.5 (2011), pp. 101–107. DOI: 10.1145/1941487.1941511.

[ABIKN23] Damiano Abram, Amos Beimel, Yuval Ishai, Eyal Kushilevitz, and Varun Narayanan. "Cryptography from Planted Graphs: Security with Logarithmic-Size Messages". In: *Proceedings of the Theory of Cryptography Conference (TCC)*. 2023, pp. 286–315. DOI: 10.1007/978-3-031-48615-9_11.

[ABW10] Benny Applebaum, Boaz Barak, and Avi Wigderson. "Public-key cryptography from different assumptions". In: *Proceedings of the Symposium on Theory of Computing (STOC)*. 2010, pp. 171–180. DOI: 10.1145/1806689.1806715.

[AFMV06] Luis Antunes, Lance Fortnow, Dieter van Melkebeek, and N. V. Vinodchandran. "Computational depth: Concept and applications". In: *Theor. Comput. Sci.* 354.3 (2006), pp. 391–404. DOI: 10.1016/j.tcs.2005.11.033.

[AGGM06] Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. "On basing one-way functions on NP-hardness". In: *Proceedings of the Symposium on Theory of Computing (STOC)*. 2006, pp. 701–710. DOI: 10.1145/1132516.1132614.

[Ajt96] Miklós Ajtai. "Generating Hard Instances of Lattice Problems (Extended Abstract)". In: *Proceedings of the Symposium on the Theory of Computing (STOC)*. 1996, pp. 99–108. DOI: 10.1145/237814.237838.

[AKS98] Noga Alon, Michael Krivelevich, and Benny Sudakov. "Finding a large hidden clique in a random graph". In: *Random Struct. Algorithms* 13.3-4 (1998), pp. 457–466. DOI: 10.1002/(SICI)1098-2418(199810/12)13:3/4\%3C457::AID-RSA14\%3E3.0.CO;2-W.

[AR05] Dorit Aharonov and Oded Regev. "Lattice problems in NP ∩ coNP". In: *J. ACM* 52.5 (2005), pp. 749–765. DOI: 10.1145/1089023.1089025.

[BB15] Andrej Bogdanov and Christina Brzuska. "On Basing Size-Verifiable One-Way Functions on NP-Hardness". In: *Proceedings of the Theory of Cryptography Conference (TCC)*. 2015, pp. 1–6. DOI: 10.1007/978-3-662-46494-6_1.

[BB20] Matthew S. Brennan and Guy Bresler. "Reducibility and Statistical-Computational Gaps from Secret Leakage". In: *Proceedings of the Conference on Learning Theory (COLT)*. 2020, pp. 648–847.

[BBH18] Matthew S. Brennan, Guy Bresler, and Wasim Huleihel. "Reducibility and Computational Lower Bounds for Problems with Planted Sparse Structure". In: *Proceedings of the Conference On Learning Theory (COLT)*. 2018, pp. 48–166.

[BCKR21] Chris Brzuska, Geoffroy Couteau, Pihla Karanko, and Felix Rohrbach. "On Derandomizing Yao's Weak-to-Strong OWF Construction". In: *Proceedings of the Theory of Cryptography Conference (TCC)*. 2021, pp. 429–456. DOI: 10.1007/978-3-030-90453-1_15.

[BGP23] Guy Bresler, Chenghao Guo, and Yury Polyanskiy. "Algorithmic Decorrelation and Planted Clique in Dependent Random Graphs: The Case of Extra Triangles". In: *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*. 2023. DOI: 10.1109/FOCS57990.2023.00132.

[BHKKMP19] Boaz Barak, Samuel B. Hopkins, Jonathan A. Kelner, Pravesh K. Kothari, Ankur Moitra, and Aaron Potechin. "A Nearly Tight Sum-of-Squares Lower Bound for the Planted Clique Problem". In: *SIAM J. Comput.* 48.2 (2019), pp. 687–735. DOI: 10.1137/17M1138236.

[BJ23] Guy Bresler and Tianze Jiang. "Detection-Recovery and Detection-Refutation Gaps via Reductions from Planted Clique". In: *Proceedings of the Conference on Learning Theory (COLT)*. 2023, pp. 5850–5889.

[BKS23] Rares-Darius Buhai, Pravesh K. Kothari, and David Steurer. "Algorithms Approaching the Threshold for Semi-random Planted Clique". In: *Proceedings of the Symposium on Theory of Computing (STOC)*. 2023, pp. 1918–1926. DOI: 10.1145/3564246.3585184.

[BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities*. Oxford University Press, Feb. 2013. ISBN: 9780199535255. DOI: 10.1093/acprof:oso/9780199535255.001.0001.

[BR13a] Quentin Berthet and Philippe Rigollet. "Complexity Theoretic Lower Bounds for Sparse Principal Component Detection". In: *Proceedings of the Conference on Learning Theory (COLT)*. 2013, pp. 1046–1066.

[BR13b] Andrej Bogdanov and Alon Rosen. "Input Locality and Hardness Amplification". In: *J. Cryptology* 26.1 (2013), pp. 144–171. DOI: 10.1007/s00145-011-9117-y.

[BT06a] Andrej Bogdanov and Luca Trevisan. "Average-Case Complexity". In: *Foundations and Trends in Theoretical Computer Science* 2.1 (2006). DOI: 10.1561/0400000004.

[BT06b] Andrej Bogdanov and Luca Trevisan. "On Worst-Case to Average-Case Reductions for NP Problems". In: *SIAM J. Comput.* 36.4 (2006), pp. 1119–1159. DOI: 10.1137/S0097539705446974.

[CI99] Giovanni Di Crescenzo and Russell Impagliazzo. "Security-Preserving Hardness-Amplification for Any Regular One-Way Function". In: *Proceedings of the Symposium on Theory of Computing (STOC)*. 1999, pp. 169–178. DOI: 10.1145/301250.301296.

[CMZ23] Zongchen Chen, Elchanan Mossel, and Ilias Zadik. "Almost-Linear Planted Cliques Elude the Metropolis Process". In: *Proceedings of the Symposium on Discrete Algorithms (SODA)*. 2023, pp. 4504–4539. DOI: 10.1137/1.9781611977554.CH171.

[CT06] Thomas M. Cover and Joy A. Thomas. *Elements of information theory (2. ed.)* Wiley, 2006. ISBN: 978-0-471-24195-9.

[DF80] P Diaconis and D Freedman. "Finite Exchangeable Sequences". en. In: *aop* 8.4 (Aug. 1980), pp. 745–764. ISSN: 0091-1798, 2168-894X. DOI: 10.1214/aop/1176994663.

[DGP14] Yael Dekel, Ori Gurel-Gurevich, and Yuval Peres. "Finding Hidden Cliques in Linear Time with High Probability". In: *Comb. Probab. Comput.* 23.1 (2014), pp. 29–49. DOI: 10.1017/S096354831300045X.

[DM15] Yash Deshpande and Andrea Montanari. "Finding Hidden Cliques of Size $\sqrt{N/e}$ in Nearly Linear Time". In: *Found. Comput. Math.* 15.4 (2015), pp. 1069–1128. DOI: 10.1007/s10208-014-9215-y.

[ERSY22] Reyad Abed Elrazik, Robert Robere, Assaf Schuster, and Gal Yehuda. "Pseudorandom Self-Reductions for NP-Complete Problems". In: *Proceedings of the Innovations in Theoretical Computer Science Conference (ITCS)*. 2022, 65:1–65:12. DOI: 10.4230/LIPIcs.ITCS.2022.65.

[FF93] Joan Feigenbaum and Lance Fortnow. "Random-Self-Reducibility of Complete Sets". In: *SIAM J. Comput.* 22.5 (1993), pp. 994–1005. DOI: 10.1137/0222061.

[FGRVX17] Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh S. Vempala, and Ying Xiao. "Statistical Algorithms and a Lower Bound for Detecting Planted Cliques". In: *J. ACM* 64.2 (2017), 8:1–8:37. DOI: 10.1145/3046674.

[FK00] Uriel Feige and Robert Krauthgamer. "Finding and certifying a large hidden clique in a semirandom graph". In: *Random Struct. Algorithms* 16.2 (2000), pp. 195–208. DOI: 10.1002/(SICI)1098-2418(200003)16:2\%3C195::AID-RSA5\%3E3.0.CO;2-A.

[FK03] Uriel Feige and Robert Krauthgamer. "The Probable Value of the Lovász–Schrijver Relaxations for Maximum Independent Set". In: *SIAM J. Comput.* 32.2 (2003), pp. 345–370. DOI: 10.1137/S009753970240118X.

[FR10] Uriel Feige and Dorit Ron. "Finding hidden cliques in linear time". In: *Discrete Mathematics & Theoretical Computer Science* Proceedings (2010). DOI: 10.46298/dmtcs.2802.

[GILVZ90] Oded Goldreich, Russell Impagliazzo, Leonid A. Levin, Ramarathnam Venkatesan, and David Zuckerman. "Security Preserving Amplification of Hardness". In: *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*. 1990, pp. 318–326. DOI: 10.1109/FSCS.1990.89550.

[GLSS15] Dmitry Gavinsky, Shachar Lovett, Michael Saks, and Srikanth Srinivasan. "A tail bound for read-k families of functions". en. In: *Random Structures & Algorithms* 47.1 (Aug. 2015), pp. 99–108. ISSN: 1042-9832, 1098-2418. DOI: 10.1002/rsa.20532.

[Gol11] Oded Goldreich. "On Security Preserving Reductions - Revised Terminology". In: *Studies in Complexity and Cryptography*. 2011, pp. 540–546. DOI: 10.1007/978-3-642-22670-0_34.

[Hås99] Johan Håstad. "Clique is hard to approximate within $n^{1-\varepsilon}$". In: *Acta Mathematica* 182.1 (Mar. 1999), pp. 105–142. ISSN: 1871-2509. DOI: 10.1007/BF02392825.

[HHR11] Iftach Haitner, Danny Harnik, and Omer Reingold. "On the Power of the Randomized Iterate". In: *SIAM J. Comput.* 40.6 (2011), pp. 1486–1528. DOI: 10.1137/080721820.

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. "A Pseudorandom Generator from any One-way Function". In: *SIAM J. Comput.* 28.4 (1999), pp. 1364–1396. DOI: 10.1137/S0097539793244708.

[Hir15] Shuichi Hirahara. "Identifying an Honest $\mathsf{EXP}^{\mathsf{NP}}$ Oracle Among Many". In: *Proceedings of the Conference on Computational Complexity (CCC)*. 2015, pp. 244–263. DOI: 10.4230/LIPIcs.CCC.2015.244.

[Hir18] Shuichi Hirahara. "Non-Black-Box Worst-Case to Average-Case Reductions within NP". In: *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*. 2018, pp. 247–258. DOI: 10.1109/FOCS.2018.00032.

[Hir22] Shuichi Hirahara. "Meta-Computational Average-Case Complexity: A New Paradigm Toward Excluding Heuristica". In: *Bull. EATCS* 136 (2022).

[HK11] Elad Hazan and Robert Krauthgamer. "How Hard Is It to Approximate the Best Nash Equilibrium?" In: *SIAM J. Comput.* 40.1 (2011), pp. 79–91. DOI: 10.1137/090766991.

[HN23] Shuichi Hirahara and Mikito Nanashima. "Learning in Pessiland via Inductive Inference". In: *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*. 2023. DOI: 10.1109/FOCS57990.2023.00033.

[HS23] Shuichi Hirahara and Nobutaka Shimizu. "Hardness Self-Amplification: Simplified, Optimized, and Unified". In: *Symposium on Theory of Computing (STOC)*. STOC 2023. Orlando, FL, USA: Association for Computing Machinery, June 2023, pp. 70–83. ISBN: 9781450399135. DOI: 10.1145/3564246.3585189.

[HWX15] Bruce E. Hajek, Yihong Wu, and Jiaming Xu. "Computational Lower Bounds for Community Detection on Random Graphs". In: *Proceedings of the Conference on Learning Theory (COLT)*. 2015, pp. 899–928.

[IJK09] Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. "Chernoff-Type Direct Product Theorems". In: *J. Cryptol.* 22.1 (2009), pp. 75–92. DOI: 10.1007/s00145-008-9029-7.

[IJKW10] Russell Impagliazzo, Ragesh Jaiswal, Valentine Kabanets, and Avi Wigderson. "Uniform Direct Product Theorems: Simplified, Optimized, and Derandomized". In: *SIAM J. Comput.* 39.4 (2010), pp. 1637–1665. DOI: 10.1137/080734030.

[Jan04] Svante Janson. "Large deviations for sums of partly dependent random variables". In: *Random Structures & Algorithms* 24.3 (2004), pp. 234–248.

[Jer92] Mark Jerrum. "Large Cliques Elude the Metropolis Process". In: *Random Struct. Algorithms* 3.4 (1992), pp. 347–360. DOI: 10.1002/rsa.3240030402.

[JP00] Ari Juels and Marcus Peinado. "Hiding Cliques for Cryptographic Security". In: *Des. Codes Cryptography* 20.3 (2000), pp. 269–280. DOI: 10.1023/A:1008374125234.

[Kuč95] Luděk Kučera. "Expected Complexity of Graph Partitioning Problems". In: *Discrete Applied Mathematics* 57.2-3 (1995), pp. 193–212. DOI: 10.1016/0166-218X(94)00103-K.

[KV02] Michael Krivelevich and Van H. Vu. "Approximating the Independence Number and the Chromatic Number in Expected Polynomial Time". In: *J. Comb. Optim.* 6.2 (2002), pp. 143–155. DOI: 10.1023/A:1013899527204.

[KVWX23] Pravesh Kothari, Santosh S. Vempala, Alexander S. Wein, and Jeff Xu. "Is Planted Coloring Easier than Planted Clique?" In: *Proceedings of the Conference on Learning Theory (COLT)*. 2023, pp. 5343–5372.

[KZ14] Pascal Koiran and Anastasios Zouzias. "Hidden Cliques and the Certification of the Restricted Isometry Property". In: *IEEE Trans. Inf. Theory* 60.8 (2014), pp. 4999–5006. DOI: 10.1109/TIT.2014.2331341.

[Lev86] Leonid A. Levin. "Average Case Complete Problems". In: *SIAM J. Comput.* 15.1 (1986), pp. 285–286. DOI: 10.1137/0215020.

[LP23] Yanyi Liu and Rafael Pass. *On One-way Functions and the Worst-case Hardness of Time-Bounded Kolmogorov Complexity*. Cryptology ePrint Archive, Paper 2023/1086. https://eprint.iacr.org/2023/1086. 2023. URL: https://eprint.iacr.org/2023/1086.

[LTW05] Henry C. Lin, Luca Trevisan, and Hoeteck Wee. "On Hardness Amplification of One-Way Functions". In: *Proceedings of the Theory of Cryptography Conference (TCC)*. 2005, pp. 34–49. DOI: 10.1007/978-3-540-30576-7_3.

[LV19] Ming Li and Paul M. B. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications, 4th Edition*. Texts in Computer Science. Springer, 2019. ISBN: 978-3-030-11297-4. DOI: 10.1007/978-3-030-11298-1.

[Mar21] Jay Mardia. "Is the Space Complexity of Planted Clique Recovery the Same as That of Detection?" In: *Proceedings of the Innovations in Theoretical Computer Science Conference (ITCS)*. 2021, 34:1–34:17. DOI: 10.4230/LIPIcs.ITCS.2021.34.

[McD89] Colin McDiarmid. "On the method of bounded differences". In: *Surveys in combinatorics* 141.1 (1989), pp. 148–188. DOI: 10.1017/CBO9781107359949.008.

[MRS21] Pasin Manurangsi, Aviad Rubinstein, and Tselil Schramm. "The Strongish Planted Clique Hypothesis and Its Consequences". In: *Proceedings of the Innovations in Theoretical Computer Science Conference (ITCS)*. 2021, 10:1–10:21. DOI: 10.4230/LIPIcs.ITCS.2021.10.

[MW15] Zongming Ma and Yihong Wu. "Computational barriers in minimax submatrix detection". In: *The Annals of Statistics* 43.3 (2015), pp. 1089–1116. DOI: 10.1214/14-AOS1300. URL: https://doi.org/10.1214/14-AOS1300.

[Pan21] Shuo Pang. "SOS Lower Bound for Exact Planted Clique". In: *Proceedings of the Computational Complexity Conference (CCC)*. 2021, 26:1–26:63. DOI: 10.4230/LIPICS.CCC.2021.26.

[Ros08] Benjamin Rossman. "On the constant-depth complexity of k-clique". In: *Proceedings of the Symposium on Theory of Computing (STOC)*. 2008, pp. 721–730. DOI: 10.1145/1374376.1374480.

[San20] Rahul Santhanam. "Pseudorandomness and the Minimum Circuit Size Problem". In: *Proceedings of the Innovations in Theoretical Computer Science Conference (ITCS)*. 2020, 68:1–68:26. DOI: 10.4230/LIPIcs.ITCS.2020.68.

[SW22] Tselil Schramm and Alexander S. Wein. "Computational barriers to estimation from low-degree polynomials". In: *The Annals of Statistics* 50.3 (June 2022). ISSN: 0090-5364. DOI: 10.1214/22-aos2179.

[Tre18] Luca Trevisan. *Search vs Decision vs Certification for Planted Problems*. https://lucatrevisan.wordpress.com/2018/05/06/search-vs-decision-vs-certification-for-planted-problems/. [Online; accessed 2023]. 2018.

[Yao82] Andrew Chi-Chih Yao. "Theory and Applications of Trapdoor Functions (Extended Abstract)". In: *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*. 1982, pp. 80–91. DOI: 10.1109/SFCS.1982.45.

[Zuc07] David Zuckerman. "Linear Degree Extractors and the Inapproximability of Max Clique and Chromatic Number". In: *Theory of Computing* 3.1 (2007), pp. 103–128. DOI: 10.4086/toc.2007.v003a006.

A Concentration from the Transportation Method

We prove a variant of Theorem 4.1 using the transportation method.

Theorem 4.3. *Under the same settings of Theorem 4.1, for any $\lambda \geq 0$, we have*

$$\mathbb{E}\left[e^{\lambda(f(X)-\mathbb{E}[f(X)])}\right] \leq \exp\left(\frac{\rho}{8}\lambda^2\right).$$

The proof of Theorem 4.3 is based on the following transportation method. For completeness, we include a proof here.

Lemma A.1 (Special Case of Lemma 4.18 of [BLM13])**.** *Let X be a random variable that takes values in a finite set Ω and $v > 0$ be a parameter. Let $f\colon \Omega \to \mathbb{R}$ be a function. Suppose that for any random variable Y whose support is contained in the support of X,*

$$\mathbb{E}[f(Y)] - \mathbb{E}[f(X)] \leq \sqrt{2v \cdot \mathrm{KL}(Y \parallel X)}, \tag{1}$$

where $\mathrm{KL}(Y \parallel X)$ denotes the Kullback–Leibler divergence between Y and X. Then, for every $\lambda > 0$, we have

$$\mathbb{E}\left[e^{\lambda(f(X)-\mathbb{E}[f(X)])}\right] \leq \exp\left(\frac{v\lambda^2}{2}\right).$$

Proof. Fix any $\lambda > 0$ and define a random variable U by

$$U = \lambda(f(X) - \mathbb{E}[f(X)]) - \max_Y\{\lambda(\mathbb{E}[f(Y)] - \mathbb{E}[f(X)]) - \mathrm{KL}(Y \parallel X)\},$$

where the maximum is taken over all random variables Y such that $\mathsf{supp}(Y) \subseteq \mathsf{supp}(X)$. Note that the set of distributions over a finite set is compact. Let μ be the distribution of X. The proof consists of three steps.

1. For any nonnegative random variable Z with $\mathbb{E}_\mu[Z] = 1$, $\mathbb{E}_\mu[UZ] \le \mathbb{E}_\mu[Z \log Z]$.
2. $\log \mathbb{E}_\mu[\mathrm{e}^U] \le 0$.
3. $\log \mathbb{E}\left[\mathrm{e}^{\lambda(f(X)-\mathbb{E}[f(X)])}\right] \le \frac{v\lambda^2}{2}$.

Step 1. Let Z be any nonnegative random variable such that $\mathbb{E}_\mu[Z] = 1$ and Y' be the random variable such that $\Pr[Y' = x] = \mu(x) \cdot Z(x)$. Note that $\sum_x \Pr[Y' = x] = \mathbb{E}_\mu[Z] = 1$ and $Y' \ll X$. Then, we have

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mu}[UZ] &\le \mathop{\mathbb{E}}_{\mu}\left[\left(\lambda(f(X) - \mathbb{E}[f(X)]) - \lambda(\mathbb{E}[f(Y')] - \mathbb{E}[f(X)]) + \mathrm{KL}\left(Y' \,\middle\|\, X\right)\right) \cdot Z\right] \\
&= \lambda\left(\mathop{\mathbb{E}}_{\mu}[f(X)Z] - \mathbb{E}[f(Y')]\mathop{\mathbb{E}}_{\mu}[Z]\right) + \mathrm{KL}\left(Y' \,\middle\|\, X\right) \cdot \mathop{\mathbb{E}}_{\mu}[Z] \\
&= \mathrm{KL}\left(Y' \,\middle\|\, X\right) \\
&= \mathop{\mathbb{E}}_{\mu}[Z \log Z].
\end{aligned}$$

Step 2. By substituting $Z = \mathrm{e}^U / \mathbb{E}_\mu[\mathrm{e}^U]$ to $\mathbb{E}_\mu[UZ] \le \mathbb{E}_\mu[Z \log Z]$, we obtain

$$\begin{aligned}
\frac{1}{\mathbb{E}[\mathrm{e}^U]}\,\mathbb{E}[U\mathrm{e}^U] &\le \frac{1}{\mathbb{E}[\mathrm{e}^U]}\,\mathbb{E}[\mathrm{e}^U \log(\mathrm{e}^U / \mathbb{E}[\mathrm{e}^U])] \\
&= \frac{\mathbb{E}[U\mathrm{e}^U] - \mathbb{E}[\mathrm{e}^U] \log \mathbb{E}[\mathrm{e}^U]}{\mathbb{E}[\mathrm{e}^U]}.
\end{aligned}$$

Since $\mathbb{E}\left[\mathrm{e}^U\right] > 0$, we have $\log \mathbb{E}\left[\mathrm{e}^U\right] \le 0$.

Step 3. Since $\log \mathbb{E}\left[\mathrm{e}^U\right] \le 0$, we have

$$\begin{aligned}
\log \mathbb{E}\left[\mathrm{e}^{\lambda(f(X)-\mathbb{E}[f(X)])}\right] &\le \max_{Y \ll X} \{\lambda(\mathbb{E}[f(Y)] - \mathbb{E}[f(X)]) - \mathrm{KL}(Y \parallel X)\} \\
&\le \lambda\sqrt{2v\mathrm{KL}(Y \parallel X)} - \mathrm{KL}(Y \parallel X) \\
&= -\left(\sqrt{\mathrm{KL}(Y \parallel X)} - \lambda\sqrt{\frac{v}{2}}\right)^2 + \frac{v}{2}\lambda^2 \\
&\le \frac{v}{2}\lambda^2.
\end{aligned}$$

□

Proof of Theorem 4.3. Let Y be a random variable such that $\mathsf{supp}(Y) \subseteq \mathsf{supp}(X)$. Denote by $d_{\mathrm{TV}}(\cdot, \cdot)$ the total variation distance. Then, we have

$$\begin{aligned}
\mathbb{E}[f(Y)] - \mathbb{E}[f(X)] &= \mathop{\mathbb{E}}_{I \sim \mathcal{I}}[\mathbb{E}[S_I(Y_I) - S_I(X_I)]] \\
&\le \mathop{\mathbb{E}}_{I \sim \mathcal{I}}[d_{\mathrm{TV}}(Y_I, X_I)] && S_I \text{ is } [0,1]\text{-valued} \\
&\le \mathop{\mathbb{E}}_{I \sim \mathcal{I}}\left[\sqrt{\frac{\mathrm{KL}(Y_I \parallel X_I)}{2}}\right] && \text{Pinsker's inequality} \\
&\le \sqrt{\frac{\mathbb{E}_I[\mathrm{KL}(Y_I \parallel X_I)]}{2}} && \text{concavity of } \sqrt{\cdot} \\
&\le \sqrt{\frac{\rho}{2}\mathrm{KL}(Y \parallel X)},
\end{aligned}$$

where the last inequality follows from Lemma 4.2. Applying Lemma A.1 for $v = \frac{\rho}{4}$, we obtain the result. □

B A Decision Algorithm by Edge Counting

We show how to distinguish $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\Omega(k^2/n)$.

Proposition B.1. *Let A be the algorithm that, given a graph G as input, outputs 1 if and only if the number of the edges in G is at least $\frac{n^2}{4} + \frac{k^2}{4}$. Then, A distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage $\Omega\left(\frac{k^2}{n}\right)$ for every $n \in \mathbb{N}$ and every $k \leq \sqrt{n}$.*

Proof. Let M be the number of edges of the given graph. If $G \sim \mathcal{G}(n, 1/2, k)$, then $M - \binom{k}{2} \sim \mathrm{Bin}\left(\binom{n}{2} - \binom{k}{2}, 1/2\right)$. Therefore, we have $\Pr[A(\mathcal{G}(n, 1/2, k)) = 1] = \frac{1}{2} - o(1)$.

Let $\Phi(c) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{c} \mathrm{e}^{-x^2/2} \mathrm{d}x$ denote the cumulative distribution function of the standard normal distribution. If $G \sim \mathcal{G}(n, 1/2)$, then $M \sim \mathrm{Bin}\left(\binom{n}{2}, 1/2\right)$. By the Berry–Esseen theorem[7], with probability $1 - \Phi(c) - O(1/n)$, we have $M \geq \frac{n^2}{2} + c\sqrt{\frac{n^2}{2}}$. If we set $c = \frac{\sqrt{2}k^2}{4n}$, then $\Pr[A(\mathcal{G}(n, 1/2)) = 1] = 1 - \Phi(c) + O(1/n)$.

Then, $c = o(1)$ and by the Taylor expansion, we obtain $\Phi(c) = \Phi(0) + \Phi'(0)c + O(c^2) = \frac{1}{2} + \frac{c}{\sqrt{2\pi}} + O(c^2)$. Therefore, the advantage of A is $\frac{c}{\sqrt{2\pi}} - O(c^2) - O(1/n) = \Omega\left(\frac{k^2}{n}\right)$. □

C Proof of Exchange Lemma

In this section, we prove the exchange lemma for samplers, which asserts that if (Y, X) is a sampler, then so is (X, Y).

Lemma 3.12 (Exchange Lemma)**.** *If (Y, X) is a $\left(\frac{\varepsilon}{2}, \frac{\delta\varepsilon}{8}\right)$-sampler, then, (X, Y) is a (δ, ε)-sampler.*

Proof. We first show that (X, Y) is a "one-side" sampler, that is, for any function $S\colon \mathsf{supp}(Y) \to [0, 1]$,

$$\Pr_{x \sim X}[\mathbb{E}[S(Y) \mid X = x] - \mathbb{E}[S(Y)] \leq -\varepsilon] \leq \frac{\delta}{2}. \tag{2}$$

Then, we prove the "two-side" version using (2).

Let $\delta' = \frac{\varepsilon}{2}$ and $\varepsilon' = \frac{\delta\varepsilon}{8}$ and thus (Y, X) is a (δ', ε')-sampler. Fix a function $S\colon \mathsf{supp}(Y) \to [0, 1]$ and define a function $H\colon \mathsf{supp}(X) \to \{0, 1\}$ by

$$H(x) = 1 \iff \mathbb{E}[S(Y) \mid X = x] - \mathbb{E}[S(Y)] \leq -\varepsilon.$$

By definition of H, we have

$$\begin{aligned}\mathbb{E}[H(X)S(Y)] &= \mathbb{E}[S(Y) \mid H(X) = 1]\,\mathbb{E}[H(X)] \\ &\leq (\mathbb{E}[S(Y)] - \varepsilon)\,\mathbb{E}[H(X)].\end{aligned}$$

[7] The Berry–Esseen theorem asserts that $\left|\Pr\left[\mathrm{Bin}(\ell, p) - \ell p \leq x\sqrt{\ell p(1-p)}\right] - \Phi(x)\right| = O\left(\ell^{-1/2}\right)$ (cf. Theorem A.1 of [DGP14]).

On the other hand, since (Y, X) is a (δ', ε')-sampler, for a $(1-\delta')$-fraction of $y \sim Y$, we have $\mathbb{E}[H(X) \mid Y = y] > \mathbb{E}[H(X)] - \varepsilon'$. Let $T \subseteq \mathsf{supp}(Y)$ be the set of such y. Then, we have

$$\begin{aligned}\mathbb{E}[H(X)S(Y)] &\geq \sum_{y \in T} \mathbb{E}[H(X) \mid Y = y] S(y) \Pr[Y = y] \\ &\geq \big(\mathbb{E}[H(X)] - \varepsilon'\big)\big(\mathbb{E}[S(Y)] - \Pr[Y \in T]\big) \\ &\geq \big(\mathbb{E}[H(X)] - \varepsilon'\big)\big(\mathbb{E}[S(Y)] - \delta'\big).\end{aligned}$$

Therefore, we obtain

$$\mathbb{E}[H(X)](\mathbb{E}[S(Y)] - \varepsilon) > \big(\mathbb{E}[H(X)] - \varepsilon'\big)\big(\mathbb{E}[S(Y)] - \delta'\big). \tag{3}$$

To prove (2), it suffices to show that $\mathbb{E}[H(X)] \leq \frac{\delta}{2}$. Suppose for contradiction that $\mathbb{E}[H(X)] \geq \frac{\delta}{2}$. Note that we may assume that $\mathbb{E}[S(Y)] \geq \varepsilon$ (otherwise, (2) is trivial).

If $\mathbb{E}[S(Y)] = \varepsilon$, then we obtain

$$\left(\mathbb{E}[H(X)] - \frac{\delta\varepsilon}{8}\right)\left(\mathbb{E}[S(Y)] - \frac{\varepsilon}{2}\right) < 0,$$

which contradicts with $\mathbb{E}[H(X)] \geq \frac{\delta}{2}$.

If $\mathbb{E}[S(Y)] > \varepsilon$, then we obtain

$$\begin{aligned}\frac{(\mathbb{E}[H(X)] - \varepsilon')(\mathbb{E}[S(Y)] - \delta')}{\mathbb{E}[H(X)](\mathbb{E}[S(Y)] - \varepsilon)} &= \left(1 - \frac{\delta\varepsilon}{8\,\mathbb{E}[H(X)]}\right)\left(1 + \frac{\varepsilon/2}{\mathbb{E}[S(Y)] - \varepsilon}\right) \\ &\geq \left(1 - \frac{\varepsilon}{4}\right)\left(1 + \frac{\varepsilon}{2}\right) \\ &\geq 1,\end{aligned}$$

which contradicts with (3). This proves (2).

Now we prove that (X, Y) is a (δ, ε)-sampler. Fix $S \colon \mathsf{supp}(Y) \to [0, 1]$ and apply (2) for S and $1 - S$. Then, we have

$$\begin{aligned}\Pr_{x \sim X}[\mathbb{E}[S(Y) \mid X = x] - \mathbb{E}[S(Y)] \leq -\varepsilon] &\leq \frac{\delta}{2}, \\ \Pr_{x \sim X}[-\mathbb{E}[S(Y) \mid X = x] + \mathbb{E}[S(Y)] \leq -\varepsilon] &\leq \frac{\delta}{2}.\end{aligned}$$

By the union bound, we obtain the claim. □

We also prove the exchange lemma for one-sided multiplicative samplers. This result is known in [HS23, Lemma 3.7].

Lemma 6.12. *If (Y, X) is a one-sided multiplicative $\left(\frac{c\varepsilon}{2}, \frac{c}{2}\right)$-sampler for density δ, then (X, Y) is a one-sided multiplicative (δ, c)-sampler for density ε.*

Proof. Let $c' = \frac{c}{2}$, $\delta' = \frac{c\varepsilon}{2}$, and $\varepsilon' = \delta$. Fix a function $S \colon \mathsf{supp}(Y) \to [0, 1]$ with $\mathbb{E}[S(Y)] \geq \varepsilon$ and define $H \colon \mathsf{supp}(X) \to \{0, 1\}$ by

$$H(x) = 1 \iff \mathbb{E}[S(Y) \mid X = x] \leq (1 - c)E[S(Y)].$$

It suffices to show that $\mathbb{E}[H(X)] \le \delta$. Suppose for contradiction that $\mathbb{E}[H(X)] > \delta$. By definition of H, we have

$$\mathbb{E}[S(Y)H(X)] = \mathbb{E}[H(X)]\,\mathbb{E}[S(Y) \mid H(X) = 1] \le (1-c)\,\mathbb{E}[H(X)]\,\mathbb{E}[S(Y)].$$

On the other hand, since (Y, X) is a one-sided multiplicative (δ', c')-sampler for density ε', for a $(1-\delta')$-fraction of $y \sim Y$, we have $\mathbb{E}[H(X) \mid Y = y] \ge (1-c')\,\mathbb{E}[H(X)]$. Let T be the set of such $y \in \mathsf{supp}(Y)$. Then, we have

$$\begin{aligned}
\mathbb{E}[S(Y)H(X)] &\ge \sum_{y \in T} \mathbb{E}[H(X) \mid Y = y] S(y) \Pr[Y = y] \\
&\ge (1-c')\,\mathbb{E}[H(X)](\mathbb{E}[S(Y)] - \Pr[Y \in T]) \\
&\ge (1-c')\,\mathbb{E}[H(X)](\mathbb{E}[S(Y)] - \delta') \\
&\ge (1-c')\left(1 - \frac{\delta'}{\varepsilon}\right)\mathbb{E}[H(X)]\,\mathbb{E}[S(Y)].
\end{aligned}$$

Therefore, we obtain

$$(1-c')\left(1 - \frac{\delta'}{\varepsilon}\right) \le 1 - c.$$

However, by the choice of parameters,

$$(1-c')\left(1 - \frac{\delta'}{\varepsilon}\right) = \left(1 - \frac{c}{2}\right)^2 > 1 - c$$

and we obtain the contradiction. Therefore, we have $\mathbb{E}[H(X)] \le \delta$. □

D Proof of Previous Results

D.1 Search to Decision Reduction by Alon et al.

Lemma 5.11 ([AAKMRX07])**.** *Let $k \ge 18 \log n$ for a sufficiently large constant $c > 0$. Suppose there exists a randomized polynomial-time algorithm A that, for all $k' \ge k/3$, distinguishes $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ with advantage $1 - \frac{\delta}{n}$. Then, there exists a randomized polynomial-time algorithm A' that, for every $k' \ge k$, finds a k'-clique in $\mathcal{G}(n, 1/2, k')$ with probability $1 - 2\delta - n\mathrm{e}^{-k/18}$.*

Proof. Suppose for simplicity that the algorithm A satisfies, for all $k' \ge k/3$,

$$\mathop{\mathbb{E}}_{A, G \sim \mathcal{G}(n,1/2,k')}[A(G)] - \mathop{\mathbb{E}}_{A, G \sim \mathcal{G}(n,1/2)}[A(G)] \ge 1 - \frac{\delta}{n}.$$

Let $\mathrm{Resample}(G, F)$ be the algorithm of Definition 5.6. Our algorithm A' runs as follows on input $G \sim \mathcal{G}(n, 1/2, k)$:

1. Let $S \leftarrow \emptyset$.
2. For each $i = 1, \ldots, n$, do the following:
   (a) Let π_i be a uniformly random permutation over $[n]$.
   (b) Let $G_i \leftarrow \pi_i(\mathrm{Resample}(G, E(U, U)))$, where $U = \{i\} \cup \Gamma_G(i)$.
   (c) If $A(G_i) = 0$, let $S \leftarrow S \cup \{i\}$.

3. Output S.

We prove the correctness of A'. Let $G \sim \mathcal{G}(n, 1/2, k')$ be the input and C be the planted location. Consider the marginal distribution of each G_i.

Conditioned on $i \in C$, the marginal distribution of each G_i is $\mathcal{G}(n, 1/2)$ since the edges inside C are resampled. Therefore, the output S satisfies

$$\Pr[i \in S \mid i \in C] = \Pr[A(G_i) = 0 \mid i \in C] = \Pr_{G \sim \mathcal{G}(n,1/2)}[A(G) = 0] \geq 1 - \frac{\delta}{n}.$$

Let $d_C(i) = |C \setminus \Gamma_G(i)|$. Conditioned on $i \notin C$ and $d_C(i)$, the marginal distribution of G_i is identical to $\mathcal{G}(n, 1/2, d_C(i))$ since edges inside $C \cap \Gamma_G(i)$ are resampled and the location of the remaining clique $C \setminus \Gamma_G(i)$ is uniformly random in G_i due to the random shuffle π_i. Moreover, for $i \notin C$, the marginal distribution of $d_C(i)$ is $\mathrm{Bin}(k', 1/2)$ (over the random choice of $G \sim \mathcal{G}(n, 1/2, k')$). By the Chernoff bound (Lemma 3.3), we have $\Pr[d_C(i) \leq k/3 \mid i \notin C] \leq \Pr[\mathrm{Bin}(k, 1/2) \leq k/3] \leq \mathrm{e}^{-k/18}$. Therefore, the output S satisfies

$$\begin{aligned}
\Pr[i \in S \mid i \notin C] &= \Pr[A(G_i) = 0 \mid i \notin C] \\
&\leq \Pr[A(G_i) = 0 \mid d_C(i) \geq k/3 \text{ and } i \notin C] + \mathrm{e}^{-k/18} \\
&\leq \frac{\delta}{n} + \mathrm{e}^{-k/18}.
\end{aligned}$$

By the union bound over $i \in [n]$, we have $\Pr[A'(G) = C] \leq 1 - 2\delta - n\mathrm{e}^{-k/18}$. □

E Boosting via Random Partition

We present a simple boosting technique, which is inspired by the proof of Hazan and Krauthgamer [HK11, Lemma 2.2]. This technique transforms an algorithm that distinguishes $\mathcal{G}(n, 1/2, k/t)$ from $\mathcal{G}(n, 1/2)$ with advantage $\varepsilon \gg 1/\sqrt{k}$ into another algorithm that distinguishes $\mathcal{G}(n, 1/2, k)$ from $\mathcal{G}(n, 1/2)$ with advantage≈ 1, where $t \in \mathbb{N}$ is a parameter that will be specified later.

Definition E.1. *An algorithm A* predicts $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, 1/2)$ with advantage γ *if*

$$\Pr_{G \sim \mathcal{G}(n,1/2,k), A}[A(G) = 1] \geq \frac{1+\gamma}{2} \text{ and } \Pr_{G \sim \mathcal{G}(n,1/2), A}[A(G) = 0] \geq \frac{1+\gamma}{2}.$$

For a given graph $G = (V, E)$, we partition V into t parts $V_1, \ldots, V_t$ randomly. By the Chernoff bound, each part has size $|V_i| \approx n/t$. For each induced subgraph $G[V_i]$, we add $n - |V_i|$ vertices with random edges and then shuffle vertices by a random permutation. Let G_i be the resulting graph (see Figure 2).

We briefly explain how to amplify the success probability. If $G \sim \mathcal{G}(n, 1/2)$, then $G_1, \ldots, G_t \sim \mathcal{G}(n, 1/2)$ are independent. If $G \sim \mathcal{G}(n, 1/2, k)$, then the distribution of each G_i is $\mathcal{G}(n, 1/2, \ell)$ for $\ell \approx k/t$. Moreover, conditioned on the size of planted cliques in $G_1, \ldots, G_t$, these t graphs are independent. Therefore, if $A(G_i)$ correctly decides $G \sim \mathcal{G}(n, 1/2, k/t)$ or $G \sim \mathcal{G}(n, 1/2)$ with success probability $1/2 + \varepsilon$, by taking the majority among $A(G_1), \ldots, A(G_t)$, we can amplify the success probability to $1 - \mathrm{e}^{-\Theta(\varepsilon^2 t)}$. This argument is formalized as follows.

Theorem E.2. *Let $t \in \mathbb{N}$ be a parameter. Suppose there exists a randomized polynomial-time algorithm that predicts $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, k')$ with advantage ε for all $k' \geq \frac{k}{2t}$. Then, there exists a randomized polynomial-time algorithm that distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, k)$ with advantage $1 - \mathrm{e}^{-\Theta(\varepsilon^2 t)} - \mathrm{e}^{-\Theta(k/t)}$.*

Remark E.3. *Theorem E.2 makes sense if $\varepsilon^{-2} \ll t \ll k$. In other words, we can apply Theorem E.2 if the advantage satisfies $\varepsilon \gg 1/\sqrt{k}$.*

Proof of Theorem E.2. Let A be the algorithm that distinguishes $\mathcal{G}(n, 1/2, k')$ from $\mathcal{G}(n, 1/2)$ with advantage $\varepsilon(n)$ for all $k' \geq \frac{k}{2t}$. Our aim is to construct another algorithm A' that distinguishes $\mathcal{G}(n, 1/2, k)$ and $\mathcal{G}(n, k)$ with a high advantage.

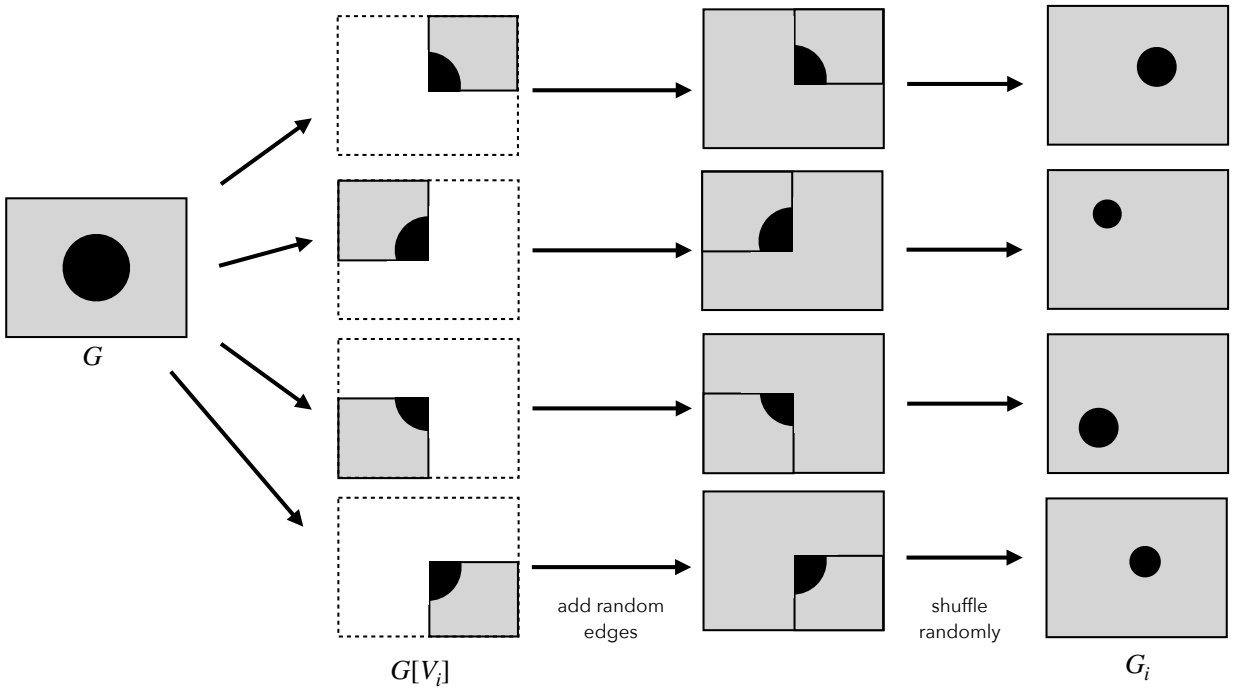

Figure 2: Graphs $G_1, \ldots, G_t$ produced by $A'(G)$. The size of planted clique in each G_i is concentrated on k/t.

Let A' be the algorithm that, given an n-vertex graph $G = ([n], E)$ as input, runs as follows:

1. Let $V_1, \ldots, V_t \subseteq [n]$ be a random partition of $[n]$. Formally, for each $v \in [n]$ is assigned to V_i for independently random $i \sim [t]$.
2. For every $i \in [t]$, let $G_i = \pi_i(\text{Resample}(G, E(V_i, [n])))$, where π_i is a uniformly random permutation over $[n]$ and Resample is the algorithm of Definition 5.6.
3. Output the majority among $A(G_1), \ldots, A(G_t)$.

Suppose $G \sim \mathcal{G}(n, 1/2)$. Since $V_1, \ldots, V_t$ are disjoint, the marginal distribution of $(G_1, \ldots, G_t)$ is the t-wise direct product of $\mathcal{G}(n, 1/2)$. Therefore, by the Chernoff bound, $A'(G)$ outputs 0 with probability $1 - \mathrm{e}^{-\Theta(\varepsilon^2 t)}$ over the choice of $G \sim \mathcal{G}(n, 1/2)$ and the internal randomness of A'.

Suppose that the input G is drawn from $\mathcal{G}(n, 1/2, k)$ and let C be the planted location. Let $V_1, \ldots, V_t$ be the random partition considered in Step 1 and let $C_i = C \cap V_i$. Conditioned on $|C_i|$ for all $i = 1, \ldots, t$, the location of each C_i in G_i is independently and uniformly distributed. Similarly, edges of G_i outside C_i are uniformly distributed. Therefore, $G_i \sim \mathcal{G}(n, 1/2, |C_i|)$ are independent conditioned on $|C_i|$.

Let $\mathcal{E}$ be the event that $|V_i| \geq \frac{k}{2t}$ for all $i \in [t]$. By the Chernoff bound and union bound over $i \in [t]$, the event $\mathcal{E}$ occurs with probability $1 - t \cdot \mathrm{e}^{-\frac{k}{8t}}$. Let $Z_i = A(G_i)$ be the binary random variable. Since A predicts $\mathcal{G}(n, 1/2, k')$ and $\mathcal{G}(n, 1/2)$ with advantage ε for all $k' \geq \frac{k}{2t}$, conditioned on $\mathcal{E}$, the random variables $Z_1, \ldots, Z_t$ are indepednent and $\mathbb{E}[Z_i] \geq \frac{1+\varepsilon}{2}$. Therefore, we obtain

$$\begin{aligned}\Pr\left[Z_1 + \cdots + Z_t \leq \frac{t}{2} \middle| C\right] &\leq \Pr\left[Z_1 + \cdots + Z_t \leq \frac{t}{2} \middle| \mathcal{E}\right] + t \cdot \exp\left(-\frac{k}{8t}\right) \\ &\leq \exp\left(-\frac{\varepsilon^2 t}{64}\right) + t \cdot \exp\left(-\frac{k}{8t}\right).\end{aligned}$$

The claim follows by taking the expectation over $C \sim \binom{[n]}{k}$. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il