

Algorithmic Decorrelation and Planted Clique in Dependent Random Graphs: The Case of Extra Triangles

Guy Bresler* Chenghao Guo† Yury Polyanskiy‡

Massachusetts Institute of Technology

Abstract

We aim to understand the extent to which the noise distribution in a planted signal-plus-noise problem impacts its computational complexity. To that end, we consider the planted clique and planted dense subgraph problems, but in a different ambient graph. Instead of Erdős-Rényi $G(n, p)$, which has independent edges, we take the ambient graph to be the *random graph with triangles* (RGT) obtained by adding triangles to $G(n, p)$. We show that the RGT can be efficiently mapped to the corresponding $G(n, p)$, and moreover, that the planted clique (or dense subgraph) is approximately preserved under this mapping. This constitutes the first average-case reduction transforming dependent noise to independent noise. Together with the easier direction of mapping the ambient graph from Erdős-Rényi to RGT, our results yield a strong equivalence between models. In order to prove our results, we develop a new general framework for reasoning about the validity of average-case reductions based on *low sensitivity to perturbations*.

*Supported in part by NSF Career award CCF-1940205.

†Supported in part by NSF TRIPODS grant DMS-2022448, NSF Career award CCF-1940205, CCF-2131115 and the MIT-IBM Watson AI Lab.

‡Supported in part by the NSF grant CCF-2131115 and the MIT-IBM Watson AI Lab.

Contents

1	Introduction	1
1.1	The Problems We Consider	2
1.2	Our results	3
2	General Result and Application to RGT	5
2.1	Reductions and Sensitivity to Perturbation	5
2.2	Mapping between Erdős-Rényi and RGT	8
2.3	Applying Theorem 2.1 to Triangle Removal Algorithm	9
2.4	Showing Perturbation Insensitivity: Main Ideas Behind Lemma 2.5	11
3	Strong Equivalence Between Parameterized Problems	13
4	Preliminaries	17
4.1	Probability Notation	17
4.2	Graph Theory Notation	17
4.3	Total Variation Distance	18
4.4	Chernoff Bound and Mcdiarmid's Inequality	20
5	Forward Map Preserves Planted Signal	20
5.1	Class of Graphs \mathcal{G}	21
5.2	Checking Conditions in Theorem 2.1	21
5.3	Perturbation Insensitivity of \mathcal{A}_Δ and Proof of Lemma 5.3	22
6	Properties of Triangle Distribution μ_G	26
7	Reverse Map Preserves Planted Signal	28
7.1	Influence of an Edge on Non-Containing Triangles	29
7.2	Replacing X^+ and Y^+ by X^* and Y^*	35
7.3	Perturbation Insensitivity and Proof of Lemma 2.5	36
7.4	Proof of Lemma 7.12 via Likelihood Concentration	39
7.5	The class of graphs \mathcal{G} has high-probability	45
8	Low Influences for Marginally Small Distributions	47
8.1	Low Influences for Marginally Small Distributions	48
8.2	Influences in $\mathcal{L}_G(Y)$	50
9	Concentration and Mixing for Marginally Small Distributions	53
9.1	Glauber Dynamics and Path Coupling	54
9.2	Concentration from Mixing	55
A	Relation Between Models	61
A.1	RGT vs Erdős-Rényi	61
A.2	RGT vs RIG	65
B	Deferred Proofs	66
B.1	Proof of Lemma 7.13 on the distance between $K_{uv}(X)$ and X	66
B.2	Proof of Lemma A.7	71
B.3	Proof of Lemma 7.2	72
C	Technical Lemmas	72

1 Introduction

Most modern statistical inference problems exhibit a striking phenomenon: the best efficient algorithm requires substantially more data, or lower noise level, than the information-theoretic limit achieved by inefficient algorithms. Such problems are said to exhibit a *statistical-computational gap*. In this paper we are interested in *average-case* planted problems, meaning that their inputs are sampled according to some probability distribution with planted structure. In order to develop methodology for rigorously reasoning about the computational complexity of these average-case problems, a vibrant line of research at the interface of statistics, probability, and computational complexity has emerged. There are two main approaches to substantiating computational limits: hardness against specific classes of algorithms and average-case reductions.

The first approach attempts to determine the best possible performance for restricted classes of algorithms, such as low-degree polynomials [Hop18, Wei22], sum-of-squares (SoS) relaxations [BHK⁺16], [HKP⁺17], statistical query algorithms [FGR⁺13, FPV15, DKS17], first-order methods [CMW20], or classes of circuits [Ros08, Ros14]. Beyond hardness against specific classes of algorithms, the overlap gap method introduced in [GS14] and refined in [HS22] has revealed how structural properties of a solution landscape can rule out multiple classes of algorithms. The power of distinct classes of algorithms for solving families of statistical problems have been related: low-degree polynomials vs. statistical query [BBH⁺20], low-degree polynomials vs. approximate message passing [MW22], and SoS vs. spectral applied to matrices of low-degree polynomials [HKP⁺17]. This broad approach continues to see intense activity and yields insight into the limits of current algorithms. However, it has the drawbacks that (1) each given class of algorithms is known to be suboptimal in certain settings [KM22, WEAM19, ZSWB22], which reduces confidence in lower bounds against any specific class, and (2) one must prove fresh lower bounds for each problem of interest and for each new class of algorithms that emerges.

The second approach entails devising average-case reductions that map one statistical problem to another using a polynomial-time algorithm, whereby hardness of the first problem is transferred to the second. The reduction approach is a foundational tool across complexity theory and cryptography for elucidating relationships between problems and constructing complexity hierarchies. Moreover, it provides insights applicable to all algorithms. The main drawback of this approach is that in the average-case setting, one must precisely map not just problem instances, but rather entire probability distributions over inputs. This is notoriously challenging due to a lack of techniques, as emphasized by Goldreich [Gol11], Barak [Bar17], and Bogdanov and Trevisan [BT06].

A common starting hardness assumption in the literature on average-case reductions is the Planted Clique (PC) Conjecture. This conjecture posits that no polynomial-time algorithm can detect a planted clique of size k in an Erdős-Rényi random graph with edge density p , when p is constant and $k = o(\sqrt{n})$. Many problems have been shown to be computationally hard based on this assumption including sparse principle component analysis [BR13a, WBS16, GMZ17, BBH18, BB19], submatrix detection or biclustering [MW15, BBH18, BBH19], planted dense subgraph [HWX15, BBH18]. A variety reduction techniques were introduced in [BBH18] and [BB20], resulting in a web of reductions from PC to problems including robust sparse mean estimation, tensor PCA, general planted dense subgraph, dense stochastic block model, mixtures of sparse linear regressions, robust sparse mean estimation, and many others. Some works show reductions from hypergraph planted clique or hypergraph planted dense subgraph to other problems, including tensor clustering [LZ22] and SVD for random 3-tensors [ZX17]. Another line of research has also explored reductions to statistical problems based on cryptographic assumptions [GVV22, SZB21, CGKM22, DV21].

Despite these advances, the current literature on average-case reductions for statistical problems exhibits a notable limitation: the absence of two-way equivalences. Most reductions rely on a single hardness assumption, such as the Planted Clique Conjecture or Learning Parity with Noise (LPN), and demonstrate a statistical-computational gap for other problems. Going beyond this, one might dream of partitioning statistical problems into equivalence classes, just as has been done so fruitfully in worst-case complexity. We mention here, only briefly, that the notion of equivalence is itself nontrivial because each statistical “problem” is itself an entire *parameterized family* of problems. We rigorously define what we mean by strong computational equivalence in Section 3.

The primary obstacle to proving equivalence between statistical problems, just as for one-way reductions, is that average-case reductions are challenging: we must transform one noise distribution into another while preserving the planted signal. To date, all reductions have converted models with independent noise to models with either independent or dependent noise. The challenge in establishing equivalence between planted problems is therefore in developing techniques for removing dependence in the noise of high-dimensional distributions. In this paper, we introduce a general framework for the validity of reductions based on low sensitivity to perturbations and apply it to show the first non-trivial computational equivalence between two statistical problems.

A second (related) motivation for the present paper is towards developing an understanding of how general are the phenomena observed in planted statistical problems. For planted matrices or tensors with independent entries, a reduction was devised from planted clique or planted dense subgraph to a broad class of entry distributions by [BBH19], thereby showing that the observed phenomena are *universal* for this class. Thus, we understand how to change the distribution of the entries, but as of yet there are no general techniques for connecting problems with different *dependence* between the entries¹.

1.1 The Problems We Consider

General Hypothesis Testing Setup In a simple-versus-simple hypothesis testing problem there is a pair of distributions P_0 and P_1 on space Ω , and one observes X generated according to one of the distributions. The task is to decide between the two hypotheses

$$H_0 : X \sim P_0 \quad \text{versus} \quad H_1 : X \sim P_1 .$$

We consider noisy signal detection problems indexed by problem size parameter n , where the space is Ω_n , $P_{0,n}$ represents the pure noise distribution, and $P_{1,n}$ is the distribution with planted signal. We often keep the dependence on n implicit. An algorithm $\Phi : \Omega \rightarrow \{0, 1\}$ *solves* the hypothesis testing problem if the sum of type I and type II errors asymptotically vanish, i.e.,

$$\mathbb{P}_{X \sim P_0}(\Phi(X) = 1) + \mathbb{P}_{X \sim P_1}(\Phi(X) = 0) \rightarrow 0 \quad \text{as} \quad n \rightarrow \infty .$$

Planted Dense Subgraph Problem In this paper, we consider a slight generalization of Planted Clique known as the *Planted Dense Subgraph (PDS)* problem. Here, the pure noise distribution is the simple Erdős-Rényi random graph, $G(n, p)$. The alternative hypothesis, denoted by $G(n, p, k, q)$, is generated by starting with $G(n, p)$, choosing a uniformly random subset S of size k from $[n]$, erasing all edges in S , and independently resampling the edges in S to be included with probability

¹The exception is sparse PCA, which has a very specific dependence structure.

q . The PDS problem is, given a graph G , to distinguish between the two hypotheses

$$H_0 : G \sim G(n, p) \quad \text{and} \quad H_1 : G \sim G(n, p, k, q).$$

It is believed that the PDS problem with constant p and q has the same $k = o(\sqrt{n})$ computational threshold as the planted clique problem [BHK⁺16], [Hop18].

Random Graph with Triangles Next, instead of the ambient graph being Erdős-Rényi, we consider a simple graph distribution with dependent edges.

Definition 1 (Random Graph with Triangles). The random graph with triangles distribution $\text{RGT}(n, p, p')$ is the law of a graph generated as follows. Let $G \sim G(n, p)$ be an Erdős-Rényi random graph and for every triple of nodes $\{i, j, k\} \in \binom{[n]}{3}$ with probability p' independently add the triangle consisting of the three edges (i, j) , (j, k) , and (i, k) to $E(G)$.

Remark 1. The expected edge density of $G(n, p, p')$ is $p + (1 - p)(1 - (1 - p')^{n-2})$. It is always assumed that $p' = O(1/n)$ so that the edge density remains bounded away from 1.

Remark 2. In certain parameter regimes, the RGT model is close in total variation distance to the Random Intersection Graph, a model inspired by real-world networks. This will be explained in more depth in Section 1.2.

Planted Dense Subgraph in Random Graph with Triangles We plant a signal in the random graph with triangles (*planted RGT*) to obtain $\text{RGT}(n, p, p', k, q)$, the law of a graph generated as follows. Start with $\text{RGT}(n, p, p')$, choose a random subset S of k vertices, erase all edges within S and then independently include each edge within S with probability q . The Planted RGT problem is, given a graph G , to distinguish between the two hypotheses

$$H_0 : G \sim \text{RGT}(n, p, p') \quad \text{and} \quad H_1 : G \sim \text{RGT}(n, p, p', k, q).$$

Remark 3. There are multiple options for the definition of Planted RGT. We will discuss a substantial generalization in the next section. One could also, for example, plant a smaller RGT with different edge or triangle densities within a larger RGT. We leave this for future work.

1.2 Our results

Our main result is a strong computational equivalence between the planted dense subgraph (PDS) and planted random graph with triangles (Planted RGT) problems. To achieve this, we must design transformations that connect these problems in both directions. Here, and throughout the paper, for two distributions μ and ν we use $\mu \approx_\epsilon \nu$ as shorthand for $d_{\text{TV}}(\mu, \nu) \leq \epsilon$.

Theorem 1.1. *For any $k = o(n^{1/4} \log^{-17/4} n)$, $p' \leq 1/(n \log n)$ and $0 < p, q < 1$, there are efficient algorithms \mathcal{A}_Δ and \mathcal{A} that satisfy the following transition properties.*

1. \mathcal{A}_Δ maps both H_0 and H_1 hypotheses of PDS to planted RGT with parameter map f :

$$\mathcal{A}_\Delta(G(n, p)) \approx_{o_n(1)} \text{RGT}(n, p, p'), \quad \text{and} \quad \mathcal{A}_\Delta(G(n, p, k, q)) \approx_{o_n(1)} \text{RGT}(n, p, p', k, f(q)),$$

2. \mathcal{A} maps both H_0 and H_1 hypotheses of planted RGT to PDS with parameter map g :

$$\mathcal{A}(\text{RGT}(n, p, p')) \approx_{o_n(1)} G(n, p), \quad \text{and} \quad \mathcal{A}(\text{RGT}(n, p, p', k, q)) \approx_{o_n(1)} G(n, p, k, g(q)).$$

Moreover, the parameter mappings f and g satisfy $f \circ g(q) = q + o_n(1)$ and $g \circ f(q) = q + o_n(1)$.

Part 1 entails transforming from PDS to planted RGT, i.e., an independent noise model to a dependent noise model. In the case of $G(n, p)$ to $\text{RGT}(n, p, p')$, the mapping itself is straightforward based on the definition of RGT: one simply adds each triangle with probability p' . However, the fact that the planted distributions are mapped correctly via the very same procedure does still require a nontrivial argument.

In Part 2, the other direction, we show how to map from RGT to Erdős-Rényi. Here the mapping even in the unplanted pure-noise case is not obvious, and we describe it in Section 2.2. It is obtained by viewing the triangle addition procedure as a Markov transition kernel on graphs, and implementing its time-reversal.

To constitute a valid reduction from Planted RGT to PDS, the same mapping must also work when a dense subgraph has been planted. We face a delicate balance between two competing objectives: (1) maintaining the algorithm's behavior outside of the dense subgraph, ensuring that its performance remains approximately invariant to the placement of the dense subgraph, and (2) removing the dependence between triangles, which inherently necessitates dependence in the algorithm's behavior.

Validity of both reductions is proved via application of a general framework that we develop, stated as Theorem 2.1 in Section 2. Applying our general theorem presents a number of technical challenges in the form of bounding total variation distances between perturbed distributions. Methods for proving the proximity of two high-dimensional distributions in total variation distance are notably limited, and even more so when neither of them is a product distribution. To this end, we introduce several technical innovations that will be discussed in detail in the next section.

We conjecture that the statement of Theorem 1.1 holds for k up to $\tilde{O}(\sqrt{n})$, while our result only applies for $\tilde{O}(n^{1/4})$. The range of p' is optimal up to log factors: if $p' \geq 2.1n^{-1} \ln n$, then with high probability the RGT is the complete graph and planted RGT is information-theoretically impossible, so the planted RGT cannot be equivalent to PDS in this regime. We leave open the problem of obtaining a similar result for the full range of k, p, q , and p' .

General Planted Signal Our results not only demonstrate that the detection of planted dense subgraphs problem is equivalent in Erdős-Rényi graphs and RGT, but also reveal that a wide range of planted signal detection problems are equivalent in both models. Let us define the general planted signal (GPS) model.

Definition 2 (General Planted Signal (GPS)). For edge sequence $\vec{e} = (e_1, \dots, e_K)$ and $\vec{p} = (p_1, \dots, p_K)$ let $\text{GPS}_{\vec{e}, \vec{p}}$ take as input a graph, and for each $1 \leq i \leq K$, edge e_i is resampled to be included in the graph with probability p_i . For a pure-noise distribution \mathbf{P}_0 , let $\text{GPS}(\mathbf{P}_0)$ denote the distribution of $\text{GPS}(G)$ where $G \sim \mathbf{P}_0$.

Planted dense subgraph is a special case of GPS with $e_1, \dots, e_{\binom{k}{2}}$ being the edges inside the dense subgraph, and $p_i = q$ is the planted edge density. The GPS problem in random graph distribution \mathbf{P}_0 is the hypothesis testing problem between \mathbf{P}_0 and $\text{GPS}(\mathbf{P}_0)$.

Our bi-directional reduction result holds also in this general setting.

Theorem 1.2. Let $\text{GPS}_n = \text{GPS}_{\vec{e}, \vec{p}}$ be a sequence of general planted signal problems. For any $K = o(\sqrt{n} \log^{-17/2} n)$, $p' \leq 1/(n \log n)$ and p_1, \dots, p_K each bounded away from 0 and 1 by a constant, there are distributions $\text{GPS}_n^g = \text{GPS}_{\vec{e}, g(\vec{p})}$ and $\text{GPS}_n^f = \text{GPS}_{\vec{e}, f(\vec{p})}$ (where $f(\vec{p})$ stands for $f : \mathbb{R} \rightarrow \mathbb{R}$ applied to all elements of the vector) and efficient algorithms \mathcal{A} , \mathcal{A}' satisfying the transition properties.

1. $\mathcal{A}(G(n, p)) \approx_{o_n(1)} \text{RGT}(n, p, p')$ and $\mathcal{A}(\text{GPS}_n(G(n, p))) \approx_{o_n(1)} \text{GPS}_n^f(\text{RGT}(n, p, p'))$, and
2. $\mathcal{A}'(\text{RGT}(n, p, p')) \approx_{o_n(1)} G(n, p)$ and $\mathcal{A}'(\text{GPS}_n(\text{RGT}(n, p, p'))) \approx_{o_n(1)} \text{GPS}_n^g(G(n, p))$.

Moreover, the parameter mappings f and g satisfy $f \circ g(q) = q + o_n(1)$ and $g \circ f(q) = q + o_n(1)$.

Remark 4. GPS includes Subgraph Stochastic Block Model (SSBM), which has a rank-1 signal (i.e., a small SBM) on k vertices planted inside an Erdős-Rényi graph. Its complete phase diagram was determined by [BBH18], with computationally hard region obtained via reduction from the planted clique conjecture. Our result shows that for $k = o(n^{1/4} \log^{-17/4} n)$, the computational complexity of SSBM is the same in ambient graph being Erdős-Rényi or RGT.

Approximation of Random Intersection Graph Aside from RGT being intrinsically interesting as a natural random graph model with low-order dependence, it turns out that the RGT non-trivially approximates the random intersection graph.

Definition 3 (Random Intersection Graph). A sample from $\text{RIG}(n, d, \delta)$ is defined by sampling n sets $S_1, \dots, S_n \subset [d]$ where each S_i includes each element of $[d]$ independently with probability δ . Vertices i and j are connected in G if and only if S_i and S_j have non-empty intersection.

Theorem 1.3. For an RIG with constant edge density, i.e., $d = \Theta(1/\delta^2)$, if $d \gg n^{2.5}$, then

$$d_{\text{TV}}(\text{RIG}(n, d, \delta), \text{RGT}(n, 1 - e^{-d\delta^2 + (n-2)d\delta^3(1-\delta)^{n-3}}, 1 - e^{-d\delta^3(1-\delta)^{n-3}})) = o(1).$$

It was shown in [BBN19] that the threshold for distinguishing RIG from Erdős-Rényi occurs at $d \sim n^3$. Thus, the RGT is close in total variation in the range $n^{2.5} \ll d \ll n^3$, whereas in this range the RIG and Erdős-Rényi have total variation distance close to 1.

2 General Result and Application to RGT

We start by introducing a general theorem that shows how a mapping between unplanted distributions yields a mapping also in the planted case if the mapping satisfies a certain *perturbation invariance*. We then describe our specific mappings between Erdős-Rényi and RGT. After that, in the following section, we give the main ideas for showing that our mappings satisfy the conditions of our general perturbation theorem, which constitutes the bulk of the paper's technical innovation.

2.1 Reductions and Sensitivity to Perturbation

We introduce a novel and general framework for validity of a reduction from one general graph model to another. The general result will be used to show that planted dense subgraph is computationally equivalent in Erdős-Rényi and in RGT. In this section, the graph models and transformations are abstract and specific constructions will be discussed in Section 2.2.

Let P_0 and P_0' be two distributions of random graphs on vertex set $[n]$ and consider the general planted signal $\text{GPS}_{\vec{e}, \vec{p}}$ on both random graphs.

Suppose that \mathcal{A} is a randomized algorithm (equivalently, a Markov kernel) satisfying $\mathcal{A}(P_0) = P_0'$. The following theorem gives a general sufficient condition for \mathcal{A} to approximately correctly map the *planted distribution* $P_1 := \text{GPS}_{\vec{e}, \vec{p}}(P_0)$ to the planted distribution $P_1' := \text{GPS}_{\vec{e}, \vec{q}}(P_0')$. For each $0 \leq i \leq K$, let P_{1i} (or P_{1i}') be the planted version of P_0 (or P_0'), defined by starting with P_0 (or P_0') and resampling edges $e_1, e_2, \dots, e_i \in \binom{[n]}{2}$, respectively, with probability p_1, p_2, \dots, p_i (or q_1, q_2, \dots, q_i).

Theorem 2.1. *Let \mathcal{G} be a set of graphs on n vertices such that for any $0 \leq i \leq K$, $\mathbb{P}_{G \sim P_{1i}}(G \in \mathcal{G}) \geq 1 - \delta$. Let \mathcal{A} be a randomized mapping between graphs on n vertices satisfying $\mathcal{A}(P_0) = P_0'$. Suppose that for each $e \in \binom{[n]}{2}$, there exists $p_-^e, p_+^e \in [0, 1]$ such that for every $G \in \mathcal{G}$:*

C1 *Presence of edge e in the input to \mathcal{A} has little influence on the other edges in the output of \mathcal{A} :*

$$\mathcal{A}(G - e)|_{\sim e} \approx_\epsilon \mathcal{A}(G + e)|_{\sim e}.$$

C2 *In the output graph, edge e is approximately independent of the rest of the edges:*

$$\mathcal{A}(G - e) \approx_\epsilon \mathcal{A}(G - e)|_{\sim e} \times \mathcal{A}(G - e)|_e \quad \text{and} \quad \mathcal{A}(G + e) \approx_\epsilon \mathcal{A}(G + e)|_{\sim e} \times \mathcal{A}(G + e)|_e.$$

C3 *The marginal probability on edge e is approximately constant as a function of the input graph:*

$$|\mathbb{P}(e \in \mathcal{A}(G - e)) - p_-^e| \leq \epsilon \quad \text{and} \quad |\mathbb{P}(e \in \mathcal{A}(G + e)) - p_+^e| \leq \epsilon.$$

Then we have

$$d_{\text{TV}}(\mathcal{A}(\text{GPS}_{\vec{e}, \vec{p}}(P_0)), \text{GPS}_{\vec{e}, \vec{q}}(P_0')) = O(K(\epsilon + \delta)),$$

where $q_i = p_i p_+^{e_i} + (1 - p_i) p_-^{e_i}$ for each $i \in [K]$.

The main technical challenge of this paper is to prove that our proposed mapping \mathcal{A} from RGT to Erdős-Rényi satisfies conditions **C1** and **C2**. The mapping is given in Section 2.2.

Remark 5. It turns out that the conditions in the theorem are equivalent (up to constant factors) to the following more compact conditions:

$$d_{\text{TV}}(\mathcal{A}(G - e), \mathcal{A}(G - e)|_{\sim e} \times \text{Bern}(p_-^e)) \leq \epsilon \quad \text{and} \quad d_{\text{TV}}(\mathcal{A}(G + e), \mathcal{A}(G + e)|_{\sim e} \times \text{Bern}(p_+^e)) \leq \epsilon.$$

We state Theorem 2.1 in the form above because it corresponds to how we apply it.

2.1.1 Proof of Theorem 2.1 via Single Edge Lemma

The idea is to resample edges of the input one at a time and bound the effect of each step on the output of \mathcal{A} . Let Res_e^p be the operation of resampling edge e to be present with probability p .

Lemma 2.1 (Single Edge Lemma). *Under the conditions of Theorem 2.1, for any $1 \leq i \leq K$,*

$$d_{\text{TV}}(\mathcal{A} \circ \text{Res}_{e_i}^{p_i}(P_{1_{i-1}}), \text{Res}_{e_i}^{q_i} \circ \mathcal{A}(P_{1_{i-1}})) = O(\epsilon + \delta).$$

Thus, the operations of applying \mathcal{A} and resampling edge e_i *approximately commute*, with p_i being replaced by q_i .

Note that $P_{1i} = \text{Res}_{e_i}^{p_i} P_{1i-1}$, so Lemma 2.1 equivalently states that

$$d_{\text{TV}}(\mathcal{A}(P_{1i}), \text{Res}_{e_i}^{q_i} \circ \mathcal{A}(P_{1i-1})) = O(\epsilon + \delta).$$

Applying the triangle inequality K times and using the last display for each term yields

$$d_{\text{TV}}(\text{Res}_{e_1}^{q_1} \circ \dots \circ \text{Res}_{e_K}^{q_K} \circ \mathcal{A}(P_0), \mathcal{A}(P_{1K})) = O(K(\epsilon + \delta)).$$

Note that $\mathcal{A}(P_0) = P_0'$, and by definition $\text{Res}_{e_1}^{q_1} \dots \text{Res}_{e_K}^{q_K} P_0' = P_{1K}'$, so Theorem 2.1 is proved. It remains to prove Lemma 2.1.

2.1.2 Proof of Single Edge Lemma 2.1

Proof. Suppose $G_{i-1} \sim P_{1i-1}$ and $G_i = \text{Res}_{e_i}^{p_i} G_{i-1}$. Let $\mathcal{A}^{e_i}(G)$ be the following transformation that decouples edge e_i with rest of the edges,

1. Sample a graph from $\mathcal{A}(G - e)$.
2. If $e \notin G$, resample edge e to be present with probability p_-^e . Otherwise, resample edge e to be present with probability p_+^e .

In other words, $\mathcal{A}^{e_i}(G) = \text{Res}_{e_i}^{p_-^{e_i} \mathbb{1}\{e \notin G\} + p_+^{e_i} \mathbb{1}\{e \in G\}} \circ \mathcal{A}(G - e)$. From C2, $\mathcal{A}(G) \approx_\epsilon \mathcal{A}(G)|_{\sim e_i} \times \mathcal{A}(G)|_{e_i}$. From C3, if $e \in G$, $\mathcal{A}(G)|_{e_i} \approx_\epsilon \text{Bern}(p_+^{e_i})$, if $e \notin G$, $\mathcal{A}(G)|_{e_i} \approx_\epsilon \text{Bern}(p_-^{e_i})$. Therefore, $d_{\text{TV}}(\mathcal{A}^{e_i}(G), \mathcal{A}(G)) \leq 2\epsilon$. If $G_{i-1} \in \mathcal{G}$, we have

$$\text{Res}_{e_i}^{q_i} \circ \mathcal{A}(G_{i-1}) \approx_{2\epsilon} \text{Res}_{e_i}^{q_i} \circ \text{Res}_{e_i}^{p_-^{e_i} \mathbb{1}\{e_i \notin G_{i-1}\} + p_+^{e_i} \mathbb{1}\{e_i \in G_{i-1}\}} \circ \mathcal{A}(G_{i-1} - e) = \text{Res}_{e_i}^{q_i} \circ \mathcal{A}(G_{i-1} - e).$$

Because $G_{i-1} \in \mathcal{G}$ with probability at least $1 - \delta$, $\text{Res}_{e_i}^{q_i} \circ \mathcal{A}(P_{1i-1}) \approx_{2\epsilon + \delta} \text{Res}_{e_i}^{q_i} \circ \mathcal{A}(P_{1i-1} - e)$.

Similarly, fixing $G_i \in \mathcal{G}$, we have $\mathcal{A}(G_i) \approx_\epsilon \mathcal{A}^{e_i}(G_i)$. Therefore,

$$\mathcal{A}(P_{1i}) \approx_{2\epsilon + \delta} \mathcal{A}^{e_i}(P_{1i}).$$

Use $P_{1i} = \text{Res}_{e_i}^{p_i}(P_{1i-1}) = \mathbb{E}_{G_{i-1}}[\text{Res}_{e_i}^{p_i} G_{i-1}]$ and the definition of \mathcal{A}^{e_i} , the left-hand side is equal to $\mathcal{A} \circ \text{Res}_{e_i}^{p_i}(P_{1i-1})$ and the right-hand side is equal to

$$\begin{aligned} & \mathbb{E}_{G_{i-1}} \left[\text{Res}_{e_i}^{p_-^{e_i} \mathbb{1}\{e \notin \text{Res}_{e_i}^{p_i}(G_{i-1})\} + p_+^{e_i} \mathbb{1}\{e \in \text{Res}_{e_i}^{p_i}(G_{i-1})\}} \circ \mathcal{A}(G_i - e) \right] \\ &= \mathbb{E}_{G_{i-1}} \left[\text{Res}_{e_i}^{p_i p_+^{e_i} + (1-p_i) p_-^{e_i}} \circ \mathcal{A}(G_{i-1} - e) \right] \\ &= \mathbb{E}_{G_{i-1}} \left[\text{Res}_{e_i}^{q_i} \circ \mathcal{A}(G_{i-1} - e) \right] \\ &= \text{Res}_{e_i}^{q_i} \circ \mathcal{A}(P_{1,i-1} - e). \end{aligned}$$

In other words, we have

$$\mathcal{A} \circ \text{Res}_{e_i}^{p_i}(P_{1,i-1}) \approx_{4\epsilon + 2\delta} \text{Res}_{e_i}^{q_i} \circ \mathcal{A}(P_{1,i-1} - e). \quad \square$$

2.2 Mapping between Erdős-Rényi and RGT

To use Theorem 2.1 to relate ordinary PDS to the version in RGT, we first need to specify transformations between the ambient graph models. Our mapping from $G(n, p)$ to $\text{RGT}(n, p, p')$ simply applies the definition of RGT.

Definition 4 (Forward Transition). Given any graph G , let $\mathcal{A}_\Delta(G)$ be the graph obtained from G by the following process: independently for each set of three vertices, add the three edges between them with probability p' . This defines a Markov transition kernel on the space of graphs.

By the definition of random graph with triangles, \mathcal{A}_Δ transfers $G(n, p)$ to $\text{RGT}(n, p, p')$. The reverse transition is more complicated. We will first describe the distribution of the set of triangles that were added to G_0 , conditioned on observing $G = \mathcal{A}_\Delta(G_0)$. Let $X \in \{0, 1\}^{\binom{[n]}{3}}$ be an indicator of a set of triangles. We use $|X|$ to denote the size of this set, $E(X)$ to denote the set of edges included in at least one of the triangles, and $e(X) = |E(X)|$ to denote the total number of edges.

Definition 5 (Triangle Distribution). For a given graph G , the *triangle distribution* μ_G is a distribution over subsets of triangles $x \in \{0, 1\}^{\binom{[n]}{3}}$. The probability mass function μ_G is given by

$$\mu_G(x) = \frac{1}{Z_G} \left(\frac{p'}{1-p'} \right)^{|x|} p^{-e(x)},$$

if $E(x) \subseteq G$ and $\mu_G(x) = 0$ if $E(x) \not\subseteq G$. Here $Z_G = \sum_{x: E(x) \subseteq G} \left(\frac{p'}{1-p'} \right)^{|x|} p^{-e(x)}$ normalizes μ_G .

It will be convenient to let $Y \in \{0, 1\}^{\binom{[n]}{2}}$ be the indicator vector of the edge set $E(X)$,

$$Y_e = \mathbf{1}\{e \in E(X)\} \quad \text{for each } e \in \binom{[n]}{2}.$$

To distinguish between different graphs, we use $\mathcal{L}_G(Y)$ to denote the law of Y for a given graph G .

Definition 6 (Reverse Transition). Given any graph G , let $G' = \mathcal{A}(G)$ be the graph obtained from G by the following process:

1. Sample a set of triangles $X \sim \mu_G$
2. Let G' be equal to G on the set $E(X)^c$
3. For each $e \in E(X)$, include e in G' independently with probability p .

Remark 6. The reverse transition is only analyzed for $p' = 1/(n \log n)$, not for any $p' \leq 1/(n \log n)$. Nonetheless, we can still construct a reverse map that works for any $p' \leq 1/(n \log n)$ by first adding triangles to increase p' to $1/(n \log n)$, and then applying \mathcal{A} . This is formalized in Corollary 7.1.

Lemma 2.2. *The reverse transition \mathcal{A} maps $\text{RGT}(n, p, p')$ to the Erdős-Rényi $G(n, p)$ distribution.*

Proof. Recall the definition of $\text{RGT}(n, p, p')$: we start with $G(n, p)$ and add triangles. Let T denote the set of triangles that are selected, and G denote the random graph after adding triangles. Let $|G|$ be the number of edges in graph G .

Let $P(G|G')$ denote the distribution on graphs G obtained by applying the forward transition on G' , i.e., the distribution of $\mathcal{A}_\Delta(G')$. We will show that the reverse transition is exactly the chain given by the posterior distribution $P(G'|G)$. To start,

$$\begin{aligned} P(G'|G) &\propto P(G|G')P(G') = \sum_X \mathbb{1}\{G = G' + E(X)\} p^{|X|} (1-p')^{\binom{n}{3}-|X|} p^{|E(G')|} (1-p)^{\binom{n}{2}-|E(G')|} \\ &\propto \sum_X \mathbb{1}\{G = G' + E(X)\} \left(\frac{p'}{1-p'}\right)^{|X|} \left(\frac{p}{1-p}\right)^{|E(G')|}. \end{aligned}$$

Now, let $Q(G'|G)$ be the distribution given by applying the reverse transition with input G , i.e., the distribution of $\mathcal{A}(G)$:

$$Q(G'|G) = \sum_X \mathbb{1}\{G = G' + E(X)\} \mu_G(X) p^{\mathbf{e}(X)-|E(G)|+|E(G')|} (1-p)^{|E(G)|-|E(G')|},$$

where $\mu_G(X) \propto \left(\frac{p'}{1-p'}\right)^{|X|} p^{-\mathbf{e}(X)}$. Dropping the factor depending only on $|G'|$ shows that $Q(G'|G) = P(G'|G)$ and proves the lemma. \square

For the purpose of reasoning about polynomial-time algorithms, it is crucial that our reduction \mathcal{A} can be implemented in polynomial time. Fortunately, producing a sample $X \sim \mu_G$ can be done efficiently via the Glauber dynamics Markov chain.

Lemma 2.3. *For any fixed graph G over n vertices, $p' \ll 1/n$ and constant p , the Glauber dynamics on μ_G mixes in $O(n^3 \log n)$ time.*

The Glauber dynamics is defined in Section 9.1 and the lemma is proved in Section 9.2.

2.3 Applying Theorem 2.1 to Triangle Removal Algorithm

Having Theorem 2.1 and the transformations in hand, we are ready to prove Theorem 1.1. In this section we focus on the reverse transition, \mathcal{A} , as it contains a wider range of technical ideas. The proof for the forward transition \mathcal{A}_Δ , provided in Section 5, is significantly easier and follows the same high-level outline.

To begin, for the unplanted case Lemma 2.2 states that $\mathcal{A}(\text{RGT}(n, p, p')) = G(n, p)$, so it remains to prove that $\mathcal{A}(\text{RGT}(n, p, p', k, q)) \approx_{o_n(1)} G(n, p, k, g(q))$. We focus here on the case $p' = \tilde{\Theta}(1/n)$.

Theorem 2.2 (Triangle Removal in Planted Case). *For $k = o(n^{1/4} \log^{-17/4} n)$, $p' = 1/(n \log n)$ and $0 < p < q < 1$ being constant,*

$$d_{\text{TV}}(\mathcal{A}(\text{RGT}(n, p, p', S, q)), G(n, p, S, q \cdot p_e)) = o_n(1),$$

where $p_e = \mathbb{E}_{G \sim \text{RGT}(n, p, p')} \mathbb{P}_{\mathcal{A}}(e \in \mathcal{A}(G + e))$ for an arbitrary edge e .

Here $\text{RGT}(n, p, p', S, q)$ stands for the random graph generated by planting a dense subgraph $G(k, q)$ at vertex set S in $\text{RGT}(n, p, p')$. Similarly, $G(n, p, S, q \cdot p_e)$ stands for the random graph generated by planting a dense subgraph $G(k, q \cdot p_e)$ at vertex set S in $G(n, p)$.

To apply our general Theorem 2.1, we need to specify two items: the intermediate planted RGT models and the class of graphs \mathcal{G} . We will then check conditions C1, C2, and C3.

2.3.1 Intermediate Planted Models

Define $\text{RGT}_i(n, p, p'S, q)$ to be a random graph generated by starting with $\text{RGT}(n, p)$ and independently resampling each edge e_1, e_2, \dots, e_i to be included with probability q .

2.3.2 Defining Class of Graphs \mathcal{G}

It is not hard to devise examples of graphs G for which each of the conditions **C1**, **C2**, **C3** are violated. We define the class of graphs \mathcal{G} to avoid these bad examples.

The class of graphs \mathcal{G} is chosen to be $\mathcal{G}_1 \cap \mathcal{G}_2$, where these are as follows:

\mathcal{G}_1 is the set of graphs that are $p^2/3$ -uniformly 2-star dense, where a graph is c -uniformly 2-star dense if for any pair of nodes i, j , there are at least $c(n-2)$ nodes k such that both (i, k) and (j, k) are in the graph; and

\mathcal{G}_2 is the set of graphs G that satisfy for every $e \in \binom{[n]}{2}$ that

$$|\mu_{G+e}(Y_e = 1) - p_e^+| = C_p n^{-5/2} p'^{-2} \sqrt{\log n}, \quad \text{where } p_e^+ = \mathbb{E}_{G' \sim \text{RGT}(n, p, p')} [\mu_{G'+e}(Y_e = 1)],$$

and C_p is a fixed constant depending on p .

The following lemma states that \mathcal{G} is a high probability set for each RGT_i as required by Theorem 2.1. It is proved in Section 7.5.

Lemma 2.4. *For any RGT_i , $\mathbb{P}_{G \sim \text{RGT}_i}(G \in \mathcal{G}) = 1 - o(1/n)$.*

2.3.3 Checking the Conditions of Theorem 2.1

We now state the main lemma needed to verify the conditions in Theorem 2.1. The lemma shows insensitivity of \mathcal{A} to perturbing the input by a single edge, and contains the bulk of our technical contributions. The key ideas will be presented in Section 2.4 with the full proof deferred to Section 7.

Lemma 2.5 (Perturbation Insensitivity). *If $G \in \mathcal{G}_1$, then*

$$\mathcal{L}_{G-e}(Y_{\sim e}) = \mathcal{L}_{G+e}(Y_{\sim e} | Y_e = 0) \approx_{O(\log^{17/2} n / \sqrt{n})} \mathcal{L}_{G+e}(Y_{\sim e} | Y_e = 1),$$

and from this it follows that

$$\mathcal{L}_{G+e}(Y) \approx_{O(\log^{17/2} n / \sqrt{n})} \mathcal{L}_{G+e}(Y_{\sim e}) \times \mathcal{L}_{G+e}(Y_e).$$

We first check that condition **C1** follows from the lemma. Since $\mathcal{L}_{G+e}(Y_{\sim e})$ is a mixture of the laws $\mathcal{L}_{G+e}(Y_{\sim e} | Y_e = 0)$ and $\mathcal{L}_{G+e}(Y_{\sim e} | Y_e = 1)$, Lemma 2.5 implies $\mathcal{L}_{G+e}(Y_{\sim e}) \approx_{\tilde{O}(1/\sqrt{n})} \mathcal{L}_{G-e}(Y_{\sim e})$. Note that \mathcal{A} simply resamples edges in Y , so by the data processing inequality for TV,

$$\mathcal{A}(G+e)|_{\sim e} \approx_{\tilde{O}(1/\sqrt{n})} \mathcal{A}(G-e)|_{\sim e}.$$

Condition **C2** for $G-e$ is trivial: $\mathcal{A}(G-e)$ has no edge in position e , so for any G ,

$$\mathcal{A}(G-e) = \mathcal{A}(G-e)|_{\sim e} \times \text{Bern}(0).$$

For $G + e$, the second display of Lemma 2.5 implies, again by data processing inequality, that

$$\mathcal{A}(G + e) \approx_{\tilde{O}(1/\sqrt{n})} \mathcal{A}(G + e)|_{\sim e} \times \mathcal{A}(G + e)|_e.$$

Lastly, condition C3 is immediate for $G - e$, since $\mathcal{A}(G - e)|_e = \text{Bern}(0)$ for all G . For $G + e$, it follows from the definition of \mathcal{G}_2 : For any $G \in \mathcal{G}_2$, we have by conditioning on the value of Y_e that

$$\begin{aligned} \mathcal{A}(G + e)|_e &\sim \mu_{G+e}(Y_e = 1) \cdot \text{Bern}(p) + [1 - \mu_{G+e}(Y_e = 1)] \cdot \text{Bern}(1) \\ &= \text{Bern}(1 - (1 - p)\mu_{G+e}(Y_e = 1)) \\ &\approx_{\tilde{O}(1/\sqrt{n})} \text{Bern}\left(1 - (1 - p) \cdot \mathbb{E}_{G' \sim \text{RGT}(n, p, p')} [\mu_{G'+e}(Y_e = 1)]\right). \end{aligned}$$

The last step used the fact that $d_{\text{TV}}(\text{Bern}(a), \text{Bern}(b)) = |a - b|$. Let $p_e = \mathbb{E}_{G' \sim \text{RGT}(n, p, p')} \mathbb{P}(e \in \mathcal{A}(G' + e)) = 1 - (1 - p) \cdot \mathbb{E}_{G' \sim \text{RGT}(n, p, p')} [\mu_{G'+e}(Y_e = 1)]$. Combining the last two displayed equations,

$$\mathcal{A}(G + e) \approx_{\tilde{O}(1/\sqrt{n})} \mathcal{A}(G + e)|_{\sim e} \times \text{Bern}(p_e).$$

Therefore, the conditions of Theorem 2.1 hold with $q' = q \cdot p_e$, proving Theorem 2.2.

2.4 Showing Perturbation Insensitivity: Main Ideas Behind Lemma 2.5

Let us examine more closely the generation of variables in Lemma 2.5. We fix G and let $X \sim \mu_{G+e}(\cdot | Y_e = 0) = \mu_{G-e}$ and $X^+ \sim \mu_{G+e}(\cdot | Y_e = 1)$, where recall that Y and Y^+ are the corresponding edge indicator vectors as defined in Section 2.2. Without loss of generality, assume $e \notin G$, so $G - e = G$. Our objective is to show that $Y_{\sim e}$ and $Y_{\sim e}^+$ are close in total variation.

Let $X_{T(e)}$, $X_{T(e)}^+$, and $X_{\sim T(e)}$, $X_{\sim T(e)}^+$ be the triangle indicator vectors restricted to the set of triangles that contain and do not contain e , respectively. The proof of Lemma 2.5 is divided into two conceptual parts: (1) insensitivity to perturbing Y_e of triangles $X_{\sim T(e)}$ that do not include e , and (2) conditioning on $X_{\sim T(e)}$, addressing the difference of edges brought about by triangles $X_{T(e)}$ that include e . We decompose $d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^+)$ into two terms (using Lemma 4.3 in Section 4.3):

$$d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^+) \leq d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^+) + d_{\text{TV}}(Y_{\sim e}, \mathbb{E}_{X' \sim X_{\sim T(e)}} \mathcal{L}(Y_{\sim e}^+ | X_{\sim T(e)}^+ = X')). \quad (1)$$

We next describe how to bound each of the two terms on the right-hand side.

2.4.1 Edge e Has Low Influence on Non-Containing Triangles

We first bound $d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^+)$. Only in Section 2.4.1 we will pretend that $X^+ \sim \mu_{G+e}$ rather than $\mu_{G+e}(\cdot | Y_e = 1)$, and this turns out to be valid as the two distributions have roughly the same variation distance to $X_{\sim T(e)}$ because μ_{G+e} is a constant-weight mixture of $\mu_{G+e}(\cdot | Y_e = 1)$ and μ_{G-e} .

Note that because $e \notin G$, $X_{T(e)}$ is always 0, so $X_{\sim T(e)}$ has the same distribution as X , defined by μ_G . As for $X_{\sim T(e)}^+$, letting T_e be the set of triangles containing e ,

$$\mathcal{L}(X_{\sim T(e)}^+) = \sum_{x \in \{0,1\}^{T_e}} \mathbb{P}(X_{T(e)}^+ = x) \mathcal{L}(X_{\sim T(e)}^+ | X_{T(e)}^+ = x)$$

is a mixture of Gibbs distributions indexed by the value of $X_{T(e)}^+$. By Pinsker's inequality we have $d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^+)^2 \leq \chi^2(X_{\sim T(e)} \| X_{\sim T(e)}^+)$, and we bound the latter quantity via Ingster's 2nd moment method [IS03]. We emphasize that both $\mathcal{L}(X_{\sim T(e)})$ and $\mathcal{L}(X_{\sim T(e)}^+)$ are complicated dependent distributions, while all prior works to the best of our knowledge have always shown bounds between mixtures of *product distributions* (see, for instance, [IS03, BR13b, HWX15, BBN19]). We show that despite this dependence it is still possible to derive a tractable bound.

Lemma 2.6 (χ^2 -divergence for mixture of Gibbs measures). *Let P be a distribution defined by*

$$P(X) = f(X)/Z.$$

Let U be a discrete random variable and Q be a mixture of Gibbs distributions defined by

$$Q(X) = \mathbb{E}_U[f_U(X)/Z_U].$$

Letting $\rho_U(X) = f_U(X)/f(X)$, we have that

$$\chi^2(Q \| P) + 1 = \mathbb{E}_{U, U'} \frac{\mathbb{E}_{X \sim P} \rho_U(X) \rho_{U'}(X)}{\mathbb{E}_{X \sim P} \rho_U(X) \mathbb{E}_{X \sim P} \rho_{U'}(X)},$$

where U' is an independent and identically distributed copy of U .

It turns out that when applying Lemma 2.6 to $X_{\sim T(e)}$ versus $X_{\sim T(e)}^+$, we can bound the χ^2 -divergence via marginal influences of the edge distribution $\mathcal{L}(Y)$. Here the marginal influence $\mathbf{I}_{A' \rightarrow e}^M$ characterizes how much configurations on A' can affect the conditional marginal probability on edge e . A formal definition of marginal influence can be found in Definition 9.

Lemma 2.7. *Suppose $G \in \mathcal{G}$ and consider $\mathcal{L}_G(Y)$ for Y as in Section 2.2 and $p' = o(1/n)$. Let $\mathbf{I}_{A' \rightarrow e}^M$ be the marginal influence of A' on e for $\mathcal{L}_G(Y)$. Letting $\tilde{X}_{T(e)}^+$ be an independent copy of $X_{T(e)}^+$, we have*

$$\chi^2(X_{\sim T(e)}^+ \| X_{\sim T(e)}) + 1 \leq \mathbb{E}_{X_{T(e)}^+, \tilde{X}_{T(e)}^+} \min \left\{ \left(1 + \frac{1-p}{p} \sup_{\substack{A' \subset A \cup B \\ e \in (A \cup B) \setminus A'}} \mathbf{I}_{A' \rightarrow e}^M \right)^{|B|} (1/p)^{|A \cap B|}, (1/p)^{|A|+|B|} \right\},$$

where $A = E(\tilde{X}_{T(e)}^+) - e$ and $B = E(X_{T(e)}^+) - e$.

The marginal influence $\mathbf{I}_{A' \rightarrow e}^M$ can be bounded by exploiting the fact that μ_G has small marginal probabilities under arbitrary conditioning, and we prove a general bound to this effect in Section 8.

Lemma 2.8. *Suppose $G \in \mathcal{G}$ and consider $\mathcal{L}_G(Y)$ for Y as in Section 2.2 and $p' = o(1/n)$. Let $\mathbf{I}_{A \rightarrow e}^M$ be the marginal influence of A on e for $\mathcal{L}_G(Y)$. For an edge set $A \subset E(G)$ with $|A| = O(n)$,*

$$\mathbf{I}_{A \rightarrow e}^M = \tilde{O}(|A|/n).$$

With high probability, $|A|$ and $|B|$ each have size $\tilde{O}(1)$ and $|A \cap B| = 0$, so the right hand of Lemma 2.7 is $1 + \tilde{O}(1/n)$ with high probability. Of course, the tail distribution of $|A|$ and $|B|$ is important, and we will make the bound rigorous in Section 7 to get that $\chi^2(X_{\sim T(e)} \| X_{\sim T(e)}^+) = \tilde{O}(1/n)$. By Pinsker's inequality, this shows $d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^+) = \tilde{O}(1/\sqrt{n})$.

2.4.2 Triangles Including e and TV Between Projections of Distributions

Recalling the decomposition (1), we now aim to bound the second term,

$$d_{\text{TV}}(Y_{\sim e}, \mathbb{E}_{X' \sim X_{\sim T(e)}} \mathcal{L}(Y_{\sim e}^+ | X_{\sim T(e)}^+ = X')) . \quad (2)$$

This can be rewritten as the more symmetric expression

$$d_{\text{TV}}\left(\mathbb{E}_{X' \sim X_{\sim T(e)}} \mathcal{L}(Y_{\sim e} | X_{\sim T(e)} = X'), \mathbb{E}_{X' \sim X_{\sim T(e)}} \mathcal{L}(Y_{\sim e}^+ | X_{\sim T(e)}^+ = X')\right) .$$

Since $Y_{\sim e}|X$ and $Y_{\sim e}^+|X^+$ are projections of triangle variables X and X^+ onto the edge space, the most natural approach to establish their closeness is to show that $\mathcal{L}(X_{T(e)}|X_{\sim T(e)})$ and $\mathcal{L}(X_{T(e)}^+|X_{\sim T(e)}^+)$ are close and appeal to data processing inequality. However, this is not true: As mentioned earlier, $X_{T(e)} = \vec{0}$ since $e \notin G$, and in contrast $X_{T(e)}^+$ is non-zero with non-negligible probability, since there are $\Theta(n)$ triangles containing to e and each is selected with probability approximately $p' = \tilde{\Theta}(1/n)$.

Nevertheless, the fact that we are proving a statement about projection onto the edge space, $Y_{\sim e}$ and $Y_{\sim e}^+$, allows us to carry out manipulations in the triangle space before projection. We will design an auxiliary distribution X^{aux} over the same support as $X_{\sim T(e)}$ that when added to $X_{\sim T(e)}^+$ results in the *identical* edge projection as X_e^+ , i.e.,

$$E(X^{\text{aux}} \vee X_{\sim T(e)}^+) - e = E(X_{T(e)}^+ \vee X_{\sim T(e)}^+) - e .$$

This means the edge indicator vector of $X^{\text{aux}} \vee X_{\sim T(e)}^+$, which we denote by $\tilde{Y}_{\sim e}$, is the same as $Y_{\sim e}^+$. Thus, instead of comparing $Y_{\sim e}$ and $Y_{\sim e}^+$ in (2), it suffices to compare $Y_{\sim e}$ and $\tilde{Y}_{\sim e}$, and by data processing inequality it in turn suffices to compare $X_{\sim T(e)}$ and $X_{\sim T(e)}^+ \vee X^{\text{aux}}$ (see Fig. 1):

$$d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^+) \leq d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^+ \vee X^{\text{aux}}) \quad (3)$$

We now define X^{aux} . Each triangle in $X_{T(e)}^+$ adds at most two edges to $Y_{\sim e}^+$. To simulate this change without using any triangles containing e , we add a triangle to X^{aux} for each new edge introduced by $X_{T(e)}^+$. Crucially, this is done without adding other new edges, see Fig. 2. Avoiding adding new edges is possible with high probability as long as the graph G is sufficiently dense (here c -uniformly 2-star dense plays a role) and $p' = \tilde{\Theta}(1/n)$, which implies that the edge set of $X_{\sim T(e)}^+$ has sufficient coverage of relevant triangles that we might potentially add to X^{aux} .

From the previous section, we have $X_{\sim T(e)} \approx_{\tilde{O}(1/\sqrt{n})} X_{\sim T(e)}^+$, so to show (3) it remains to prove that $X_{\sim T(e)} \approx X_{\sim T(e)} \vee X^{\text{aux}}$, i.e., the addition of X^{aux} must be undetectable. We show this via concentration of the likelihood ratio between $X_{\sim T(e)} \vee X^{\text{aux}}$ and $X_{\sim T(e)}$. In Section 8, we prove that Lipschitz functions over μ_G concentrate. The likelihood ratio under consideration is not quite Lipschitz – it is Lipschitz only over a high probability subset – but this turns out to suffice for it to concentrate.

3 Strong Equivalence Between Parameterized Problems

Each problem we consider is actually an entire *family* of problems. Our results turn out to imply a strong form of computational equivalence between families of problems. This form of equivalence is new, so we carefully introduce the relevant definitions.




Figure 1: Relation between variables. By the data processing inequality, which implies that total variation shrinks under projection, closeness of (1) implies closeness of (2).




Figure 2: Design of X^{aux} . For each edge in $E(X_{T(e)}^+) - E(X_{\sim T(e)}^+) - e$ (purple edges), X^{aux} includes a triangle (in red) containing the edge, without adding other edges not already in $E(X_{\sim T(e)}^+)$ (in green).

Average-Case Reduction in Total Variation Let $\mathcal{P} = \{(P_{0,n}, P_{1,n}), n \in \mathbb{N}\}$ be a hypothesis testing problem between

$$H_0 : X \sim P_{0,n} \quad \text{and} \quad H_1 : X \sim P_{1,n},$$

and $\mathcal{P}' = \{(P'_{0,n}, P'_{1,n}), n \in \mathbb{N}\}$ be another hypothesis testing problem between

$$H_0 : X \sim P'_{0,n} \quad \text{and} \quad H_1 : X \sim P'_{1,n}.$$

An *average-case reduction* from \mathcal{P} to \mathcal{P}' is a randomized mapping $\mathcal{A} : \Omega \rightarrow \Omega$, computable in polynomial time, such that $d_{\text{TV}}(\mathcal{A}(P_{0,n}), P'_{0,n}) = o_n(1)$ and $d_{\text{TV}}(\mathcal{A}(P_{1,n}), P'_{1,n}) = o_n(1)$.²

It is clear that if there exists an average case reduction \mathcal{A} from \mathcal{P} to \mathcal{P}' , then any polynomial-time algorithm Φ that has vanishing error on \mathcal{P}' yields a polynomial-time algorithm $\Phi \circ \mathcal{A}$ with vanishing error on \mathcal{P} . Therefore, if \mathcal{P} is computationally hard (or information-theoretically impossible), then \mathcal{P}' is also computationally hard (or information-theoretically impossible).

Parameterized Problems and Reductions Most statistical problems have several parameters. In our case of planted RGT, for a given graph size n we have parameters for the dense-subgraph size k , ambient edge probability p , triangle probability p' , and edge density q in the dense subgraph. Each tuple of parameters describes a distinct problem instance from within the larger family, and a reduction maps between a problem at a specific parameter to another problem at some parameter. We use $\mathcal{P}(\vec{\alpha}_n) = \{(P_{0,n}(\vec{\alpha}_n), P_{1,n}(\vec{\alpha}_n)), \vec{\alpha}_n \in S, n \in \mathbb{N}\}$ to denote a family of parameterized problems. Here $S \subset \mathbb{R}^a$ is assumed to be compact, which can be achieved for all problems we know of by using an appropriate parameterization. For instance, for PDS with constant p and q , one can take $\vec{\alpha}_n = (\log_n k, p, q)$. For the regime $p = cq$ and $q = n^{-b}$ for $b > 0$ as studied in [HWX15] one can take $\vec{\alpha}_n = (\log_n k, \log_n(1/q), c)$.

Phase Diagram In the literature on the computational complexity of statistics problems, feasibility of solving a problem is usually described in the language of phase diagrams (see, e.g., [MW15, HWX15, BBH18, SW22]). These diagrams divide the parameter space into regions corresponding to efficiently solvable, computationally hard, or information-theoretically impossible. In order to describe the relation between families of problems, we need to introduce a precise definition that captures the meaning of existing phase diagrams.

For a parameterized family of hypothesis testing problems $\mathcal{P} = \{(P_{0,n}(\vec{\alpha}_n), P_{1,n}(\vec{\alpha}_n)), \vec{\alpha}_n \in S, n \in \mathbb{N}\}$, the *Phase Diagram* is a mapping $\text{PD}_{\mathcal{P}} : S \rightarrow \{\text{easy}, \text{hard}, \text{impossible}, \perp\}$:

- $\text{PD}_{\mathcal{P}}(\vec{\alpha}) = \text{easy}$ if and only if for any sequence $\vec{\alpha}_n$ converging to $\vec{\alpha}$, $\mathcal{P}(\vec{\alpha}_n)$ can be solved efficiently with vanishing error as $n \rightarrow \infty$.
- $\text{PD}_{\mathcal{P}}(\vec{\alpha}) = \text{hard}$ if and only if for any sequence $\vec{\alpha}_n$ converging to $\vec{\alpha}$, $\mathcal{P}(\vec{\alpha}_n)$ can only be solved by super-polynomial algorithms.
- $\text{PD}_{\mathcal{P}}(\vec{\alpha}) = \text{impossible}$ if and only if for any sequence $\vec{\alpha}_n$ converging to $\vec{\alpha}$, solving $\mathcal{P}(\vec{\alpha}_n)$ is information-theoretically impossible.

²This notion of reduction is based on *Le Cam deficiency* [LC12] and has been (sometimes implicitly) used in the recent line of work on average-case reductions for statistical problems, e.g., [BR13a, MW15, HWX15, CLM16, BBH18, BB19, BB20]. Levin's [Lev86] notion of average-case reduction is similar but tailored for dist-NP completeness results, and is more stringent in the error probability and less stringent in the distance between distributions.

- $\text{PD}_{\mathcal{P}}(\vec{\alpha}) = \perp$, otherwise.

Remark 7. It is important to note that the phase diagram of a problem depends on the parameterization: Some parameterizations lead to meaningless phase diagrams.

For example, for the PDS problem with constants $p \neq q$, if the parameters are chosen to be $\vec{\alpha}_n = (k/n, p, q)$, then the phase diagram becomes trivial. For any $0 < c \leq 1$, $\text{PD}_{\mathcal{P}}(c, p, q) = \text{easy}$ as it suffices to count the number of edges and tell the difference between the two hypotheses. And $\text{PD}_{\mathcal{P}}(0, p, q) = \perp$, because we know the problem is information-theoretically impossible for constant k and easy for $k = \Theta(\sqrt{n})$, and k/n converges to 0 in both regimes. On the other hand, if we use the parameterization $\vec{\alpha}_n = (\log_n k, p, q)$, the conjectured non-trivial phase transition will appear when $\log_n k = 1/2$. If we believe the computational threshold happens at $k = \tilde{\Theta}(\sqrt{n})$, then

$$\text{PD}_{\mathcal{P}}(c, p, q) = \begin{cases} \text{hard} & \text{if } 0 < c < 1/2 \\ \text{easy} & \text{if } 1/2 < c \leq 1 \\ \perp & \text{if } c = 1/2. \end{cases}$$

One can also study the more refined behavior near the phase transition by introducing extra parameters. For instance, in the case of PDS, one can use the parameterization $\vec{\alpha}_n = (a, b, p, q)$ where $k = n^a \log^b n$. This allows the study of more fine-grained phase transition at $k = \sqrt{n} \log^b n$.

Remark 8. Although $\vec{\alpha}_n$ belongs to S for every n , it is crucial that S remains a compact set, independent of n . In the case of PDS, the size of the planted subgraph, k , cannot be selected directly as a parameter, as it would diverge with n . To ensure that the asymptotic behavior of the problem is represented by a single point on the phase diagram, $\vec{\alpha}_n$ should remain constant or converge to a constant as n increases.

Strong Equivalence Between Problems To describe the relation between the phase diagrams of two parameterized families of problems, $\mathcal{P}(\vec{\alpha}_n)$ where $\vec{\alpha}_n \in S$ and $\mathcal{P}'(\vec{\beta}_n)$ where $\vec{\beta}_n \in S'$, we say that $\vec{\alpha} \in S$ maps to $\vec{\beta} \in S'$ if

- for any sequence $\vec{\alpha}_n$ converging to $\vec{\alpha}$, there exists a sequence $\vec{\beta}_n$ converging to $\vec{\beta}$ with an average case reduction from $\mathcal{P}(\vec{\alpha}_n)$ to $\mathcal{P}'(\vec{\beta}_n)$, and
- for any sequence $\vec{\beta}_n$ converging to $\vec{\beta}$, there exists a sequence $\vec{\alpha}_n$ converging to $\vec{\alpha}$ with an average case reduction from $\mathcal{P}(\vec{\alpha}_n)$ to $\mathcal{P}'(\vec{\beta}_n)$.

This definition yields the following correspondence of phase diagrams.

Lemma 3.1. *Let $\mathcal{P}(\vec{\alpha}_n)$, where $\vec{\alpha}_n \in S$, be a parameterized family of problems with phase diagram $\text{PD}_{\mathcal{P}}$. And $\mathcal{P}'(\vec{\beta}_n)$, where $\vec{\beta}_n \in S'$, is a parameterized family of problems with phase diagram $\text{PD}_{\mathcal{P}'}$. If $\vec{\alpha} \in S$ maps to and can be mapped from $\vec{\beta} \in S'$, then*

$$\text{PD}_{\mathcal{P}}(\vec{\alpha}) = \text{PD}_{\mathcal{P}'}(\vec{\beta}).$$

Proof. In this proof, we use the convention that information-theoretically impossible is harder than computationally hard, which is harder than efficiently solvable. If two problem are in the same category, they are considered to have same hardness.

Let $\mathcal{P}(\vec{\alpha}_n^h)$ (or $\mathcal{P}(\vec{\alpha}_n^e)$) be one of the hardest (or easiest) problems where $\vec{\alpha}_n$ converges to $\vec{\alpha}$. $\mathcal{P}'(\vec{\beta}_n^h)$ and $\mathcal{P}'(\vec{\beta}_n^e)$ are defined in the same way. Since $\vec{\alpha}$ maps to $\vec{\beta}$, there exists a problem $\mathcal{P}'(\vec{\beta}_n)$

that can be reduced from $\mathcal{P}(\vec{\alpha}_n^h)$ with $\vec{\beta}_n$ converging to $\vec{\beta}$. So $\mathcal{P}'(\vec{\beta}_n^h)$ is at least as hard as $\mathcal{P}'(\vec{\beta}_n)$, and therefore, at least as hard as $\mathcal{P}(\vec{\alpha}_n^h)$. Similarly, we can conclude that $\mathcal{P}'(\vec{\beta}_n^h)$ is at least as hard as $\mathcal{P}(\vec{\alpha}_n^h)$, which means they have the same hardness.

On the other hand, since $\vec{\alpha}$ maps to $\vec{\beta}$, there exists a problem $\mathcal{P}(\vec{\alpha}_n)$ where $\vec{\alpha}_n$ converges to $\vec{\alpha}$ that $\mathcal{P}(\vec{\alpha}_n)$ reduces to $\mathcal{P}'(\vec{\beta}_n^e)$. So $\mathcal{P}'(\vec{\beta}_n^e)$ is at least as hard as $\mathcal{P}(\vec{\alpha}_n)$, and therefore, at least as hard as $\mathcal{P}(\vec{\alpha}_n^e)$. We can also conclude that $\mathcal{P}(\vec{\alpha}_n^e)$ and $\mathcal{P}'(\vec{\beta}_n^e)$ have the same hardness. It is not hard to see from the definition of PD that $\text{PD}_{\mathcal{P}}(\vec{\alpha}) = \text{PD}_{\mathcal{P}'}(\vec{\beta})$. \square

We say two families of problems \mathcal{P} and \mathcal{P}' are *strongly equivalent* if for any point $\vec{\alpha} \in S$, there exists $\vec{\beta} \in S'$ such that $\vec{\alpha}$ maps to and can be mapped from $\vec{\beta}$, and similarly, for any point $\vec{\beta} \in S'$, there exists $\vec{\alpha} \in S$ such that $\vec{\beta}$ maps to and can be mapped from $\vec{\alpha}$.

Fact 1. *If two problems are strongly equivalent, then the phase diagram of either one of the problems determines the phase diagram of the other problem.*

Remark 9. The phase diagrams of the same problem under different parameterizations may not be strongly equivalent. One such example was discussed in Remark 7.

The main result of the paper, Theorem 1.1, can be re-stated as the following.

Theorem 3.1. *Consider the PDS problem, i.e., hypothesis testing between $G(n, p)$ and $G(n, p, k, q)$, parameterized by $\vec{\alpha}_n = (\log_n k, p, q)$, restricted to the parameter region $\log_n k < 1/4$, $0 < p, q < 1$. The Planted RGT problem has hypotheses $\text{RGT}(n, p, p')$ and $\text{RGT}(n, p, p', k, q)$ and is parameterized by $\vec{\beta} = (p, \log_n p', \log_n k, q)$ where the parameter region is $\log_n p' < -1$, $\log_n k < 1/4$, $0 < p, q < 1$. The PDS and Planted RGT problems are strongly equivalent in the given parameter regimes.*

4 Preliminaries

4.1 Probability Notation

The probability law of a random variable X is denoted by $\mathcal{L}(X)$. It will often be necessary to condition on a graph G for a jointly distributed X and G , and we use the shorthand $\mathcal{L}_G(X)$ in place of $\mathcal{L}(X|G)$.

Any randomized algorithm \mathcal{A} defines a Markov kernel where the distribution of the output for given input x follows the law of $\mathcal{A}(x)$. When an algorithm \mathcal{A} is applied to a distribution μ , $\mathcal{A}(\mu)$ stands for the law of $\mathcal{A}(X)$ where $X \sim \mu$. For random variables X, Y we use $d_{\text{TV}}(X, Y)$ in place of $d_{\text{TV}}(\mathcal{L}(X), \mathcal{L}(Y))$. For the ease of presentation, we use $d_{\text{TV}}(P, Q) \leq \epsilon$ and $P \approx_\epsilon Q$ interchangeably.

For two vectors X and Y , $X \vee Y$ is the coordinate-wise max.

We say that an event occurs with high probability if its probability is at least $1 - \frac{1}{n^c}$ for a large enough constant c , i.e., large enough for subsequent use (which in our case entails losing some numerical constant in c and still having the probability tend to one).

4.2 Graph Theory Notation

G is always used to denote a graph. In general we call a pair of vertices $e = \{i, j\} \in \binom{[n]}{2} = K_n$ an *edge*, regardless of whether e is actually in the particular graph G under consideration. If it is, then we say that edge e is present, or included, in G . A graph G is considered a subset of edges in K_n ,

or a set of indicators in $\{0, 1\}^{K_n}$ for whether an edge is included in G . We also use the convention that for any $E \subset K_n$, $G|_E$ stands for the intersection $G \cap E$.

Let $N_k(v)$ stands for the set all vertices in G with distance at most k to v without v itself. Two triangles in the graph are considered adjacent or having distance 1 if they share an edge. The distance between two triangles are defined by the length of a shortest path between them. Similarly, two edges in a graph are adjacent or having distance 1 if they share a node. The distance between two edges are defined by the length of a shortest path between them.

For a set of triangles T (as a binary vector with 1 indicating that a triangle is present), $E(T)$ denotes the edge set of the union of the triangles in T .

For a graph G and a pair of vertices u, v , the *set of wedge vertices* $W_G(u, v)$ is the set of vertices that are adjacent to both u and v (excluding u and v themselves):

$$W_G(u, v) = \{w \in [n] | (w, u), (w, v) \in G\}. \quad (4)$$

4.3 Total Variation Distance

The *total variation distance* between two distributions P and Q on the same space Ω is defined as $d_{\text{TV}}(P, Q) = \sup_{A \subseteq \Omega} P(A) - Q(A)$. The total variation is convex as function of the pair of distributions (P, Q) . We state several lemmas that will be used to bound total variation distance.

Lemma 4.1 (Data Processing Inequality (DPI)). *For any two random variables X, Y on space Ω and any randomized algorithm (Markov transition kernel) $R : \Omega \rightarrow \Omega'$ for arbitrary space Ω' , we have*

$$d_{\text{TV}}(R(X), R(Y)) \leq d_{\text{TV}}(X, Y).$$

Lemma 4.2 (Convexity and Mixtures). *Let P and Q be two distributions over the same space Ω . If there is a joint distribution (X, X', Y) where the marginal distributions of X and X' are P and Q , respectively, we have*

$$d_{\text{TV}}(P, Q) \leq \mathbb{E}_Y d_{\text{TV}}(P_{X|Y}, Q_{X'|Y}).$$

Proof. Since $\mathbb{E}_Y P_{X|Y} = P$ and $\mathbb{E}_Y P_{X'|Y} = Q$, the statement follows from convexity of d_{TV} . \square

Lemma 4.3 (TV Decomposition). *Consider two pairs of random variables, (X, Y) and (X', Y') . We have*

$$d_{\text{TV}}(Y, Y') \leq d_{\text{TV}}(X, X') + d_{\text{TV}}\left(Y, \mathbb{E}_{Z \sim \mathcal{L}(X)} \mathcal{L}(Y' | X' = Z)\right).$$

Proof. Consider an intermediate distribution $\mathbb{E}_{Z \sim \mathcal{L}(X)} \mathcal{L}(Y' | X' = Z)$. By the triangle inequality,

$$\begin{aligned} d_{\text{TV}}(Y, Y') &\leq d_{\text{TV}}\left(Y, \mathbb{E}_{Z \sim \mathcal{L}(X)} \mathcal{L}(Y' | X' = Z)\right) \\ &\quad + d_{\text{TV}}\left(\mathbb{E}_{Z \sim \mathcal{L}(X')} \mathcal{L}(Y' | X' = Z), \mathbb{E}_{Z \sim \mathcal{L}(X)} \mathcal{L}(Y' | X' = Z)\right). \end{aligned}$$

By the data processing inequality, the second term is bounded by $d_{\text{TV}}(X, X')$. \square

Lemma 4.4. *Let P and Q be two distributions over the same space Ω , and suppose that Q is absolutely continuous with respect to P . For any event $A \subseteq \Omega$,*

$$d_{\text{TV}}(P, Q) \leq 2(P(A) + Q(A)) + (1 - Q(A))d_{\text{TV}}(P|_{A^c}, Q|_{A^c}).$$

Proof. Several applications of the triangle inequality yield

$$\begin{aligned}
d_{\text{TV}}(P, Q) &= \mathbb{E}_P \left| \frac{dQ}{dP} - 1 \right| \\
&= \mathbb{E}_P \left[\mathbb{1}_A \left| \frac{dQ}{dP} - 1 \right| \right] + \mathbb{E}_P \left[\mathbb{1}_{A^c} \left| \frac{dQ}{dP} - 1 \right| \right] \\
&= \mathbb{E}_P \left[\mathbb{1}_A \left| \frac{dQ}{dP} - 1 \right| \right] + \mathbb{E}_P \left[\mathbb{1}_{A^c} \left| \frac{Q(A^c)}{P(A^c)} \cdot \frac{dQ|_{A^c}}{dP|_{A^c}} - 1 \right| \right] \\
&\leq Q(A) + P(A) + \mathbb{E}_P \left[\mathbb{1}_{A^c} \frac{Q(A^c)}{P(A^c)} \cdot \left| \frac{dQ|_{A^c}}{dP|_{A^c}} - 1 \right| + \mathbb{1}_{A^c} \left| \frac{Q(A^c)}{P(A^c)} - 1 \right| \right] \\
&\leq Q(A) + P(A) + Q(A^c) \cdot \mathbb{E}_{P|_{A^c}} \left| \frac{dQ|_{A^c}}{dP|_{A^c}} - 1 \right| + |P(A) - Q(A)| \\
&\leq 2(P(A) + Q(A)) + Q(A^c) \cdot d_{\text{TV}}(P|_{A^c}, Q|_{A^c}) \quad \square
\end{aligned}$$

Lemma 4.5 (Bounding TV via Concentration of Likelihood Ratio). *Let P and Q be two distributions over the same space Ω , with Q absolutely continuous with respect to P . Let $\frac{dQ}{dP}$ be the likelihood ratio (or Radon-Nikodym derivative) between the two distributions. If $\left| \frac{dQ}{dP}(X) - 1 \right| \leq \epsilon$ with probability at least $1 - \delta$ for $X \sim P$, then*

$$d_{\text{TV}}(P, Q) \leq 2\epsilon + 2\delta.$$

Proof. Let A be the event that $\left| \frac{dQ}{dP}(X) - 1 \right| > \epsilon$, so $P(A) \leq \delta$ from the assumption. Since $\frac{dQ}{dP}(X) \geq 1 - \epsilon$ on A^c , $Q(A^c) \geq (1 - \epsilon)P(A^c) \geq 1 - \epsilon - \delta$. So $Q(A) \leq \epsilon + \delta$. We have

$$\begin{aligned}
d_{\text{TV}}(P, Q) &= \mathbb{E} \left| \frac{dQ}{dP}(X) - 1 \right| = \mathbb{E} \left[\mathbb{1}_A \left| \frac{dQ}{dP}(X) - 1 \right| \right] + \mathbb{E} \left[\mathbb{1}_{A^c} \left| \frac{dQ}{dP}(X) - 1 \right| \right] \\
&\leq P(A) + Q(A) + \epsilon \leq 2\epsilon + 2\delta. \quad \square
\end{aligned}$$

Lemma 4.6 (Many Steps of Markov Kernel in Terms of One Step). *Let X be a random variable over space Ω , K be a Markov kernel from Ω to Ω . Then for any $m \in \mathbb{Z}^+$,*

$$d_{\text{TV}}(X, K^m(X)) \leq m \cdot d_{\text{TV}}(X, K(X)).$$

Proof. We do an induction on m . The case of $m = 1$ is trivial. For $m > 1$, by triangle inequality,

$$\begin{aligned}
d_{\text{TV}}(X, K^m(X)) &\leq d_{\text{TV}}(X, K^{m-1}(X)) + d_{\text{TV}}(K^{m-1}(X), K^m(X)) \\
&\leq (m-1)d_{\text{TV}}(X, K(X)) + d_{\text{TV}}(K^{m-1}(X), K^m(X)).
\end{aligned}$$

From data processing inequality, the second term can also be bounded by

$$d_{\text{TV}}(K^{m-1}(X), K^m(X)) \leq d_{\text{TV}}(X, K(X)). \quad \square$$

Lemma 4.7. *Let X be a real-valued random variable satisfying $0 \leq X \leq M$ almost surely. If there exists $a \in \mathbb{R}$ that $|X - a| \leq \sigma$ with probability at least $1 - \delta$, then*

$$|X - \mathbb{E} X| \leq 2\sigma + \delta M$$

with probability at least $1 - \delta$.

Proof. We have

$$a - \sigma \leq \mathbb{E} X \leq (a + \sigma)(1 - \delta) + \delta M \leq a + \sigma + \delta M.$$

Combining this with $|X - a| \leq \sigma$, we get $-2\sigma - \delta M \leq X - \mathbb{E} X \leq 2\sigma$. \square

4.4 Chernoff Bound and Mcdiarmid's Inequality

Lemma 4.8 (Chernoff Bound for Independent Binary Variables). *Let X_1, \dots, X_n be random binary variables. Let X denote their sum and $\mu = \mathbb{E}[X]$. We have for any $\delta \geq 0$,*

$$\mathbb{P}(X \geq (1 + \delta)\mu) \leq e^{-\delta^2\mu/(2+\delta)} \quad \text{and} \quad \mathbb{P}(X \leq (1 - \delta)\mu) \leq e^{-\delta^2\mu/2}.$$

Lemma 4.9 (Mcdiarmid's Inequality [McD98]). *Consider a real-valued martingale X_0, X_1, \dots, X_n adapted to filtration \mathcal{F} . If $|X_k - X_{k-1}| \leq R$ and $\text{Var}(X_k | \mathcal{F}_{k-1}) \leq \sigma_i$ almost surely for any k , then*

$$\mathbb{P}(X_n - \mathbb{E} X_n \geq t) \leq \exp\left(\frac{-t^2/2}{\sum \sigma_i^2 + Rt/3}\right).$$

Lemma 4.10 (Alternative form of McDiarmid's Inequality). *Consider $f : \mathcal{X}^n \rightarrow \mathbb{R}$. Let $X = (X_1, \dots, X_n)$ with $X_i \sim P_i$ sampled independently, where P_i has probability mass at least p_i on one element. If there exists \mathcal{Y} such that for any $x, x' \in \mathcal{Y}$,*

$$|f(x) - f(x')| \leq \sum_{i: x_i \neq x'_i} c_i$$

and $\mathbb{P}(X \in \mathcal{Y}) = 1 - p$, then for any $t \geq p \sum_i c_i$,

$$\mathbb{P}(f(X) - \mathbb{E} f(X) > t) \leq p + \exp\left(\frac{-(t - p \sum_i c_i)^2/2}{\sigma^2 + \max_i c_i t/3}\right),$$

where $\sigma^2 = \sum_i p_i(1 - p_i)c_i^2$.

Proof. The proof is the same as [Com15] but instead of using the version of Mcdiarmid's inequality in [Com15], we use the more general version of Mcdiarmid's inequality stated in Lemma 4.9. \square

5 Forward Map Preserves Planted Signal

This section aims to show the first part of Theorem 1.1 by applying Theorem 2.1 on the forward process \mathcal{A}_Δ described in Defn. 4.

In this section only, we use $X \in \{0, 1\}^{\binom{[n]}{3}}$ to denote the indicator vector of the set of triangles added by \mathcal{A}_Δ and use $Z \in \{0, 1\}^{\binom{[n]}{2}}$ to denote the indicator for the set of edges in $\mathcal{A}_\Delta(G)$, i.e., the indicator of $G \cup E(X)$. From the definition of \mathcal{A}_Δ , $X \sim \text{Bern}(p')^{\otimes \binom{n}{3}}$.

Theorem 5.1 (\mathcal{A}_Δ Increases Triangle Density in Planted RGT). *Let $p'_1 < p'_2 \leq 1/n$, $k := |S| = O(n^{1/4} \log^{-5/4} n)$, and $0 < p, q < 1$ constant. If \mathcal{A}_Δ includes each triangle with probability $p' = \frac{p'_2 - p'_1}{1 - p'_1}$, then*

$$d_{\text{TV}}(\mathcal{A}_\Delta(\text{RGT}(n, p, p'_1, S, q)), \text{RGT}(n, p, p'_2, S, q')) = o_n(1),$$

where $q' = q + (1 - q)(1 - (1 - p')^{n-2})$.

By choosing $p'_1 = 0$ we get the reduction from planted dense subgraph to planted RGT.

Corollary 5.1. For $p' \leq 1/n$, $k := |S| = O(n^{1/4} \log^{-5/4} n)$ and p, q being constant,

$$d_{\text{TV}}(\mathcal{A}_{\Delta}(G(n, p, S, q)), \text{RGT}(n, p, p'_2, S, q')) = o_n(1),$$

where $q' = q + (1 - q)(1 - (1 - p')^{n-2})$.

Remark 10. The statement of the corollary is actually true all the way up to $k = O(\sqrt{n})$. This can be proved by calculating the χ^2 -divergence induced by adding a single triangle to an Erdős-Rényi random graph, with or without resampling edges in S . (A similar calculation appears in [BBN19].) We instead use our general result, Theorem 2.1, as this gives a unified view on how to deal with correlation in planted models.

Remark 11. In Section 1.2, Theorem 1.1 is stated without explicitly giving the mapping function $f(q)$. This corollary shows that $f(q) = q + (1 - q)(1 - (1 - p')^{n-2})$. Because we assume $p' \ll 1/n$, we have $f(q) = q + O(np') = q + o_n(1)$. The parameter mapping g for the reverse map also satisfies $g(q) = q + o_n(1)$ (see Cor. 7.1 in Section 7). So as claimed in Theorem 1.1, $f \circ g(q) = q + o_n(1)$.

Proof of Theorem 5.1. We proof follows by checking the conditions of Theorem 2.1, which we do in the next three subsections. Here the two pure-noise models are $\text{RGT}(n, p, p'_1)$ and $\text{RGT}(n, p, p'_2)$, and the transformation between them is \mathcal{A}_{Δ} . We also need to describe the partially planted version of the model, which we do similarly to Section 2.3.1. Let $e_1, e_2, \dots, e_{\binom{k}{2}}$ be an enumeration of edges with both endpoints in S . Let $\text{RGT}_i(n, p, p'_1, S, q)$ be a random graph generated by starting with $\text{RGT}(n, p, p'_1)$ and resampling edges e_1, e_2, \dots, e_i to be each included with probability q .

5.1 Class of Graphs \mathcal{G}

Define the class of graphs \mathcal{G} to be the set of graphs that are $p^2/2$ -uniformly 2-star dense.

Lemma 5.2. Assume that $k = o(\sqrt{n})$. Both $G(n, p)$ and all $\text{RGT}_i(n, p, p')$ for $1 \leq i \leq \binom{k}{2}$ are $p^2/2$ -uniformly 2-star dense with probability at least $1 - e^{-\Omega(n)}$.

Proof. Let i, j be an arbitrary pair of nodes. For any $k \notin \{i, j\}$, the probability that both (i, k) and (j, k) are in the graph is p^2 . Recalling the wedge vertices $W_G(i, j) = \{w \in [n] \mid (w, i), (w, j) \in G\}$ from Section 4.2, $\mathbb{E}|W_G(i, j)| = p^2(n - 2)$. Since the event that both (i, k) and (j, k) are in the graph is independent for each k , we can use a Chernoff bound (Lemma 4.8) to get

$$\mathbb{P}(|W_G(i, j)| < 2p^2(n - 2)/3) \leq e^{-p^2(n-2)/18}.$$

By the union bound over all pairs, we get that the probability that there is exists i, j with $|W_G(i, j)| < 2p^2(n - 2)/3$ is at most $\binom{n}{2}e^{-p^2(n-2)/18} = e^{-\Omega(n)}$. The same holds for $\text{RGT}(n, p, p')$ because the property of being c -uniformly 2-star dense is preserved under addition of edges. Finally, note that changing $o(n)$ edges preserves the property, so the conclusion holds for $\text{RGT}_i(n, p, p')$. \square

5.2 Checking Conditions in Theorem 2.1

C1 of Theorem 2.1 is immediate, as $\mathcal{A}_{\Delta}(G - e) = G - e + E(X)$ and $\mathcal{A}_{\Delta}(G + e) = G + e + E(X)$ are the same on $\binom{[n]}{2} - e$. Condition **C3** is also easy to check and holds with $\epsilon = 0$: $\mathbb{P}(e \in \mathcal{A}_{\Delta}(G - e)) = \mathbb{P}(e \in E(X)) = 1 - (1 - p')^{n-2} =: p_-^e$. Since \mathcal{A}_{Δ} only adds edges, $\mathbb{P}(e \in \mathcal{A}_{\Delta}(G + e)) = 1 =: p_+^e$.

Verifying condition **C2** is more technical. For $G+e$, $e \in \mathcal{A}_\Delta(G+e)$ always holds so $\mathcal{A}_\Delta(G+e) = \mathcal{A}_\Delta(G+e)|_{\sim e} \times \text{Bern}(1)$. For $G-e$, we will show condition **C2** assuming the following lemma. Let $Z \in \{0,1\}^{\binom{[n]}{2}}$ be the indicator vector for the set of edges in $\mathcal{A}_\Delta(G-e)$.

Lemma 5.3. *Fix $0 < c < 1$. For a c -uniformly 2-star-dense graph G with $e \notin G$,*

$$d_{\text{TV}}(\mathcal{L}(Z_{\sim e}), \mathcal{L}(Z_{\sim e}|Z_e = 0)) = O(n^{-1/2} \log^{5/2} n).$$

The proof is deferred until the next subsection, and we proceed with checking condition **C2**. Note that

$$\mathcal{L}(Z_{\sim e}) = p_-^e \mathcal{L}(Z_{\sim e}|Z_e = 1) + (1 - p_-^e) \mathcal{L}(Z_{\sim e}|Z_e = 0).$$

Manipulating the relevant p.m.f.s, it follows that

$$d_{\text{TV}}(\mathcal{L}(Z_{\sim e}), \mathcal{L}(Z_{\sim e}|Z_e = 1)) \leq \frac{1}{p_-^e} \cdot d_{\text{TV}}(\mathcal{L}(Z_{\sim e}), \mathcal{L}(Z_{\sim e}|Z_e = 0)) = \tilde{O}\left(\frac{1}{p_-^e \sqrt{n}}\right).$$

Combining with Lemma 5.3,

$$\begin{aligned} \mathcal{L}(Z) &= p_-^e \mathcal{L}(Z|Z_e = 1) + (1 - p_-^e) \mathcal{L}(Z|Z_e = 0) \\ &\approx_{O(n^{-1/2} \log^{5/2} n)} p_-^e \mathcal{L}(Z_{\sim e}) \times \text{Bern}(1) + p_-^e \mathcal{L}(Z_{\sim e}) \times \text{Bern}(0) \\ &= \mathcal{L}(Z_{\sim e}) \times \text{Bern}(p_-^e). \end{aligned}$$

This completes the proof for condition **C2**.

Therefore, Theorem 2.1 applies with $\epsilon = O(n^{-1/2} \log^{5/2} n)$ and $\delta = O(e^{-\Omega(n)})$. When $k = O(n^{1/4} \log^{-5/4} n)$, there are at most $O(n^{1/2} \log^{-5/2} n)$ planted edges and the theorem follows from Theorem 2.1 with $q' = qp_+^e + (1-q)p_-^e = q + (1-q)(1 - (1-p')^{n-2})$. \square

5.3 Perturbation Insensitivity of \mathcal{A}_Δ and Proof of Lemma 5.3

It only remains to prove the lemma. The idea of the proof is similar to Section 2.4.2. In the forward mapping, we do not need to worry about the influence of non-incident triangles in Section 2.4.1, as all the triangles are independent. This section thus serves as a warm up for the proof of Lemma 2.5.

In this proof, we will apply the projection idea as in Section 2.4.2 and design an auxiliary triangle distribution X^{aux} that will allow us to compare $Z_{\sim e}$ and $\mathcal{L}(Z_{\sim e}|Z_e = 0)$ indirectly, by comparing certain triangle distributions. Finally, these triangle distributions will be bounded by showing that their likelihood ratio concentrates.

Proof of Lemma 5.3. Let T_e be the set of triangles in G that contain e . Define $Z = \mathcal{A}_\Delta(G) = G + E(X)$ to be the output graph, so that $\mathcal{L}(Z_{\sim e}|Z_e = 0)$ is the conditional distribution of the output graph given that all triangles in T_e are not selected in X . Let X' be the same as X on $\binom{[n]}{3} \setminus T_e$ and 0 on T_e . We have

$$E(Z_{\sim e}) \stackrel{d}{=} G + E(X) - \{e\} \quad \text{and} \quad E(Z_{\sim e}|Z_e = 0) \stackrel{d}{=} G + E(X') - \{e\}. \quad (5)$$




Figure 3: Relation between X, X', X^{aux} and the corresponding edge sets used in proving Lemma 5.3. Closeness of (1) implies closeness of (2).

Design of Auxiliary Triangle Distribution X^{aux} To prove the lemma, it would be convenient if X and X' were close in distribution, but this is unfortunately not true as X includes a triangle in T_e with constant probability. Luckily, it is the edge set $E(X) - e$ and $E(X')$ that we care about. Our approach is to construct another set of triangles X^{aux} that includes the extra edges in $E(X) - e$ that are missing in $E(X')$, by adding in triangles that could plausibly be in X' . The relation between X, X' and X^{aux} is shown in Figure 3. In order to define X^{aux} we first require a lemma.

Recall that G is c -uniformly 2-star-dense and $e = (u, v) \notin G$. We can choose for each $w \notin \{u, v\}$ a subset $\bar{W}_G(u, w)$ of $W_G(u, w)$ such that the collection satisfies two conditions.

Lemma 5.4. *Let G be a c -uniformly two-star dense graph with n vertices, u, v are vertices in G and $(u, v) \notin G$. There exists $n_0 > 0$ such that for any $n > n_0$, we can choose for every $w \notin \{u, v\}$ a set $\bar{W}_G(u, w) \subset W_G(u, w)$ such that*

- For every $w \notin \{u, v\}$, $|\bar{W}_G(u, w)| = cn/3$, and
- For every pair $w_1, w_2 \notin \{u, v\}$, either $w_1 \notin \bar{W}_G(u, w_2)$ or $w_2 \notin \bar{W}_G(u, w_1)$.

The proof of the lemma can be found in Appendix C. We apply the lemma with v in place of u to also find sets $\bar{W}_G(v, w)$ satisfying the conditions.

We can now define X^{aux} :

1. Start with $X^{\text{aux}} = \vec{0}$.
2. For each $w \notin \{u, v\}$, with probability p' do the following:
 - Sample $w_1 \sim \text{Unif}(\bar{W}_G(u, w))$ and $w_2 \sim \text{Unif}(\bar{W}_G(v, w))$.
 - Set $X^{\text{aux}}_{(u, w, w_1)} = X^{\text{aux}}_{(v, w, w_2)} = 1$.

Since $w_1 \in W_G(u, w)$ and $w_2 \in W_G(v, w)$, the only possible new edges introduced by X^{aux} are (u, w) and (v, w) , which has the same effect on $\binom{[n]}{2} \setminus \{e\}$ as adding triangle (u, v, w) to the graph. Therefore, we have

$$E(Z_{\sim e}) = G + E(X) - \{e\} \stackrel{d}{=} G + E(X' \vee X^{\text{aux}}) - \{e\}. \quad (6)$$




Figure 4: Illustration of X^{aux} . It includes edges that are in $G + E(X) - \{e\}$ but not in $G + E(X')$.

An illustration of X^{aux} is given in Figure 4. Comparing (6) with (5), we have by data processing inequality that

$$d_{\text{TV}}(\mathcal{L}(Z_{\sim e}), \mathcal{L}(Z_{\sim e}|Z_e = 0)) \leq d_{\text{TV}}(\mathcal{L}(X'), \mathcal{L}(X' \vee X^{\text{aux}})), \quad (7)$$

and we now turn to bounding the latter quantity by adding triangles in X^{aux} to X' one at a time.

Adding one triangle at a time Let K be the following channel acting on any input $X \in \{0, 1\}^{\binom{[n]}{3}}$. Sample $w \sim \text{Unif}([n] \setminus \{u, v\})$, $w_1 \sim \text{Unif}(\bar{W}_G(u, w))$ and $w_2 \sim \text{Unif}(\bar{W}_G(v, w))$. Set $X_{(v, w, w_1)}^{\text{aux}} = X_{(v, w, w_2)}^{\text{aux}} = 1$. It turns out that $X' \vee X^{\text{aux}}$ is closely approximated by repeating K on X' a $\text{Binomial}(n - 2, p')$ number of times: we show in Lemma C.1 in the appendix that

$$X' \vee X^{\text{aux}} \approx_{O(p'^2 n)} K^m(X'), \quad \text{for } m \sim \text{Binomial}(n - 2, p'). \quad (8)$$

The only difference between applying K^m and adding triangles in X^{aux} is that K^m may sample the same triangle twice, which gives rise to the $O(p'^2 n)$ error.

Let us use Q and P to denote the p.m.f. of $K(X')$ and X' , respectively. Let $A(x)$ denote the set of possible triples w, w_1, w_2 that could have been chosen by K if $x = K(X')$ for some X' , i.e.,

$$A(x) = \{(w, w_1, w_2) : w \in [n] \setminus \{u, v\}, w_1 \in \bar{W}_G(u, w), w_2 \in \bar{W}_G(v, w), x_{(u, w, w_1)} = x_{(v, w, w_2)} = 1\}.$$

Note that

$$\mathbb{E} |A(X')| = \sum_{\substack{w \in [n] \setminus \{u, v\}, \\ w_1 \in \bar{W}_G(u, w), \\ w_2 \in \bar{W}_G(v, w)}} \mathbb{P}(X_{(u, w, w_1)} = X_{(v, w, w_2)} = 1) = p'^2 (n - 2) (c/3)^2 n^2. \quad (9)$$

Using this, we can write the likelihood ratio between Q and P as

$$\frac{dQ}{dP}(X') = \frac{\sum_{(w, w_1, w_2) \in A(X')} \frac{p'^{|X'| - 2} (1 - p')^{\binom{n}{3} - |X'|}}{(n - 2)(c/3)^2 n^2}}{p'^{|X'|} (1 - p')^{\binom{n}{3} - |X'|}} = \frac{|A(X')|}{p'^2 (n - 2)(c/3)^2 n^2} = \frac{|A(X')|}{\mathbb{E} |A(X')|}. \quad (10)$$

Concentration of the Likelihood Ratio The idea is to use Lemma 4.5 to bound the total variation between P and Q by showing that $|A(X')|$ (and hence dQ/dP) concentrates.

Let $\bar{T}(u, w) = \{(u, w, w_1) : w_1 \in \bar{W}_G(u, w)\}$ and $\bar{T}(v, w) = \{(u, w, w_2) : w_2 \in \bar{W}_G(v, w)\}$. Note that

$$|A(X')| = \sum_{w \in [n] \setminus \{u, v\}} |X'_{\bar{T}(u, w)}| \cdot |X'_{\bar{T}(v, w)}| := \sum_{w \in [n] \setminus \{u, v\}} H_w.$$

By construction of $\bar{W}_G(u, w)$ and $\bar{W}_G(v, w)$ in Lemma 5.4, $\bar{T}(u, w)$ and $\bar{T}(v, w)$ are disjoint sets of triples, so the $H_w = |X'_{\bar{T}(u, w)}| \cdot |X'_{\bar{T}(v, w)}|$ variables are independent. Let H be the vector of all H_w .

To show concentration of $|A(X')|$ one natural approach is to apply McDiarmid's inequality to $|A(X')| = f(H)$. This does not work, because $X'_{\bar{T}(u, w)}$ and $X'_{\bar{T}(v, w)}$ can have size $cn/3$, which leads to a $\Theta(n^2)$ Lipschitz constant. The key is to use the fact that these sets are *typically* much smaller, since each coordinate is 1 only with probability p' . In fact, by a Chernoff bound, $X'_{\bar{T}(u, w)}$ and $X'_{\bar{T}(v, w)}$ are with high probability of size on the order of $\log n$. Each coordinate of $X'_{\bar{T}(u, w)}$ is independent $\text{Bern}(p')$, and there are $cn/3$ coordinates, so we have

$$\mathbb{P}(|X'_{\bar{T}(u, w)}| > 10 \log n) \leq \exp\left(-\frac{\delta^2 cn p'}{3(2 + \delta)}\right),$$

where $\delta = \frac{30 \log n}{cn p'} - 1$. Since $\delta > 2$, we have $\delta^2/(2 + \delta) \geq \delta/2$ and therefore,

$$\exp\left(-\frac{\delta^2 cn p'}{3(2 + \delta)}\right) \leq e^{-\delta cn p'/6} = e^{-5 \log n + cn p'/6} \leq n^{-4}.$$

By the union bound, $|X'_{\bar{T}(u, w)}| \leq 10 \log n$ and $|X'_{\bar{T}(v, w)}| \leq 10 \log n$ for all pairs u, w and v, w with probability at least $1 - 2/n^2$. Hence, with high probability, $H_w \leq 100 \log^2 n$ for all w .

Now we can use the modified McDiarmid's inequality (Lemma 4.10). Let $\mathcal{Y} = \{h : h_w \leq 100 \log^2 n \text{ for all } w\}$. By the calculation above, $\mathbb{P}(H \in \mathcal{Y}) \geq 1 - 2/n^2$. On the other hand, for any $H, H' \in \mathcal{Y}$, $|A(X')| = \sum_w H_w := f(H)$ satisfies $|f(H) - f(H')| \leq \sum_{w: H_w \neq H'_w} 100 \log^2 n$. By Markov's inequality, $\mathbb{P}(|X'_{\bar{T}(u, w)}| > 0) \leq \mathbb{E}|X'_{\bar{T}(u, w)}| = cn p'/3$, so $\mathbb{P}(H_w \neq 0) \leq (cn p'/3)^2$. Applying Lemma 4.10 now yields

$$\mathbb{P}(|A(X')| - \mathbb{E}|A(X')| \geq t) \leq 4/n^2 + 2 \exp\left(\frac{-(t - 200(n - 2) \log^2 n/n^2)^2/2}{100^2(n - 2)(cn p'/3)^2 \log^4 n + 100t \log^2 n/3}\right).$$

Choosing $t = O(n^{3/2} p' \log^{5/2} n)$ makes this probability at most $1/n$. Combining this with $\mathbb{E}|A(X')|$ in (9), we get that with probability at least $1 - 1/n$ the likelihood ratio dQ/dP in (10) is

$$\frac{dQ}{dP}(X') = \frac{|A(X')|}{\mathbb{E}|A(X')|} = 1 + O\left(\frac{n^{3/2} p' \log^{5/2} n}{n^3 p'^2}\right) = 1 + O\left(\frac{\log^{5/2} n}{n^{3/2} p'}\right). \quad (11)$$

Putting the Pieces Together By Lemma 4.5, (11) implies that

$$d_{\text{TV}}(K(X'), X') = d_{\text{TV}}(Q, P) = O\left(\frac{1}{n}\right) + O\left(\frac{\log^{5/2} n}{n^{3/2} p'}\right).$$

By Lemma 4.6 and convexity of total variation distance,

$$d_{\text{TV}}(K^m(X'), X') = O\left(\frac{\mathbb{E}[m] \log^{5/2} n}{n^{3/2} p'}\right) = O(n^{-1/2} \log^{5/2} n),$$

where in the last step we used that $m \sim \text{Binomial}(n-2, p')$. This further implies by (8) that $d_{\text{TV}}(X' \vee X^{\text{aux}}, X') = O(n^{-1/2} \log^{5/2} n)$ and hence by (7), $d_{\text{TV}}(\mathcal{L}(Z_{\sim e}), \mathcal{L}(Z_{\sim e}|Z_e = 0)) = O(n^{-1/2} \log^{5/2} n)$. \square

6 Properties of Triangle Distribution μ_G

Recall the triangle distribution μ_G from Defn. 5. To prepare for proving Theorem 2.2 in the next section, we introduce some notation associated to μ_G and show some basic properties.

In the remainder of the paper, for any events A, B over X , we will use the notation

$$\mu_G(A) := \mathbb{P}(X \in A) \quad \text{and} \quad \mu_G(A|B) := \mathbb{P}(X \in A|B)$$

to distinguish distributions defined on different graphs.

Remark 12 (Graphical Model of μ_G). The graphical model of μ_G is the graph over the set of all triangles in $\binom{[n]}{3}$ with two triangles being adjacent if and only if they share an edge. We assign variables also for triangles that are not in G for convenience in the proofs.

Remark 13 (Conditioning Property of μ_G). For any two graphs G and G' such that $E(G') \supset E(G)$,

$$\mu_G = \mu_{G'}(\cdot | Y_{E(G')-E(G)} = 0)$$

and hence also $\mathcal{L}_G(Y) = \mathcal{L}_{G'}(Y | Y_{E(G')-E(G)} = 0)$.

Below are a few properties of μ_G that will be used later in the proof. We next define two properties of a distribution and show that μ_G satisfies these properties.

Definition 7. A distribution P over $\{0, 1\}^N$ is said to be q -marginally small (or large) if the marginal at any coordinate i of P is bounded by q under arbitrary conditioning $x_{\sim i}$:

$$P(x_i = 1 | x_{\sim i}) \leq q \quad (\text{or } P(x_i = 1 | x_{\sim i}) \geq q).$$

Lemma 6.1 (μ_G has bounded marginals). μ_G is $O(p')$ -marginally small and $\Omega(p')$ -marginally large.

Proof. Let t be a triangle in T_G and let $X_{\sim t}$ be an arbitrary $\{0, 1\}$ configuration on the other triangles. Let X^{+t} be defined by setting $(X^{+t})_t = 1$, $(X^{+t})_{\sim t} = X_{\sim t}$, and let X^{-t} be defined analogously with $(X^{-t})_t = 0$. Then

$$\mu_G(X_t = 1 | X_{\sim t}) = \frac{\left(\frac{p'}{1-p'}\right)^{|X^{+t}|} p^{-e(X^{+t})}}{\left(\frac{p'}{1-p'}\right)^{|X^{-t}|} p^{-e(X^{-t})} + \left(\frac{p'}{1-p'}\right)^{|X^{+t}|} p^{-e(X^{+t})}}.$$

Note that $|X^{+t}| = |X^{-t}| + 1$ and $|e(X^{+t}) - e(X^{-t})| \leq 3$. It follows that

$$p' \leq \mu_G(X_t = 1 | X_{\sim t}) \leq \frac{p' p^{-3}}{1 - p' + p' p^{-3}}.$$

Here we used that p is a constant with respect to n . \square

We restate the definition of uniformly 2-star dense graphs.

Definition 8 (Uniformly 2-star Dense). A graph G is c -uniformly 2-star dense if for any pair $\{i, j\} \in \binom{[n]}{2}$, the wedge set $W_G(i, j) = \{w \in [n] \mid (w, i), (w, j) \in G\}$ has size $|W_G(i, j)| \geq c(n-2)$.

Corollary 6.2. For any c -uniformly 2-star dense graph G and for any $e \in G$, $\mu_G(Y_e = 1) = \Theta(np')$.

Proof. For any edge $e \in G$, let $T_e = \{t_1, t_2, \dots, t_k\}$ be the set of triangles in G that contain e . Since G is c -uniformly 2-star dense, $k = \Theta(n)$. First, $\mu_G(Y_e = 1) = \mu_G(X_{T_e} \neq \vec{0})$, and

$$\mu_G(X_{T_e} \neq \vec{0}) = \mu_G(X_{t_1} = 1) + \mu_G(X_{t_2} = 1 \mid X_{t_1} = 0) + \dots + \mu_G(X_{t_k} = 1 \mid X_{t_1} = \dots = X_{t_{k-1}} = 0).$$

From Lemma 6.1, each term is $\Theta(p')$. So the probability is $\Theta(kp') = \Theta(np')$. \square

As a consequence of the marginal smallness of μ_G , it has concentration of measure and also the edge distribution $\mathcal{L}_G(Y)$ has low influence. The proofs are deferred to Sections 8 and 9.

Definition 9 (Marginal Influence). For a binary distribution over $\{0, 1\}^{\mathcal{X}}$, the *marginal influence* of $S \subset \mathcal{X}$ on $S' \subset \mathcal{X}$ is defined by

$$\mathbf{I}_{S \rightarrow S'}^M = \sup_{x_S^a, x_S^b \in \{0, 1\}^S} d_{\text{TV}}(P_{x_{S'} \mid x_S^a}, P_{x_{S'} \mid x_S^b}).$$

For $x, y \in \mathcal{X}$, use $\mathbf{I}_{x \rightarrow y}^M$ to denote $\mathbf{I}_{\{x\} \rightarrow \{y\}}^M$.

Lemma 6.3. Suppose G is a c -uniformly 2-star dense graph for a constant $c > 0$. Consider distribution $\mathcal{L}_G(Y)$ as in Definition 5. For $A \subset E(G)$ with $|A| \leq cn/4$ and any edge $e \notin A$,

$$\mathbf{I}_{A \rightarrow e}^M = O(|A|/(n^3 p'^2)) = \tilde{O}(|A|/n).$$

For any pair of edges $e, e' \in E(G)$ that do not share common nodes,

$$\mathbf{I}_{e' \rightarrow e}^M = O(1/(n^3 p')) = \tilde{O}(1/n^2).$$

Corollary 6.4. For any c -uniformly 2-star dense graph G and any $e \neq e'$,

$$|\mu_G(Y_e = 1) - \mu_{G+e'}(Y_e = 1)| = \tilde{O}(1/n).$$

Corollary 6.5. Let $f : \{0, 1\}^{\binom{[n]}{3}} \rightarrow \mathbb{R}$ be an L -Lipschitz function. Let T be a subset of $\binom{[n]}{3}$ with size $\binom{n}{3} - m$ and $x_T \in \{0, 1\}^T$ be a configuration on T . If $X \sim \mu_G(\cdot \mid X_T = x_T)$, then

$$|f(X) - \mathbb{E} f(X)| = O(\sqrt{mp' L^2 \log m})$$

holds with high probability.




Figure 5: Exchange Diagram. Theorem 2.2 states that $\mathcal{A}(\text{RGT}(n, p, p', S, q))$ is close to $G(n, p, S, q \cdot p_e)$. An equivalent way of looking at the result is that the operations of resampling edges and mapping \mathcal{A} (approximately) commute when applied to $\text{RGT}(n, p, p')$.

7 Reverse Map Preserves Planted Signal

In this section we prove our main result, restated here for convenience.

Theorem 2.2 (Triangle Removal in Planted Case). *For $k = o(n^{1/4} \log^{-17/4} n)$, $p' = 1/(n \log n)$, and $0 < p < q < 1$ being constant,*

$$d_{\text{TV}}(\mathcal{A}(\text{RGT}(n, p, p', S, q)), G(n, p, S, q \cdot p_e)) = o_n(1),$$

where $p_e = \mathbb{E}_{G \sim \text{RGT}(n, p, p')} \mathbb{P}_{\mathcal{A}}(e \in \mathcal{A}(G + e))$ for an arbitrary edge e .

Remark 14. In Theorem 1.1, we assumed $p' = \tilde{O}(1/n)$. In order to simplify the proof, we only consider $p' = 1/(n \log n) = \tilde{\Theta}(1/n)$, which is the most challenging case. For smaller p' , we map from $\text{RGT}(n, p, p', S, q)$ to $G(n, p, S, qp_e)$ by first adding triangles independently via \mathcal{A}_{Δ} to increase p' up to $\tilde{\Theta}(1/n)$ and then apply the reverse transition \mathcal{A} . This is the content of the following corollary.

Corollary 7.1. *For $k = o(n^{1/4} \log^{-17/4} n)$, $p' = O(1/(n \log n))$ and $0 < p < q < 1$ being constant, there exists an algorithm \mathcal{A}_r (taking a graph and both p and p' as inputs) such that*

$$d_{\text{TV}}(\mathcal{A}_r(\text{RGT}(n, p, p', S, q)), G(n, p, S, q')) = o_n(1),$$

where we let $p'_* = 1/(n \log n)$ and

$$q' = \left(q + (1 - q) \left(1 - \left(1 - \frac{p'_* - p'}{1 - p'} \right)^{n-2} \right) \right) \cdot \mathbb{E}_{G \sim \text{RGT}(n, p, p'_*)} \mathbb{P}_{\mathcal{A}}(e \in \mathcal{A}(G + e)) = q + o(1).$$

Proof. We can first apply \mathcal{A}_{Δ} as stated in Theorem 5.1 with $p'_1 = p'$ and $p'_2 = p'_*$, and then apply \mathcal{A} stated in Theorem 2.2. The resulting edge density in the planted dense subgraph after \mathcal{A}_{Δ} would be $(q + (1 - q)(1 - (1 - p'_{\Delta})^{n-2}))$, where $p'_{\Delta} = \frac{p'_* - p'}{1 - p'}$. So after \mathcal{A} , the density becomes

$$q' = (q + (1 - q)(1 - (1 - p'_{\Delta})^{n-2})) \cdot \mathbb{E}_{G \sim \text{RGT}(n, p, p'_*)} \mathbb{P}_{\mathcal{A}}(e \in \mathcal{A}(G + e)). \quad \square$$

We verify in Appendix B that the resulting edge density $q' = q \cdot p_e$ is still $q + o(1)$.

Lemma 7.2. *The parameter p_e in Theorem 2.2 satisfies*

$$p_e = 1 - O(np') = 1 + o_n(1).$$

In the rest of the section, we will prove Theorem 2.2. Recall that in Section 2.3 we proved Theorem 2.2 given Lemma 2.4 (that the graphs \mathcal{G} have high probability) and Lemma 2.5 (that \mathcal{A} satisfies perturbation insensitivity). What remains is to prove both lemmas. We will prove Lemma 2.4 in Section 7.5 via a concentration argument. We start with Lemma 2.5.

Lemma 2.5. *If $G \in \mathcal{G}_1$, then $\mathcal{L}_{G-e}(Y_{\sim e}) = \mathcal{L}_{G+e}(Y_{\sim e}|Y_e = 0) \approx_{O(\log^{17/2} n/\sqrt{n})} \mathcal{L}_{G+e}(Y_{\sim e}|Y_e = 1)$.*

In this section we always focus on G and $G+e$ where $e = (u, v) \notin G$. Let $T_e = T_{G+e} \setminus T_G$ denote the set of triangles in $G+e$ containing edge e . For an indicator x of triangles, we will always use x_e to denote x_{T_e} and $x_{\sim T(e)}$ to denote x_{T_G} . Let W denote the set of edges in G that are incident to e .

Fix G and let $X \sim \mu_{G+e}(\cdot | Y_e = 0)$ and $X^+ \sim \mu_{G+e}(\cdot | Y_e = 1)$, with Y and Y^+ the corresponding edge distributions. Lemma 2.5 is equivalent to $Y_{\sim e}$ and $Y_{\sim e}^+$ being close in distribution, i.e. we want to prove that

$$d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^+) = O(\log^{17/2} n/\sqrt{n}). \quad (12)$$

Outline of Proof of Lemma 2.5. We actually do not work with X^+ and Y^+ , instead replacing these by X^* and Y^* , which we define later in Section 7.2 and will turn out to be easier to work with. We will show that

$$\begin{aligned} d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^+) &= \Theta(d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^*)/\log n) \\ d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^+) &= \Theta(d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^*)/\log n). \end{aligned} \quad (13)$$

Now, similar to the decomposition (1), we decompose $d_{\text{TV}}(Y_{\sim e}^*, Y_{\sim e})$ using Lemma 4.3 to obtain

$$\begin{aligned} d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^*) &\leq d_{\text{TV}}((Y_{\sim e}, X_{\sim T(e)}), (Y_{\sim e}^*, X_{\sim T(e)}^*)) \\ &\leq d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^*) + d_{\text{TV}}(Y_{\sim e}, \mathbb{E}_{X' \sim X_{\sim T(e)}} \mathcal{L}(Y_{\sim e}^* | X^* = X')). \end{aligned} \quad (14)$$

Section 7.1 follows the sketch in Section 2.4.1 to bound $d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^+)$ which by (13) bounds the first term in (14) by $O(\log n/\sqrt{n})$. Section 7.3 uses the ideas in Section 2.4.2 to bound the second term in (14) by $O(\log^{15/2} n/\sqrt{n})$. \square

7.1 Influence of an Edge on Non-Containing Triangles

We will now bound $d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^+)$. In Section 7.1 only, we will overload notation and take $X^+ \sim \mu_{G+e}$ rather than $\mu_{G+e}(\cdot | Y_e = 1)$. This is justified by the following lemma, which is immediate from the definitions.

Lemma 7.3. *Let $X \sim \mu_G$, $X^{+,1} \sim \mu_{G+e}(\cdot | Y_e = 1)$, and $X^{+,2} \sim \mu_{G+e}$. Then,*

$$d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^{+,1}) = \frac{1}{\mathbb{P}(Y_e = 1)} d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^{+,2}).$$

Lemma 7.4. *Let G be a c -uniformly 2-star dense graph for a constant $c > 0$. Let $e \in \binom{[n]}{2}$ be any edge and let $X \sim \mu_{G+e}(\cdot | Y_e = 0)$ and $X^+ \sim \mu_{G+e}$. Emphasizing that $X_{\sim T(e)} = X$, we have*

$$d_{\text{TV}}(X_{\sim T(e)}, X_{\sim T(e)}^+) = O(\log n/\sqrt{n}).$$

This rest of this section is structured as follows: First, in Section 7.1.1, we establish the relationship between $\chi^2(X_{T(e)}^+ \| X_{\sim T(e)})$ and the influence of the distribution $\mathcal{L}_G(Y)$, as presented in Lemma 2.7. Subsequently, in Section 7.1.2, we combine Lemma 2.7 with the bound on influence obtained in Lemma 6.3 to compute the desired upper bound in Lemma 7.4.

7.1.1 Bounding χ^2 -divergence In Terms of Influence

The entirety of Section 7.1.1 is dedicated to showing Lemma 2.7, which states that the χ^2 -divergence between $X_{\sim T(e)}$ and $X_{T(e)}^+$ is bounded in terms of influence.

Lemma 2.7. *Suppose $G \in \mathcal{G}$ and consider $\mathcal{L}_G(Y)$ for Y as in Section 2.2 and $p' = o(1/n)$. Let $\mathbf{I}_{A' \rightarrow e}^M$ be the marginal influence of A' on e for $\mathcal{L}_G(Y)$. Letting $\tilde{X}_{T(e)}^+$ be an independent copy of $X_{T(e)}^+$, we have*

$$\chi^2(X_{\sim T(e)}^+ \| X_{\sim T(e)}) + 1 \leq \mathbb{E}_{X_{T(e)}^+, \tilde{X}_{T(e)}^+} \min \left\{ \left(1 + \frac{1-p}{p} \sup_{\substack{A' \subset A \cup B \\ e \in (A \cup B) \setminus A'}} \mathbf{I}_{A' \rightarrow e}^M \right)^{|B|} (1/p)^{|A \cap B|}, (1/p)^{|A|+|B|} \right\},$$

where $A = E(\tilde{X}_{T(e)}^+) - e$ and $B = E(X_{T(e)}^+) - e$.

We restate the lemma that bounds chi-square divergence for a mixture of Gibbs distributions.

Lemma 2.6 (χ^2 -divergence for mixture of Gibbs measures). *Let P be a distribution defined by*

$$P(X) = f(X)/Z.$$

Let U be a discrete random variable and Q be a mixture of Gibbs distributions defined by

$$Q(X) = \mathbb{E}_U[f_U(X)/Z_U].$$

Letting $\rho_U(X) = f_U(X)/f(X)$, we have that

$$\chi^2(Q \| P) + 1 = \mathbb{E}_{U, U'} \frac{\mathbb{E}_{X \sim P} \rho_U(X) \rho_{U'}(X)}{\mathbb{E}_{X \sim P} \rho_U(X) \mathbb{E}_{X \sim P} \rho_{U'}(X)},$$

where U' is an independent and identically distributed copy of U .

Proof. By definition of chi-square divergence,

$$\begin{aligned} \chi^2(Q \| P) + 1 &= \mathbb{E}_{X \sim P} \left(\frac{dQ}{dP}(X) \right)^2 = \mathbb{E}_{X \sim P} \frac{\mathbb{E}_U[f_U(X)/Z_U] \mathbb{E}_{U'}[f_{U'}(X)/Z_{U'}]}{(f(X)/Z)^2} \\ &= \mathbb{E}_{U, U'} \left[\frac{Z^2}{Z_U Z_{U'}} \cdot \mathbb{E}_{X \sim P} \frac{f_U(X) f_{U'}(X)}{(f(X))^2} \right]. \end{aligned}$$

Notice that

$$\frac{Z_U}{Z} = \sum_X \frac{f_U(X)}{f(X)} \frac{f(X)}{Z} = \mathbb{E}_{X \sim P} \rho_U(X).$$

Taking this into the previous expression finishes the proof. \square

Applying the above lemma to $X_{\sim T(e)}$ and $X_{\sim T(e)}^+$ gives us the following upper bound on their chi-square distance.

Corollary 7.5. *Let $A = E(\tilde{X}_{T(e)}^+) - e$, and $B = E(X_{T(e)}^+) - e$. Recall that Y is the indicator of $E(X)$, we have*

$$\chi^2(X_{\sim T(e)}^+ \| X_{\sim T(e)}) + 1 \leq \mathbb{E}_{X_{T(e)}^+, \tilde{X}_{T(e)}^+} \frac{\mathbb{E} [p^{|Y_A|} \mathbb{E} [p^{|Y_B|} | Y_A]]}{\mathbb{E} [p^{|Y_A|}] \mathbb{E} [p^{|Y_B|}]} . \quad (15)$$

Proof. Let $P = \mathcal{L}(X_{\sim T(e)}) = \mathcal{L}(X) = \mu_G$ and $Q = \mathcal{L}(X_{\sim T(e)}^+)$. Define $f(x) := \left(\frac{p'}{1-p'}\right)^{|x|} p^{-e(x)}$ so that $P(x) = f(x)/Z_G$ by Defn. 5 of μ_G .

The distribution of $X_{\sim T(e)}^+$ can be viewed as a mixture of Gibbs distribution given $X_{T(e)}^+$. We have

$$\begin{aligned} \mu_{G+e}(X_{\sim T(e)}^+ | X_{T(e)}^+) &\propto p'^{|X_{\sim T(e)}^+| + |X_{T(e)}^+|} p^{-e(X_{\sim T(e)}^+) - e(X_{T(e)}^+)} p^{|E(X_{\sim T(e)}^+) \cap E(X_{T(e)}^+)|} \\ &\propto p'^{|X_{\sim T(e)}^+|} p^{-e(X_{\sim T(e)}^+)} p^{|E(X_{\sim T(e)}^+) \cap E(X_{T(e)}^+)|} . \end{aligned}$$

We can define

$$f_{X_{T(e)}^+}(X) = p^{n_X} p^{-e(X)} p^{|E(X) \cap E(X_{T(e)}^+)|} = f(X) p^{|E(X) \cap E(X_{T(e)}^+)|}$$

and

$$Z_{X_{T(e)}^+} = \sum_{X \in \{0,1\}^{T_G}} f_{X_{T(e)}^+}(X)$$

so that

$$\mu_{G+e}(X_{\sim T(e)}^+ | X_{T(e)}^+) = f_{X_{T(e)}^+}(X_{\sim T(e)}^+) / Z_{X_{T(e)}^+} .$$

So μ_{G+e} can be viewed as the mixture of Gibbs distributions, i.e.,

$$\mu_{G+e}(X) = \mathbb{E}_{X_{T(e)}^+} [f_{X_{T(e)}^+}(X_{\sim T(e)}^+) / Z_{X_{T(e)}^+}] .$$

Note that $f_{X_{T(e)}^+}(X) / f(X) = p^{|E(X) \cap E(X_{T(e)}^+)|}$. We can let $\tilde{X}_{T(e)}^+$ be an independent copy of $X_{T(e)}^+$, and apply Lemma 2.6 to obtain

$$\begin{aligned} \chi^2(Q \| P) + 1 &= \mathbb{E}_{X_{T(e)}^+, \tilde{X}_{T(e)}^+} \frac{\mathbb{E}_{X \sim P} p^{|E(X) \cap E(X_{T(e)}^+)|} p^{|E(X) \cap E(\tilde{X}_{T(e)}^+)|}}{\mathbb{E}_{X \sim P} p^{|E(X) \cap E(X_{T(e)}^+)|} \mathbb{E}_{X \sim P} p^{|E(X) \cap E(\tilde{X}_{T(e)}^+)|}} \\ &= \mathbb{E}_{X_{T(e)}^+, \tilde{X}_{T(e)}^+} \frac{\mathbb{E} \left[p^{|Y_{E(X_{T(e)}^+)}|} p^{|Y_{E(\tilde{X}_{T(e)}^+)}|} \right]}{\mathbb{E} p^{|Y_{E(X_{T(e)}^+)}|} \mathbb{E} p^{|Y_{E(\tilde{X}_{T(e)}^+)}|}} . \end{aligned}$$

For simplicity of notation, let $A = E(\tilde{X}_{T(e)}^+) - e$, $B = E(X_{T(e)}^+) - e$ ($e \notin E(X)$ so $Y_e = 0$). The above expression can be written as

$$\mathbb{E}_{X_{T(e)}^+, \tilde{X}_{T(e)}^+} \frac{\mathbb{E} [p^{|Y_A|} \mathbb{E} [p^{|Y_B|} | Y_A]]}{\mathbb{E} [p^{|Y_A|}] \mathbb{E} [p^{|Y_B|}]} . \quad \square$$

From Corollary 7.5, we can derive Lemma 2.7 by showing two upper bounds for the expression within $\mathbb{E}_{X_{T(e)}^+, \tilde{X}_{T(e)}^+}$ in (15).

Lemma 7.6. *Let \mathbf{I}^M denote the marginal influence of distribution $\mathcal{L}_G(Y)$. For any $A, B \subset E(G)$, we have*

$$\frac{\mathbb{E} [p^{|Y_A|} \mathbb{E} [p^{|Y_B|} | Y_A]]}{\mathbb{E} [p^{|Y_A|}] \mathbb{E} [p^{|Y_B|}]} \leq (1/p)^{|A|+|B|}$$

and

$$\frac{\mathbb{E} [p^{|Y_A|} \mathbb{E} [p^{|Y_B|} | Y_A]]}{\mathbb{E} [p^{|Y_A|}] \mathbb{E} [p^{|Y_B|}]} \leq \left(1 + \frac{1-p}{p} \sup_{\substack{A' \subset A \cup B \\ e \in (A \cup B) \setminus A'}} \mathbf{I}_{A' \rightarrow e}^M\right)^{|B|} (1/p)^{|A \cap B|}.$$

Proof. The first upper bound is straightforward. On the one hand, $\mathbb{E} [p^{|Y_A|} \mathbb{E} [p^{|Y_B|} | Y_A]] \leq 1$, and on the other hand, $\mathbb{E} [p^{|Y_A|}] \mathbb{E} [p^{|Y_B|}] \geq p^{|A|+|B|}$, so we have

$$\frac{\mathbb{E} [p^{|Y_A|} \mathbb{E} [p^{|Y_B|} | Y_A]]}{\mathbb{E} [p^{|Y_A|}] \mathbb{E} [p^{|Y_B|}]} \leq (1/p)^{|A|+|B|}.$$

For the second upper bound, we compare the difference between $\mathbb{E} [p^{|Y_B|} | Y_A]$ and $\mathbb{E} [p^{|Y_B|}]$ by constructing a coupling between Y_B and $Y_B | Y_A = y_A$. The reason that influence is useful to us, is that a good coupling can be constructed for a distribution with low influence, as stated in the following lemma.

Lemma 7.7. *Let X be a random variable on $\{0, 1\}^N$. $A, B \subset [N]$ are disjoint. If for any subset $A' \subset A \cup B$ and $i \in (A \cup B) \setminus A'$,*

$$\mathbf{I}_{A' \rightarrow i}^M \leq q,$$

then for any $x_A, x'_A \in \{0, 1\}^B$, there exists coupling (X_B, X'_B) between $\mathcal{L}(X_B | X_A = x_A)$ and $\mathcal{L}(X_B | X_A = x'_A)$ such that $|X_B - X'_B|$ is stochastically dominated by $\text{Binomial}(|B|, q)$.

We prove the lemma later in the section. Now use I_m to denote $\sup_{\substack{A' \subset A \cup B \\ e \in (A \cup B) \setminus A'}} \mathbf{I}_{A' \rightarrow e}^M$ and apply Lemma 7.7 on A and $B \setminus A$. There exists a coupling $(Y_{B \setminus A}, Y'_{B \setminus A})$ between $\mathcal{L}_G(Y_{B \setminus A} | Y_A = y_A)$ and $\mathcal{L}_G(Y_{B \setminus A})$ such that $|Y_{B \setminus A} - Y'_{B \setminus A}|$ is stochastically dominated by $\text{Binomial}(|B \setminus A|, I_m)$. Therefore, we can couple $Y_{A \cap B}$ arbitrarily and obtain a coupling (Y_B, Y'_B) between $\mathcal{L}_G(Y_B | Y_A = y_A)$ and $\mathcal{L}_G(Y_B)$ where $|Y_B - Y'_B|$ is stochastically dominated by $\text{Binomial}(|B \setminus A|, I_m) + |A \cap B|$.

Let $M(\theta) = \mathbb{E}[e^{\theta Z}]$ be the moment generating function (MGF) of $Z \sim \text{Binomial}(|B|, q^*(|A| + |B|)) + |A \cap B|$. By the MGF of binomial distribution, $M(\theta) = (1 - I_m + e^\theta I_m)^{|B|} e^{\theta |A \cap B|}$. We have

$$\mathbb{E} [p^{|Y'_B| - |Y_B|}] \leq \mathbb{E} [(1/p)^{|Y_B - Y'_B|}] \leq M(\log(1/p)) = \left(1 + \frac{1-p}{p} I_m\right)^{|B|} (1/p)^{|A \cap B|}.$$

So for any y_A ,

$$\mathbb{E} [p^{|Y_B|} | Y_A = y_A] = \mathbb{E} [p^{|Y'_B|}] \leq \mathbb{E} [p^{|Y_B|}] \left(1 + \frac{1-p}{p} I_m\right)^{|B|} (1/p)^{|A \cap B|}.$$

Hence

$$\frac{\mathbb{E} [p^{|Y_A|} \mathbb{E} [p^{|Y_B|} | Y_A]]}{\mathbb{E} [p^{|Y_A|}] \mathbb{E} [p^{|Y_B|}]} \leq \left(1 + \frac{1-p}{p} I_m\right)^{|B|} (1/p)^{|A \cap B|}. \quad \square$$

Proof of Lemma 7.7. Suppose $B = \{i_1, \dots, i_k\}$. We couple one coordinate at a time. Since $\mathbf{I}_{A \rightarrow i_1}^M \leq q$,

$$|\mathbb{P}(X_{i_1} = 1 | X_A = x_A) - \mathbb{P}(X_{i_1} = 1 | X_A = x'_A)| \leq q.$$

Let (X_{i_1}, X'_{i_1}) be the optimal coupling between the two distributions. $\mathbb{P}(X_{i_1} \neq X'_{i_1}) \leq q$. The coupling for other coordinates can be defined similarly one by one. Again from the assumption,

$$\mathbf{I}_{A \cup \{i_1, \dots, i_{j-1}\} \rightarrow i_j}^M \leq q.$$

So

$$|\mathbb{P}(X_{i_j} = 1 | X_A = x_A, X_{\{i_1, \dots, i_{j-1}\}}) - \mathbb{P}(X_{i_j} = 1 | X_A = x'_A, X'_{\{i_1, \dots, i_{j-1}\}})| \leq q,$$

so we can set (X_{i_j}, X'_{i_j}) conditioned on $X_{\{i_1, \dots, i_{j-1}\}}, X'_{\{i_1, \dots, i_{j-1}\}}$ to be the optimal coupling between the two distributions and have

$$\mathbb{P}(X_{i_j} \neq X'_{i_j} | X_{\{i_1, \dots, i_{j-1}\}}, X'_{\{i_1, \dots, i_{j-1}\}}) \leq q.$$

Equivalently, $|X_{i_j} - X'_{i_j}|$ conditioned on $X_{\{i_1, \dots, i_{j-1}\}}, X'_{\{i_1, \dots, i_{j-1}\}}$ is stochastically dominated by $\text{Bern}(q)$. So $|X_B - X'_B|$ is stochastically dominated by $\text{Binomial}(|B|, q)$. \square

7.1.2 Plugging Influence Bound into Lemma 2.7 to Obtain Lemma 7.4

Proof of Lemma 7.4. To evaluate the bound in Lemma 2.7, we use the bound on the marginal influence of $\mathcal{L}_G(Y)$ in Lemma 6.3. Let $q^* = O(1/(n^3 p'^2)) = \tilde{O}(1/n)$ be the factor in the bound in Lemma 6.3, i.e., $\mathbf{I}_{A \rightarrow e}^M \leq q^* |A|/n$ for any A such that $|A| \leq cn/4$. It then follows that for $|A| + |B| \leq cn/4$,

$$\sup_{\substack{A' \subset A \cup B \\ e \in (A \cup B) \setminus A'}} \mathbf{I}_{A' \rightarrow e}^M \leq q^* (|A| + |B|).$$

We now evaluate the bound in Lemma 2.7 in terms of q^* .

Noting that the part within $\mathbb{E}_{X_{T(e)}^+, \tilde{X}_{T(e)}^+}$ in the right-hand side is monotone in $|A|$ and $|B|$, we begin by showing an upper bound on their distributions. For any triangle t in T_e , $\mathbb{P}(t \in X_{T(e)}^+) \leq c_0 \cdot p'$ for some $c_0 > 0$ since μ_{G+e} is $O(p')$ marginally small. It follows that the distribution of $|X_{T(e)}^+|$ (and $|\tilde{X}_{T(e)}^+|$) is stochastically dominated by $\text{Binomial}(|T_e|, c_0 p')$. As $|T_e| \leq n$, they are also stochastically dominated by $\text{Binomial}(n, c_0 p')$.

Note that A and B do not include edge e , $|A| = 2|X_{T(e)}^+|$, and $|B| = 2|\tilde{X}_{T(e)}^+|$, so $|A \cap B| = 2|X_{T(e)}^+ \cap \tilde{X}_{T(e)}^+|$. By using the first bound in the minimum in the conclusion of Lemma 2.7 when

$|A|, |B|$ are not larger than $2 \log n$ and using the second bound otherwise, we obtain

$$\begin{aligned}
& \chi^2(X_{\sim T(e)}^+ \| X_{\sim T(e)}) \\
& \leq \sum_{1 \leq i, j \leq 2 \log n} \sum_{k=0}^{\min\{i, j\}} \mathbb{P}(|\tilde{X}_{T(e)}^+| = i, |X_{T(e)}^+| = j, |X_{T(e)}^+ \cap \tilde{X}_{T(e)}^+| = k) \\
& \quad \cdot \left(\left(1 + \frac{1-p}{p} q^*(2i+2j) \right)^{2j} (1/p)^{2k} - 1 \right) \\
& \quad + \sum_{\substack{i \geq 2 \log n \text{ or} \\ j \geq 2 \log n}} \mathbb{P}(|\tilde{X}_{T(e)}^+| = i, |X_{T(e)}^+| = j) \left((1/p)^{2i+2j} - 1 \right) \\
& \leq \sum_{1 \leq i, j \leq 2 \log n} \sum_{k=0}^{\min\{i, j\}} \binom{n}{k} \binom{n-k}{i-k} \binom{n-i-k}{j-k} (c_0 p')^{i+j} \left(\left(1 + \frac{1-p}{p} q^*(2i+2j) \right)^{2j} (1/p)^{2k} - 1 \right) \\
& \quad + \sum_{\substack{i \geq 2 \log n \text{ or} \\ j \geq 2 \log n}} \binom{n}{i} \binom{n}{j} (c_0 p')^{i+j} (1/p)^{2i+2j}.
\end{aligned}$$

Because $\binom{n}{i} \leq n^i$ and $\binom{n}{j} \leq n^j$, the second summation is bounded by

$$\sum_{\substack{i \geq 2 \log n \text{ or} \\ j \geq 2 \log n}} \left(\frac{c_0 n p'}{p} \right)^{i+j}.$$

Since $n p' \ll 1$, this sum is $o(1/n^2)$.

For the first summation, we can use $\binom{n}{k} \binom{n-k}{i-k} \binom{n-i-k}{j-k} \leq n^{i+j-k}$ and separate the $k=0$ summand to bound it by

$$\sum_{1 \leq i, j \leq 2 \log n} \left\{ (c_0 n p')^{i+j} \left(\left(1 + \frac{1-p}{p} q^*(2i+2j) \right)^{2j} - 1 \right) + \left(1 + \frac{1-p}{p} q^*(2i+2j) \right)^{2j} \sum_{k \geq 1} (n p^2)^{-k} \right\}. \quad (16)$$

Since $q^* = \tilde{O}(1/n)$, we have $2j \cdot \frac{1-p}{p} q^*(2i+2j) \ll 1$, so

$$\left(1 + \frac{1-p}{p} q^*(2i+2j) \right)^{2j} \leq 1 + \frac{8(1-p)}{p} q^* j(i+j) \leq 2.$$

Therefore, for large enough n , the quantity in (16) can be upper bounded by

$$\sum_{i, j \geq 1} \frac{8(1-p)}{p} q^* (c_0 n p')^{i+j} j(i+j) + \frac{8 \log^2 n}{n p^2}.$$

As $n p' \ll 1$, we have $\sum_{i \geq 1} (c_0 n p')^i i = O(n p')$ and $\sum_{i \geq 1} (c_0 n p')^i i^2 = O(n p')$. So (16) is bounded by $O(n p' q^*) + O(\log^2 n / n) = O(\log^2 n / n)$. Now that we proved $\chi^2(Q \| P) = O(\log^2 n / n)$, by Pinsker's inequality,

$$d_{\text{TV}}(P, Q) \leq \sqrt{2 \chi^2(Q \| P)} = O(n^{-1/2} \log n).$$

This finishes the proof for Lemma 7.4. \square

7.2 Replacing X^+ and Y^+ by X^* and Y^*

As compared to the sketch in Section 2.4.2, there are a number of additional steps. Instead of working with X and X^+ , we replace X^+ by a new random variable X^* constructed as a mixture of the distributions of X and X^+ so that the conditional distribution on X_e^* given $X_{\sim T(e)}^*$ is a product distribution. As shown in the proof, it will suffice to bound the total variation between $Y_{\sim e}$ and $Y_{\sim e}^*$. We first define X^* .

X^* as a mixture of $\mathcal{L}(X)$ and $\mathcal{L}(X^+)$ to make $\mathcal{L}(X_e^*|X_{\sim T(e)}^*)$ a product distribution The conditioning on edge e introduces dependence between triangles in T_e . We modify Definition 5 to yield a new distribution over sets of triangles in $G + e$. Let

$$g_{G+e}^*(x) \triangleq \left(\frac{p'}{1-p'}\right)^{|x|} p^{-|E(x)-e|} \quad \text{and} \quad \mu_{G+e}^*(X^* = x) = \frac{g_{G+e}^*(x)}{Z^*(G+e)}.$$

Here $Z^*(G+e) = \sum_{x: E(x) \subset G+e} g^*(x)$ is the normalizing factor. The benefit of defining μ^* is that now

$$\mu_{G+e}^*(X_e|X_{\sim T(e)}) \propto \left(\frac{p'}{1-p'}\right)^{n_{X_e}} p^{-|E(X_e)-E(X_{\sim T(e)})-e|}.$$

Because triangles in T_e only intersect at e , $|E(X_e) - E(X_{\sim T(e)}) - e| = \sum_{t \in T_e} X_t |E(t) - E(X_{\sim T(e)})|$. So

$$\mu_{G+e}^*(X_e|X_{\sim T(e)}) \propto \prod_{t \in T_e} \left(\frac{p' p^{-|E(t)-E(X_{\sim T(e)})|}}{1-p'}\right)^{X_t} \quad (17)$$

is now a product distribution on triangles in T_e .

Note by definition of g^* that

$$g_{G+e}^*(X) = \begin{cases} p g_{G+e}(X) & \text{if } e \in E(X) \\ g_{G+e}(X) & \text{otherwise.} \end{cases}$$

Thus, μ^* re-weights by p those X with $e \in E(X)$, which is the same as $Y_e = 1$. It follows that μ_{G+e}^* is equal to the mixture

$$\mu_{G+e}^*(\cdot) = \frac{p \mu_{G+e}(Y_e = 1) \mu_{G+e}(\cdot | Y_e = 1) + \mu_{G+e}(Y_e = 0) \mu_{G+e}(\cdot | Y_e = 0)}{p \mu_{G+e}(Y_e = 1) + \mu_{G+e}(Y_e = 0)}. \quad (18)$$

Let $X^* \sim \mu_{G+e}^*$, and Y^* be the corresponding edge distribution of X^* on $\{0, 1\}^{E(G+e)}$. Recalling that Y^+ has law $\mathcal{L}_{G+e}(Y|Y_e = 1)$ and $\mathcal{L}(Y) = \mathcal{L}_{G+e}(Y|Y_e = 0)$, we also have

$$\mathcal{L}(Y^*) = \frac{p \mu_{G+e}(Y_e = 1) \mathcal{L}(Y^+) + \mu_{G+e}(Y_e = 0) \mathcal{L}(Y)}{p \mu_{G+e}(Y_e = 1) + \mu_{G+e}(Y_e = 0)}. \quad (19)$$

Why we can replace X^+ and Y^+ with X^* and Y^* The utility of introducing X^* and Y^* is that they are easier to work with, due the product structure in (17), and we now show that we can replace X^+ and Y^+ by X^* and Y^* in the arguments that follow, as promised in (13). By Corollary 6.2, we have that $\mu_{G+e}(Y_e = 1) = \tilde{\Theta}(1)$, and it follows from (19) that

$$\begin{aligned} d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^+) &= \frac{p \mu_{G+e}(Y_e = 1) + \mu_{G+e}(Y_e = 0)}{p \mu_{G+e}(Y_e = 1)} d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^*) \\ &= \Theta(d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^*) / (np')) = \tilde{\Theta}(d_{\text{TV}}(Y_{\sim e}, Y_{\sim e}^*)). \end{aligned}$$

As far as replacing X^+ by X^* , the identical argument as above gives from (18) that

$$d_{\text{TV}}(X_{\sim T(e)}^*, X_{\sim T(e)}) = \Theta(np' \cdot d_{\text{TV}}(X_{\sim T(e)}^+, X_{\sim T(e)})) = \tilde{\Theta}(d_{\text{TV}}(X_{\sim T(e)}^+, X_{\sim T(e)})).$$

It remains is to bound the second term in (14) using the ideas from Section 2.4.2.

7.3 Perturbation Insensitivity and Proof of Lemma 2.5

In this section bound the second term of (14). The ideas sketched in Section 2.4.2 are implemented, transforming the comparison between edge variables back into a comparison between triangle distributions. Then we compare the triangle distributions by accounting for one triangle addition at a time, and finally, we bound the total variation induced by a single triangle via concentration of the likelihood ratio.

For the simplicity of discussion, let us create a random variable that follows distribution $\mathbb{E}_{X' \sim X_{\sim T(e)}} \mathcal{L}(Y_{\sim e}^* | X^* = X')$. Specifically, let X'' be generated from $X_{\sim T(e)}$ by channel $\mu_{G+e}^*(\cdot | X_{\sim T(e)})$, and Y'' be the corresponding edge distribution, so $\mathcal{L}(Y) = \mathbb{E}_{X' \sim X_{\sim T(e)}} \mathcal{L}(Y_{\sim e}^* | X^* = X')$. This way, the second term of (14) is equal to

$$d_{\text{TV}}(Y_{\sim e}, \mathbb{E}_{X' \sim X_{\sim T(e)}} \mathcal{L}(Y_{\sim e}^* | X^* = X')) = d_{\text{TV}}(Y_{\sim e}, Y''_{\sim e}).$$

The goal of this subsection is to bound the latter quantity.

Lemma 7.8. $d_{\text{TV}}(Y_{\sim e}, Y''_{\sim e}) = O\left(\frac{\log^3 n}{n^5 p^{9/2}}\right) = O(\log^{15/2} / \sqrt{n})$.

7.3.1 Step 1: Projection of Distributions and Defining X^{aux}

Although we will show that $Y_{\sim e}$ and $Y''_{\sim e}$ are close, X and X'' are actually very different distributions, in fact, they have $\tilde{\Omega}(1)$ TV-distance in general. Indeed, $X_e = \vec{0}$, but $X''_e | X_{\sim T(e)}$ follows (17). X''_e might bring new edges that makes Y'' different from Y . We will construct an auxiliary triangle distribution \tilde{X} with $Y''_{\sim e} = \tilde{Y}_{\sim e}$, and by data-processing inequality it suffices to bound $d_{\text{TV}}(X, \tilde{X})$ (see Figure 1):

Lemma 7.9. *If \tilde{X} is such that $Y''_{\sim e} = \tilde{Y}_{\sim e}$, then $d_{\text{TV}}(Y_{\sim e}, Y''_{\sim e}) = d_{\text{TV}}(Y_{\sim e}, \tilde{Y}_{\sim e}) \leq d_{\text{TV}}(X, \tilde{X})$.*

Specifically, X^{aux} will be constructed such that when added to X , it results in the same edge projection as X'' , i.e.,

$$E(X'') - \{e\} = E(X^{\text{aux}} \vee X), \quad (20)$$

We can set $\tilde{X} = X^{\text{aux}} \vee X$. Note that $Y''_{\sim e}$ is the indicator for edge set $E(X'') - \{e\}$, this means $Y''_{\sim e} = \tilde{Y}_{\sim e}$.

We now formally define X^{aux} by its distribution conditioned on $X'' = X \vee X''_e$ (see Figure 2). To satisfy (20), X^{aux} needs to include the set of new edges that are in $E(X'')$ but not in $E(X)$,

$$E^{\text{new}} := E(X'') - E(X) - \{e\} = E(X''_e) - E(X) - \{e\} \subset W.$$

Start with X^{aux} being 0. Recall the definition of wedge set in (4). For each edge (u, w) (or (v, w)) in E^{new} , pick a node w' uniformly at random from $W_{E(X)}(u, w)$ (or $W_{E(X)}(v, w)$) and let $X^{\text{aux}}_{(u, w, w')} = 1$ (or $X^{\text{aux}}_{(v, w, w')} = 1$).

Lemma 7.10. $E^{\text{new}} = E(X^{\text{aux}}) - E(X)$.

Proof. Let (u, w) be an edge in $E(X'') - E(X) - \{e\}$, and w' be the pick from $W_{E(X)}(u, w)$ that $X_{(u, w, w')}^{\text{aux}} = 1$. We have $(u, w) \in E(X'')$. By (4), $(u, w'), (w, w') \in E(X)$. So $E(X^{\text{aux}}) \subset E(X'')$. As $e \notin E(X)$, u or v are not in $W_{E(X)}(u, w)$, we have $E(X^{\text{aux}}) \subset E(X'') - \{e\}$, so

$$E(X^{\text{aux}}) - E(X) \subset E(X'') - E(X) - \{e\}.$$

Since each edge in E^{new} is included in X^{aux} , we know

$$E(X^{\text{aux}}) - E(X) \supset E^{\text{new}}.$$

□

Using this lemma we have $E(\tilde{Y}_{\sim e}) = E(X^{\text{aux}} \vee X) = E(X'') - \{e\} = E(Y''_{\sim e})$.

The remainder of this section shows that X and $\tilde{X} = X \vee X^{\text{aux}}$ are close in total variation.

Lemma 7.11. X and $\tilde{X} = X \vee X^{\text{aux}}$ satisfy the total variation bound

$$d_{\text{TV}}(X, \tilde{X}) = O\left(\frac{\log^3 n}{n^5 p^{9/2}}\right) = \tilde{O}\left(\frac{1}{\sqrt{n}}\right).$$

7.3.2 Step 2: Representation of X^{aux}

To prove the lemma, we need to know the distribution of $X^{\text{aux}}|X$. Because X^{aux} is defined through E^{new} , we first consider its distribution. By definition, upon conditioning on $X_{\sim T(e)}$, X_e'' has distribution

$$\mu_{G+e}^*(X_e|X_{\sim T(e)}) \propto \prod_{t \in T_e} \left(\frac{p' p^{-|E(t) - E(X_{\sim T(e)})|}}{1 - p'} \right)^{X_t}.$$

So for $t = (w, u, v) \in T_e$,

$$X_t''|X_{\sim T(e)} \sim \begin{cases} \text{Bern}(p_0) & \text{if none of } (w, u) \text{ and } (w, v) \text{ are in } E(X) \\ \text{Bern}(p_1) & \text{if one of } (w, u) \text{ and } (w, v) \text{ are in } E(X) \\ \text{Bern}(p_2) & \text{if both } (w, u) \text{ and } (w, v) \text{ are in } E(X). \end{cases}$$

Here $p_i = \frac{p' p^{-3+i}}{1 - p' + p' p^{-3+i}}$ is $\Theta(p')$, $\forall i = 0, 1, 2$. Define

$$\begin{aligned} V_v(X) &\triangleq \{w : (w, u) \in E(X), (w, v) \in G \setminus E(X)\} \\ V_u(X) &\triangleq \{w : (w, v) \in E(X), (w, u) \in G \setminus E(X)\} \\ V_{uv}(X) &\triangleq \{w : (w, u), (w, v) \in G \setminus E(X)\}. \end{aligned}$$

So E^{new} has the following distribution conditioned on X :

$$\begin{aligned} \text{For } w \in V_v(X), Y_{(v, w)} &\sim \text{Bern}(p_1); \\ \text{For } w \in V_u(X), Y_{(v, w)} &\sim \text{Bern}(p_1); \\ \text{For } w \in V_{uv}(X), Y_{(v, w)} = Y_{(u, w)} &\sim \text{Bern}(p_0). \end{aligned}$$




Figure 6: Illustration of K_v

Recalling the definition of X^{aux} , its conditional distribution on X is therefore generated by

$$\begin{aligned} \text{For } w \in V_v(X), \text{ sample } w' &\sim \text{Unif}(W_{E(X)}(v, w)), \quad X_{(v,w,w')}^{\text{aux}} \sim \text{Bern}(p_1); \\ \text{For } w \in V_u(X), \text{ sample } w' &\sim \text{Unif}(W_{E(X)}(u, w)), \quad X_{(u,w,w')}^{\text{aux}} \sim \text{Bern}(p_1); \\ \text{For } w \in V_{uv}(X), \text{ sample } w' &\sim \text{Unif}(W_{E(X)}(v, w)), \quad X_{(v,w,w')}^{\text{aux}} \sim \text{Bern}(p_0), \\ &\text{sample } w'' \sim \text{Unif}(W_{E(X)}(u, w)), \quad X_{(u,w,w'')}^{\text{aux}} \sim \text{Bern}(p_0). \end{aligned}$$

7.3.3 Step 3: Adding one triangle at a time

We next aim to describe X^{aux} in terms of single triangle additions. To this end we define a few kernels from $\{0, 1\}^{T_G}$ to the same space. First,

$$K_v(X) : \text{Pick } w \sim \text{Unif}(V_v(X)), w' \sim \text{Unif}(W_{E(X)}(v, w)), \text{ set } X_{(v,w,w')} = 1,$$

as depicted in Figure 6. Similarly,

$$K_u(X) : \text{Pick } w \sim \text{Unif}(V_u(X)), w' \sim \text{Unif}(W_{E(X)}(u, w)), \text{ set } X_{(u,w,w')} = 1,$$

and

$$\begin{aligned} K_{uv}(X) : \text{Pick } w &\sim \text{Unif}(V_{uv}(X)), w' \sim \text{Unif}(W_{E(X)}(v, w)), \text{ set } X_{(v,w,w')} = 1 \\ &\text{Pick } w' \sim \text{Unif}(W_{E(X)}(u, w)), \text{ set } X_{(u,w,w')} = 1. \end{aligned}$$

Then

$$X \vee X^{\text{aux}} \stackrel{d}{=} (K_{uv})^{m_3} (K_u)^{m_2} (K_v)^{m_1} (X),$$

where $m_1 \sim \text{Binomial}(|V_v(X)|, p_1)$, $m_2 \sim \text{Binomial}(|V_u(X)|, p_1)$ and $m_3 \sim \text{Binomial}(|V_{uv}(X)|, p_0)$. We will show that $K_u(X)$, $K_v(X)$ and $K_{uv}(X)$ are all $\tilde{O}(1/\sqrt{n})$ -close to X in total variation distance in Lemmas 7.12 and 7.13 just below.

Lemma 7.12. *We have the bounds $d_{\text{TV}}(K_u(X), X) = O(\log n / (n^4 p'^{7/2}))$ and $d_{\text{TV}}(K_v(X), X) = O(\log n / (n^4 p'^{7/2}))$.*

Lemma 7.13. *We have that $d_{\text{TV}}(K_{uv}(X), X) = O(\log^3 n / (n^6 p'^{11/2}))$.*

We prove Lemma 7.12 in the next subsection and defer the proof of Lemma 7.13 to Appendix B, as the proofs are similar. With these lemmas, we can prove Lemma 7.8.

Proof of Lemma 7.8. From Lemma 4.6, Lemmas 7.12 and 7.13 imply that

$$d_{\text{TV}}((K_{uv})^{m_3}(K_u)^{m_2}(K_v)^{m_1}(X), X) = O\left(\frac{(m_1 + m_2) \log n}{n^4 p'^{7/2}}\right) + O\left(\frac{m_3 \log^3 n}{n^6 p'^{11/2}}\right).$$

Note that $|V_v(X)|, |V_u(X)|, |V_{uv}(X)| \leq n$ and $p_i = \Theta(p')$, so we have $\mathbb{E}[m_i] = O(np')$. Hence from convexity of total variation,

$$d_{\text{TV}}(X, \tilde{X}) = O(np') \cdot O\left(\frac{\log^3 n}{n^6 p'^{11/2}}\right) = O\left(\frac{\log^3 n}{n^5 p'^{9/2}}\right).$$

This completes the proof of Lemma 2.5. □

It remains to prove Lemmas 7.12 and 7.13.

7.4 Proof of Lemma 7.12 via Likelihood Concentration

The definitions of K_u and K_v are symmetric with respect to swapping u and v , so it suffices to prove the claim for one of them, e.g., K_v .

The intuition behind Lemma 7.12 is that K_v adds a triangle to X on v, w, w' in order to add edge (v, w) without introducing other new edges (see Figure 6). Here (u, w) is already in $E(X)$ and $(v, w) \notin E(X)$ (but is in G), and moreover, each of the edges (v, w') and (w, w') are in $E(X)$. Now, if (i) there are $\tilde{\Omega}(n)$ such (v, w) with triangle (v, w, w') already present in X , and (ii) there are $\tilde{\Omega}(n)$ possible positions to add an additional one, then we expect adding the new triangle to induce a total variation $\tilde{O}(1/\sqrt{n})^3$. Item (i) is shown later in Lemma 7.19 and Item (ii) is shown shortly in Lemma 7.14. The conclusion, the $\tilde{O}(1/\sqrt{n})$ total variation bound, is stated as Lemma 7.15.

7.4.1 Conditioning on Non-Incident Triangles

Define a partition of T_G as follows:

$$T_G^0 \triangleq \{t \in T_G : t \text{ does not contain } u \text{ or } v\},$$

$$T_G^v \triangleq \{t \in T_G : t \text{ contains } v\}, \quad \text{and} \quad T_G^u \triangleq \{t \in T_G : t \text{ contains } u\}.$$

Before bounding the total variation difference, we will condition on triangles in T_G^0 . This will be helpful as the concentration results we use (Theorem 9.1) depend on the dimension of the distribution. By conditioning on $X_{T_G^0}$, the dimension of the distribution changes from $|T_G|$, which typically has size $\Theta(n^3)$, to $|T_G^u \cup T_G^v|$, which has size at most $O(n^2)$.

Lemma 7.14. *Let G be c -uniformly 2-star dense. With high probability over the randomness of $X_{T_G^0}$, for any $w \notin \{u, v\}$, there are $\Theta(n^2 p')$ different w' such that $(w', v) \in G$ and $(w', w) \in E(X_{T_G^0})$.*

³This is by the heuristic of computing the shift in mean and comparing to the standard deviation, which for Gaussians and other well-behaved distributions is well-known to give the correct total variation.




Figure 7: Lemma 7.14 states that with high probability, $X_{T_G^0}$ satisfies that for any $w \notin \{u, v\}$, there are $\tilde{\Omega}(n)$ different w' such that $(w', v) \in G$ and $(w', w) \in E(X_{T_G^0})$.

Proof. Fix $w \notin \{u, v\}$. By definition of c -uniformly 2-star density, there are $\Theta(n)$ different w' that both $(w', w), (w', v) \in G$. Each edge (w', w) is contained in $\Omega(n)$ triangles in T_G^0 . Since μ_G is $\Omega(p')$ -marginally large, the number of w' that $(w', w) \in E(X_{T_G^0})$ stochastically dominates $\text{Binomial}(\Omega(n), \Omega(np'))$. So it is $\Omega(n^2 p')$ with high probability by the Chernoff bound. μ_G is also $\Omega(p')$ -marginally small, so by similar argument, the number of w' that $(w', w) \in E(X_{T_G^0})$ is also $O(n^2 p')$ with high probability. \square

We will prove the following lemma, which implies Lemma 7.12.

Lemma 7.15. *For any $X_{T_G^0}$ satisfying the conclusion of Lemma 7.14,*

$$d_{\text{TV}}(K_v(X)|X_{T_G^0}, X|X_{T_G^0}) = O\left(\frac{\log n}{n^4 p^{7/2}}\right) = \tilde{O}\left(\frac{1}{\sqrt{n}}\right).$$

Proof of Lemma 7.12. The $X_{T_G^0}$ satisfying the conclusion of Lemma 7.14 occur with high probability, hence the total variation bound in Lemma 7.15 occurs with high probability, and the result follows by applying Lemma 4.4. \square

In the remainder of Section 7.4 we prove Lemma 7.15.

7.4.2 Lemma 7.15 via Likelihood Ratio Concentration

To use Lemma 4.5 to prove Lemma 7.15, we will show that the likelihood ratio between the two distributions concentrates. We start by determining the likelihood ratio.

Likelihood Ratio Note that the distribution of $X|X_{T_G^0}$ is just $\mu_G(X|X_{T_G^0})$. We can write the density of both distributions.

$$\begin{aligned} & \mathbb{P}(K_v(X) = x | X_{T_G^0}) \\ &= \sum_{\substack{t=(v,w,w'): t \in x \\ w \in V_v(x-t), w' \in W_{E(x-t)}(v,w)}} \frac{1}{|V_v(x-t)|} \frac{1}{|W_{E(x-t)}(v,w)|} \mu_G(X = x-t | X_{T_G^0}), \end{aligned} \quad (21)$$

and

$$\mathbb{P}(X = x|X_{T_G^0}) = \mu_G(X = x|X_{T_G^0}).$$

Here $x - t$ stands for the vector that is the same as x except x_t is set to 0. So the likelihood ratio is

$$\begin{aligned} L(x) &:= \frac{\mathbb{P}(K_v(X) = x|X_{T_G^0})}{\mathbb{P}(X = x|X_{T_G^0})} \\ &= \sum_{\substack{t=(v,w,w'): t \in x \\ w \in V_v(x-t), w' \in W_{E(x-t)}(v,w)}} \frac{1}{|V_v(x-t)|} \frac{1}{|W_{E(x-t)}(v,w)|} \frac{\mu_G(X = x-t)}{\mu_G(X = x)}. \end{aligned} \quad (22)$$

Lemma 7.16 (Concentration of Likelihood Ratio). *For any $X_{T_G^0}$ that satisfies the conclusion of Lemma 7.14, we have that*

$$|L(X) - 1| = O\left(\frac{\log n}{n^4 p'^{7/2}}\right) = \tilde{O}(1/\sqrt{n}),$$

with high probability.

With Lemma 4.5, Lemma 7.16 immediately implies Lemma 7.15. In the remainder of the section, our focus will be on proving Lemma 7.16.

7.4.3 Concentration of Components of Likelihood Ratio

The proof will follow from three concentration statements for factors in the likelihood ratio (22):

- (a) $\frac{\mu_G(X=x-t)}{\mu_G(X=x)}$ is constant with respect to x ,
- (b) $|V_v(X)|$ and $|W_{E(X)}(v, w)|$ concentrate around their mean,
- (c) the total number of terms in (22) also concentrates.

(a) $\frac{\mu_G(X=x-t)}{\mu_G(X=x)}$ is constant with respect to x To start, note that if t is chosen by K_v to be added to X , it always increases the number of edges by 1, i.e., $e(x) - e(x-t) = 1$ if t satisfies the summation in (22). So

$$\frac{\mu_G(X = x-t)}{\mu_G(X = x)} = \frac{p'^{|x-t|} p^{-e(x-t)}}{p'^{|x|} p^{-e(x_0)}} = \frac{p}{p'},$$

which is constant as claimed.

(b) Concentration of $|V_v(X)|$ and $|W_{E(X)}(v, w)|$ By Corollary 6.5, any Lipschitz function of a marginally small random variable concentrates. From the definitions, $|V_v(X)|$ and $|W_{E(X)}(v, w)|$ are $O(1)$ -Lipschitz, yielding the following lemma.

Lemma 7.17. *Conditioned on any $X_{T_G^0}$ satisfying the condition in Lemma 7.14, with high probability,*

$$||V_v(X)| - \mathbb{E}|V_v(X)|| = O(\sqrt{n^2 p'} \log n)$$

and

$$||W_{E(X)}(v, w)| - \mathbb{E}|W_{E(X)}(v, w)|| = O(\sqrt{n^2 p'} \log n).$$

Proof. $\mathcal{L}(X|X_{T_G^0})$ is a conditional distribution of μ_G over $O(n^2)$ triangles. So we can apply Corollary 6.5. With high probability, any $O(1)$ -Lipschitz function is within distance $O(\sqrt{n^2 p' \log n})$ from its mean with high probability. \square

To use the concentration result, we also need to show that the means of $|V_v(X)|$ and $|W_{E(X)}(v, w)|$ are large relative to their variances.

Lemma 7.18. *For any $X_{T_G^0}$ satisfying the condition in Lemma 7.14,*

$$\mathbb{E}[|V_v(X)| | X_{T_G^0}] = \Theta(n^2 p') = \tilde{\Theta}(n),$$

and for any $w \neq u, v$

$$\mathbb{E}[|W_{E(X)}(v, w)| | X_{T_G^0}] = \Theta(n^3 p'^2) = \tilde{\Theta}(n).$$

Proof. Recall that $V_v(X)$ consists of vertices w that $(w, u) \in E(X), (w, v) \notin E(X)$. Fix any $w \in W_{u,v}(G)$. From Lemma 6.2,

$$\mathbb{P}((w, v) \notin E(X)) = \Omega(1 - np') = \Theta(1).$$

By Lemma 8.2,

$$\mathbb{P}((u, w) \in E(X) | (w, v) \notin E(X)) = \Theta(np').$$

So $\mathbb{P}(w \in V_v(X)) = \Theta(np')$. We have $\mathbb{E}[|V_v(X)|] = |W_{u,v}(G)| \cdot \Theta(np') = \Theta(n^2 p')$.

Fix any $w \neq u, v$. Consider a vertex w' such that $(w, w') \in E(X_{T_G^0}), (v, w') \in G$, since G is c -uniformly 2-star dense, (v, w') is included in $\Theta(n)$ triangles, each having probability $\Theta(p')$ of being in X . So

$$\mathbb{P}((v, w') \in E(X) | X_{T_G^0}) = \Theta(np').$$

From the assumption on $X_{T_G^0}$, there are $\Theta(n^2 p')$ such w' , so

$$\mathbb{E}[|W_{E(X)}(v, w)| | X_{T_G^0}] = \Theta(n^2 p') \cdot \Theta(np') = \Theta(n^3 p'^2). \quad \square$$

Note that $W_{E(X)}(v, w) \geq |W_{E(X_{T_G^0})}(v, w)| = \Theta(n)$, so $\mathbb{E}[|W_{E(X)}(v, w)| | X_{T_G^0}] = \Theta(n)$. By Lemma 7.17 and Lemma 7.18, for any v and w , conditioned on any $X_{T_G^0}$,

$$\left| \frac{|V_v(X)| - \mathbb{E}[|V_v(X)|]}{\mathbb{E}[|V_v(X)|]} \right| = O\left(\frac{\log n}{np'^{1/2}}\right) \text{ and } \left| \frac{|W_{E(X)}(v, w)| - \mathbb{E}[|W_{E(X)}(v, w)|]}{\mathbb{E}[|W_{E(X)}(v, w)|]} \right| = O\left(\frac{\log n}{n^2 p'^{3/2}}\right)$$

holds with high probability.

Note $|V_v(X - t)|$ and $|V_v(X)|$ differ by at most one, and the same holds for $|W_{v,w}(X - t)|$ and $|W_{E(X)}(v, w)|$. So with high probability, each term in (22) concentrates,

$$\begin{aligned} & \left| \frac{1}{|V_v(X - t)|} \frac{1}{|W_{v,w}(X - t)|} \frac{\mu_G(X - t)}{\mu_G(X)} - \frac{1}{\mathbb{E}[|V_v(X)|]} \frac{1}{\mathbb{E}[|W_{E(X)}(v, w)|]} \frac{p}{p'} \right| \\ &= O\left(\frac{\log n}{n^2 p'^{3/2}}\right) \cdot \frac{1}{\mathbb{E}[|V_v(X)|]} \frac{1}{\mathbb{E}[|W_{E(X)}(v, w)|]} \frac{p}{p'} = O\left(\frac{\log n}{n^7 p'^{11/2}}\right) = \tilde{O}(n^{-3/2}). \end{aligned} \quad (23)$$




Figure 8: Definition of $S(X)$. A pair w, w' is in $S(X)$ if $t = (v, w, w') \in X, (w, u), (w, w'), (v, w') \in E(X - t), (w, v) \notin E(X - t)$.

(c) Concentration of the Number of Terms in (22). It remains to show that the number of terms also concentrates. Let S be the set of ordered vertices

$$S(X) = \{(w, w') : t = (v, w, w'), t \in X, w \in V_v(X - t), w' \in W_{v,w}(X - t)\},$$

as illustrated in Figure 8. Each triangle (v, w, w') in the summation of (22) corresponds to a pair of vertices (w, w') in $S(X)$.

Lemma 7.19. $|S(X)|$ is $O(1)$ -Lipschitz, and under $\mathcal{L}(X|X_{T_G^0})$,

$$\mathbb{E}[|S(X)| | X_{T_G^0}] = \Theta(n^5 p'^4) = \tilde{\Theta}(n).$$

Proof. From the definition of V_v and $W_{v,w}$, S can be alternatively written as

$$S(X) = \{(w; w') : t = (v, w, w'), t \in X, (w, u), (w, w'), (v, w') \in E(X - t), (w, v) \notin E(X - t)\}.$$

Suppose X and X' differ at one triangle, $X' = X \vee e_{t_0}$, we want to show $S(X')$ and $S(X)$ differ by $O(1)$ elements.

First consider $S(X) \setminus S(X')$. Suppose $(w; w') \in S(X)$ but $(w; w') \notin S(X + t_0)$. Then we must have $(w, v) \notin E(X - t)$ but $(w, v) \in E(X + t_0 - t)$. Since $E(X + t_0 - t)$ and $E(X - t)$ differ by at most 3 edges, we have $|\{w : \exists w', (w; w') \in S(X) \setminus S(X')\}| \leq 3$. Notice that each for each w , at most one w' satisfy $(w; w') \in S(X)$, otherwise $(w, v) \in E(X - t)$. So $|S(X) \setminus S(X')| \leq 3$.

Then consider $S(X') \setminus S(X)$. If $(w; w') \in S(X') \setminus S(X)$, either $(w, u) \in E(X + t_0 - t) - E(X - t)$, $(w, w') \in E(X + t_0 - t) - E(X - t)$ or $(v, w') \in E(X + t_0 - t) - E(X - t)$. Let

$$S_1 = \{(w; w') : (w; w') \in S(X'), (w, u) \in E(X + t_0 - t) - E(X - t)\},$$

$$S_2 = \{(w; w') : (w; w') \in S(X'), (w, w') \in E(X + t_0 - t) - E(X - t)\},$$

and

$$S_3 = \{(w; w') : (w; w') \in S(X'), (v, w') \in E(X + t_0 - t) - E(X - t)\}.$$

Then

$$S(X') \setminus S(X) \subset S_1 \cup S_2 \cup S_3.$$

Still note that for each w , at most one w' satisfies $(w; w') \in S(X')$, and $E(X + t_0 - t) - E(X - t)$ has at most 3 edges, we have $|S_1| \leq 3, |S_2| \leq 3$. For S_3 , notice that each w' has at most one w such

that $(w; w') \in S(X')$ and $(v, w') \in E(X + t_0 - t) - E(X - t)$. Otherwise by definition of $S(X')$, at least two triangles in X contains edge (v, w') . So (v, w') would belong to $E(X - t)$ for any t , contradicting that $(v, w') \in E(X + t_0 - t) - E(X - t)$. Therefore, $|S_3| \leq 3$ also holds. We have $|S(X') \setminus S(X)| \leq 9$.

Next we show that $\mathbb{E}|S(X)| = \Theta(n^5 p'^4)$. It is sufficient to prove there are $\Theta(n^3 p')$ pairs $(w; w')$ that $(w; w')$ is in $S(X)$ with probability $\Theta(n^2 p'^3)$. Since G is a c -uniformly 2-star dense graph, there are $\Theta(n)$ vertices w such that $(w, u), (w, v) \in G$. From the assumption that $|W_{E(X_{T_G^0})}(v, w)| = \Theta(n^2 p')$, there are $\Theta(n^3 p')$ pairs $(w; w')$ that $(w, u), (w, v) \in G$ and $(w, w') \in E(X_{T_G^0})$.

Next we show that for such $(w; w')$, it is in $S(X)$ with probability $\Theta(n^2 p'^3)$. Consider the event that $(w; w') \in S(X)$, i.e.,

$$t = (v, w, w'), t \in X, (w, u), (w, w'), (v, w') \in E(X - t), (w, v) \notin E(X - t).$$

Since $E(X - t) \supset E(X_{T_G^0})$, $(w, w') \in E(X - t)$ always holds. First, $t \in X$ with probability $\Theta(p')$, next consider the event conditioning on this. (w, v) are included in at most $n - 2$ triangles, each being included with probability $\Theta(p')$. $(w, v) \notin E(X - t)$ with probability at least

$$(1 - \Theta(p'))^{n-2} = \Theta(1).$$

Graph G is c -uniformly 2-star dense, so there are $\Theta(n)$ triangles contains (w, u) and (v, w') and do not contain (w, v) . By Lemma 6.1, conditioning on the all the events above, each triangle is included in X with probability $\Theta(p')$, so $(w, u), (v, w') \in E(X - t)$ with probability $\Theta(n^2 p'^2)$. Therefore, the event $(w; w') \in S(X)$ happens with probability $\Theta(n^2 p'^3)$. \square

7.4.4 Proof of Concentration of the Likelihood Ratio

Proof of Lemma 7.16. Because $S(X)$ is Lipschitz by Lemma 7.19, Corollary 6.5 applies to yield $||S(X)| - \mathbb{E} S(X)| = \tilde{O}(\sqrt{n})$ with high probability. So $S(X) = \tilde{O}(n)$ with high probability.

Now for any w ,

$$\frac{1}{\mathbb{E}|V_v(X)|} \frac{1}{\mathbb{E}|W_{E(X)}(v, w)|} \frac{p}{p'} = \Theta\left(\frac{1}{n^5 p'^4}\right),$$

so again because $S(X)$ is Lipschitz we have that the following function is $O(1/(n^5 p'^4))$ -Lipschitz, as each term in the right hand side is $O(1/(n^5 p'^4))$:

$$f(x) \triangleq \sum_{(w; w') \in S(x)} \frac{1}{\mathbb{E}|V_v(X)|} \frac{1}{\mathbb{E}|W_{E(X)}(v, w)|} \frac{p}{p'}.$$

Again by Corollary 6.5, $f(X)$ concentrates around its mean with distance at most $O(\log n / (n^4 p'^{7/2}))$ with high probability.

Now combining this with (23), we get for any fixed $X_{T_G^0}$ that satisfies the conclusion of Lemma 7.14, the likelihood ratio

$$L(x) = \frac{\mathbb{P}(K_v(X) = x | X_{T_G^0})}{\mathbb{P}(X = x | X_{T_G^0})}$$

concentrates as

$$|L(X) - \mathbb{E} f(X)| = O\left(\frac{\log n}{n^4 p'^{7/2}}\right) + O\left(|S(X)| \frac{\log n}{n^7 p'^{11/2}}\right).$$

This is $O(\log n/(n^4 p'^{7/2})) = \tilde{O}(1/\sqrt{n})$ with high probability since $|S(X)| = \Theta(n^5 p'^4)$ with high probability.

We can use Lemma 4.7 (which states that concentration implies concentration about the mean for bounded random variables), as it is not hard to check that the likelihood ratio is always bounded by $\tilde{O}(n^3)$. Hence, with high probability,

$$|L(X) - \mathbb{E} L(X)| = O\left(\frac{\log n}{n^4 p'^{7/2}}\right) = \tilde{O}(1/\sqrt{n}).$$

This completes the proof of Lemma 7.16, using that $\mathbb{E} L(X) = 1$. \square

7.5 The class of graphs \mathcal{G} has high-probability

The goal of this section is to show Lemma 2.4, i.e.,

$$\mathbb{P}(\text{RGT}_i \in \mathcal{G}) = 1 - o(1).$$

Recall the definition of \mathcal{G} in Section 2.3, $\mathcal{G} = \mathcal{G}_1 \cap \mathcal{G}_2$. Here \mathcal{G}_1 is the set of graphs that are $p^2/3$ -uniformly 2-star dense and \mathcal{G}_2 is the set of graphs that satisfy

$$|\mu_{G+e}(Y_e = 1) - p_e^+| = C_p n^{-5/2} p'^{-2} \sqrt{\log n}, \quad \text{where } p_e^+ = \mathbb{E}_{G' \sim \text{RGT}(n, p, p')} [\mu_{G'+e}(Y_e = 1)],$$

where C_p is a large enough constant depending on p . We can prove that \mathcal{G}_2 is a high-probability event under $\text{RGT}(n, p, p')$ through concentration.

Lemma 7.20. *If $G \sim \text{RGT}(n, p, p')$, then $\mathbb{P}(G \in \mathcal{G}_2) = 1 - o(1)$.*

Let us first show why Lemma 2.4 holds given Lemma 7.20.

Proof of Lemma 2.4. First, RGT_0 is $p^2/2$ -uniformly 2-star dense by Lemma 5.2 with probability $1 - e^{-\Omega(n)}$. For any RGT_i , $i \in [K]$, and any $j, k \in [n]$, the number of vertices l such that (j, l) or (k, l) are different from RGT_0 is at most K . So when $K = \tilde{O}(n^{1/4})$, RGT_i is $p^2/3$ -uniformly 2-star dense if RGT_0 is $p^2/2$ -uniformly 2-star dense. Thus, $\text{RGT}_i \in \mathcal{G}_1$ for all $i \in [K]$ with high probability.

Lemma 7.20 shows that $\text{RGT}_0 \in \mathcal{G}_2$ with high probability. If $\text{RGT}_i \in \mathcal{G}_1$ for any i , then by Corollary 6.4,

$$|\mu_{\text{RGT}_i+e}(Y_e = 1) - \mu_{\text{RGT}_0+e}(Y_e = 1)| = \tilde{O}(i/n).$$

Therefore, with high probability,

$$|\mu_{\text{RGT}_i+e}(Y_e = 1) - \mathbb{E}_{G' \sim \text{RGT}(n, p, p')} [\mu_{G'+e}(Y_e = 1)]| = \tilde{O}(1/\sqrt{n}) + \tilde{O}(k^2/n) = \tilde{O}(1/\sqrt{n}). \quad \square$$

Next we prove Lemma 7.20. The idea is to use McDiarmid's inequality and the fact that the marginal probability of $\mathcal{L}(Y)$ is a Lipschitz function due to bounded influence.

Proof of Lemma 7.20. Since $\text{RGT}(n, p, p')$ is generated by adding triangles to random graphs, we can view it as a function over the original random graph and the triangles added. Let $x \sim \text{Bern}(p)^{\otimes \binom{[n]}{2}}$ be the indicator of edges in $G(n, p)$ and $y \sim \text{Bern}(p')^{\otimes \binom{[n]}{3}}$ be the indicator of whether

a triangle is added to the graph. We use $\text{RGT}(x, y)$ to denote the graph generated by choice x and y . Then $\mu_{G+e}(Y_e = 1)$ can be viewed as a function of (x, y) ,

$$f(x, y) \triangleq \mu_{\text{RGT}(x, y)+e}(Y_e = 1).$$

By Lemma 2.8, $\mu_{G+e}(Y_e = 1)$ is Lipschitz as a function of a c -uniformly 2-star dense graph G . Indeed, for any $e' \notin G$,

$$\begin{aligned} \mu_{G+e+e'}(Y_e = 1) &= \mu_{G+e+e'}(Y_{e'} = 1) \mu_{G+e+e'}(Y_e = 1 | Y_{e'} = 1) \\ &\quad + \mu_{G+e+e'}(Y_{e'} = 0) \mu_{G+e+e'}(Y_e = 1 | Y_{e'} = 0). \end{aligned}$$

As $\mu_{G+e+e'}(Y_e = 1 | Y_{e'} = 0) = \mu_{G+e}(Y_e = 1)$, we have

$$\begin{aligned} |\mu_{G+e+e'}(Y_e = 1) - \mu_{G+e}(Y_e = 1)| &\leq |\mu_{G+e+e'}(Y_e = 1 | Y_{e'} = 1) - \mu_{G+e+e'}(Y_e = 1 | Y_{e'} = 0)| \\ &\leq \mathbf{I}_{e' \rightarrow e}^M. \end{aligned}$$

So for any c -uniformly 2-star dense graph G , and edge e' incident to e ,

$$|\mu_{G+e+e'}(Y_e = 1) - \mu_{G+e}(Y_e = 1)| = O(1/(n^3 p'^2)) = \tilde{O}(1/n),$$

and for other edge e' not incident to e ,

$$|\mu_{G+e+e'}(Y_e = 1) - \mu_{G+e}(Y_e = 1)| = O(1/(n^3 p')) = \tilde{O}(1/n^2).$$

This implies that if $\text{RGT}(x, y)$ is a c -uniformly 2-star dense graph, we have

$$\left| \frac{\partial f}{\partial x_{e'}}(x, y) \right| = \begin{cases} \tilde{O}(1/n) & \text{if } e' \cap e \neq \emptyset \\ \tilde{O}(1/n^2) & \text{otherwise} \end{cases},$$

and

$$\left| \frac{\partial f}{\partial y_t}(x, y) \right| = \begin{cases} \tilde{O}(1/n) & \text{if } t \cap e \neq \emptyset \\ \tilde{O}(1/n^2) & \text{otherwise} \end{cases}.$$

Denote the quantities on the right hand side by $c_{e'}$ and c_t . To use Lemma 4.10, we need to check that f is Lipschitz. For any $x_1, x_2 \in \{0, 1\}^{\binom{[n]}{2}}$ and $y_1, y_2 \in \{0, 1\}^{\binom{[n]}{3}}$ such that $\text{RGT}(x_1, y_1)$ and $\text{RGT}(x_2, y_2)$ are c -uniformly 2-star dense graphs,

$$\begin{aligned} &|f(x_1, y_1) - f(x_2, y_2)| \\ &\leq |f(x_1, y_1) - f(x_1 \vee x_2, y_1 \vee y_2)| + |f(x_2, y_2) - f(x_1 \vee x_2, y_1 \vee y_2)| \\ &\leq \sum_{e': (x_1)_{e'} \neq (x_1 \vee x_2)_{e'}} c_{e'} + \sum_{t: (y_1)_t \neq (y_1 \vee y_2)_t} c_t + \sum_{e': (x_2)_{e'} \neq (x_1 \vee x_2)_{e'}} c_{e'} + \sum_{t: (y_2)_t \neq (y_1 \vee y_2)_t} c_t \\ &= \sum_{e': (x_1)_{e'} \neq (x_2)_{e'}} c_{e'} + \sum_{t: (y_1)_t \neq (y_2)_t} c_t. \end{aligned}$$

Here we used that $\text{RGT}(x_1 \vee x_2, y_1 \vee y_2) = \text{RGT}(x_1, y_1) \cup \text{RGT}(x_2, y_2)$ is c -uniformly 2-star dense.

We can set \mathcal{Y} in Lemma 4.10 to be the set of x, y such that $\text{RGT}(x, y)$ is c -uniformly 2-star dense. The probability of $G \sim \text{RGT}(n, p, p')$ being not c -uniformly 2-star dense is bounded by $e^{-\Omega(n)}$ by Lemma 5.2. So for $t > 0$, we have

$$\mathbb{P}_G \left(\left| \mu_{G+e}(Y_e = 1) - \mathbb{E}_{G' \sim \text{RGT}(n, p, p')} \mu_{G'+e}(Y_e = 1) \right| > t \right) \leq e^{-\Omega(n)} + \exp \left(\frac{-(t - e^{-\Omega(n)})^2}{\sigma^2 + \tilde{O}(t/n)} \right). \quad (24)$$

Here

$$\begin{aligned} \sigma^2 &= p(1-p) \left(\sum_{e' \in \binom{[n]}{2}: e' \text{ shares node with } e} O\left(\frac{1}{n^6 p'^4}\right) + \sum_{e' \in \binom{[n]}{2}: e' \text{ does not node with } e} O\left(\frac{1}{n^6 p'^2}\right) \right) \\ &\quad + p'(1-p') \left(\sum_{t \in \binom{[n]}{3}: t \text{ shares nodes with } e} \left(\frac{1}{n^6 p'^4}\right) + \sum_{t \in \binom{[n]}{3}: t \text{ does not node with } e} O\left(\frac{1}{n^6 p'^2}\right) \right) \\ &= p(1-p) \left(n \cdot O\left(\frac{1}{n^6 p'^4}\right) + n^2 \cdot O\left(\frac{1}{n^6 p'^2}\right) \right) + p'(1-p') \left(n^2 \cdot O\left(\frac{1}{n^6 p'^4}\right) + n^3 \cdot O\left(\frac{1}{n^6 p'^2}\right) \right) \\ &= O(n^{-5} p'^{-4}). \end{aligned}$$

We can choose $t = \Theta(n^{-5/2} p'^{-2} \sqrt{\log n})$ so that the right hand side of (24) is $1/n^c$ for a large enough constant c . \square

8 Low Influences for Marginally Small Distributions

The distribution μ_G studied in the paper, as proven in Lemma 6.1, is marginally small (Defn. 7). In this section, we carry out a general investigation of marginally small distributions and show, under an assumption on the underlying graphical model, that they have low influences. The influence bound on distribution $\mathcal{L}_G(Y)$ is then derived as a result of low influence in μ_G .

We start by formally defining influence.

Definition 10 (Influence). For a binary distribution over $\{0, 1\}^{\mathcal{X}}$, the *influence* of $S \subset \mathcal{X}$ on $S' \subset \mathcal{X}$ is defined by

$$\mathbf{I}_{S \rightarrow S'} = \sup_{\substack{x_S^a, x_S^b \in \{0, 1\}^S \\ \omega \in \{0, 1\}^A, A \subset \mathcal{X} - S - S'}} d_{\text{TV}}(P_{x_{S'} | x_S^a, \omega}, P_{x_{S'} | x_S^b, \omega}),$$

where ω represents an arbitrary fixing (“pinning”) of variables in A .

For $x, y \in \mathcal{X}$, use $\mathbf{I}_{x \rightarrow y}$ to denote $\mathbf{I}_{\{x\} \rightarrow \{y\}}$.

Remark 15. Definition 10 is different from the marginal influence defined in Definition 9 in the sense that it allows arbitrary pinning of the variables.

Remark 16. The notion of influence plays a prominent role in the line of work using spectral independence to bound the mixing time of Markov chains, examples include [ALG21, CGŠV21]. The influence used in these works are usually the marginal influence, but a separate influence defined for each possible pinning. We instead maximize over all possible pinnings, which turns out to be more convenient. Another notion of influence appears in Dobrushin uniqueness conditions [Dob70], where the pinning is over $\mathcal{X} - S - S'$, not over an arbitrary subset as in Defn. 10.

8.1 Low Influences for Marginally Small Distributions

Theorem 8.1. *Let P be a distribution over $\{0,1\}^N$, and G be the graphical model of P with maximum degree Δ . Suppose that P is q -marginally small where $q < \frac{1}{2\Delta}$ and let*

$$C = \sup_k \sup_{i,j:d(i,j)=k} |N_{k-1}(i) \cap N_1(j)|.$$

Then for any $S \subset [N]$ and $i \in [N]$ such that $d(S,i) \geq d$, we can bound the influence of S on i as

$$\mathbf{I}_{S \rightarrow i} \leq |S| \left(\frac{2Cq}{1-2q\Delta} \right)^d.$$

In particular, if $C = O(1)$ and $1 - 2q\Delta = \Omega(1)$, then $\mathbf{I}_{j \rightarrow i} = O(q^{d(i,j)})$.

Proof. Noticing that all properties of P are maintained under arbitrary conditioning, it is without loss of generality to prove

$$\mathbf{I}_{S \rightarrow i}^M \leq |S| \left(\frac{2Cq}{1-2q\Delta} \right)^d.$$

Let us fix vertex i , consider the possible influence on i by a set S that $d(S,i) \geq d$ and also $|S \cap N_d(i)| \leq k$. Define

$$I_d(k) = \sup_{d(S,i) \geq d, |S \cap N_d(i)| \leq k} \mathbf{I}_{S \rightarrow i}^M.$$

We will show that $I_d(k) \leq k \left(\frac{2Cq}{1-2q\Delta} \right)^d$ by induction on d .

First, note that P is q -marginally bounded, by definition of the influence, $\mathbf{I}_{\sim i \rightarrow i} \leq q$. So $I_1(k) \leq q \leq \frac{2kCq}{1-2q\Delta}$ for arbitrary k .

Now assume $I_{d-1}(k) \leq k \left(\frac{2Cq}{1-2q\Delta} \right)^{d-1}$. Suppose the supremum in $I_d(k)$ is given by set S and two marginal conditions x_S^a and x_S^b , so that

$$I_d(k) = d_{\text{TV}}(P_{x_i|x_S}(\cdot | x_S^a), P_{x_i|x_S}(\cdot | x_S^b)).$$

Let $X^a \sim P_{x_{N_1(S)}|x_S}(\cdot | x_S^a)$ and $X^b \sim P_{x_{N_1(S)}|x_S}(\cdot | x_S^b)$ be independent of each other. Note that G is the graphical model of P , so x_i and x_S are independent of each other conditioned on $x_{N_1(S)}$. This means

$$P_{x_i|x_S}(\cdot | x_S^a) = \mathbb{E} [P_{x_i|x_S, x_{N_1(S)}}(\cdot | x_S^a, X^a)] = \mathbb{E} [P_{x_i|x_{N_1(S)}}(\cdot | X^a)].$$

Similarly, $P_{x_i|x_S}(\cdot | x_S^b) = \mathbb{E} [P_{x_i|x_{N_1(S)}}(\cdot | X^b)]$. Then by convexity of total variation (Lemma 4.2), we have

$$d_{\text{TV}}(P_{x_i|x_S}(\cdot | x_S^a), P_{x_i|x_S}(\cdot | x_S^b)) \leq \mathbb{E} [d_{\text{TV}}(P_{x_i|x_{N_1(S)}}(\cdot | X^a), P_{x_i|x_{N_1(S)}}(\cdot | X^b))].$$

By the induction assumption,

$$\begin{aligned}
& \mathbb{E} [d_{\text{TV}}(P_{x_i|x_{N_1(S)}}(\cdot | X^a), P_{x_i|x_{N_1(S)}}(\cdot | X^b))] \\
& \leq \sum_{j \geq 1} \mathbb{P}_{X^a, X^b}(\|X_{N_{d-1}(i)}^a - X_{N_{d-1}(i)}^b\|_1 = j) I_{d-1}(j) \\
& \quad + \sum_{j \geq 1} \mathbb{P}_{X^a, X^b}(X_{N_{d-1}(i)}^a = X_{N_{d-1}(i)}^b, \|X_{N_d(i)}^a - X_{N_d(i)}^b\|_1 = j) I_d(j) \\
& \leq \sum_{j \geq 1} \mathbb{P}_{X^a, X^b}(\|X_{N_{d-1}(i)}^a\|_1 + \|X_{N_{d-1}(i)}^b\|_1 = j) \cdot j \cdot \left(\frac{2Cq}{1-2q\Delta}\right)^{d-1} \\
& \quad + \sum_{j \geq 1} \mathbb{P}_{X^a, X^b}(\|X_{N_d(i)}^a\|_1 + \|X_{N_d(i)}^b\|_1 = j) I_d(j).
\end{aligned}$$

Because $d(S, i) = d$, $|N_1(S) \cap N_{d-1}(i)| \leq |S \cap N_d(i)|C \leq kC$, and P is q -marginally small, the distribution of $\|X_{N_{d-1}(i)}^a\|_1 + \|X_{N_{d-1}(i)}^b\|_1$ is stochastically dominated by $\text{Binomial}(2kC, q)$. Similarly, $|N_1(S)| \leq k\Delta$, so $\|X_{N_d(i)}^a\|_1 + \|X_{N_d(i)}^b\|_1 \leq \|X^a\|_1 + \|X^b\|_1$ is stochastically dominated by $\text{Binomial}(2k\Delta, q)$. We have

$$\begin{aligned}
I_d(k) & \leq \left(\frac{2Cq}{1-2q\Delta}\right)^{d-1} \sum_{j=0}^{2ka} \text{Binomial}(2kC, q, j) \cdot j + \sum_{j=0}^{2k\Delta} \text{Binomial}(2k\Delta, q, j) I_d(j) \\
& = 2kCq \left(\frac{2Cq}{1-2q\Delta}\right)^{d-1} + \sum_{j=0}^{2k\Delta} \text{Binomial}(2k\Delta, q, j) I_d(j).
\end{aligned}$$

For simplicity of notation, let $\tilde{I}_d(k) = \frac{(1-2q\Delta)^{d-1}}{(2Cq)^d} I_d(k)$. We get a recursion

$$\tilde{I}_d(k) \leq k + \sum_{j=0}^{2k\Delta} \text{Binomial}(2k\Delta, q, j) \tilde{I}_d(j).$$

We compare this recursion with percolation on a tree with branching factor 2Δ and branching probability q . Let m_k be the total number of nodes in k such independent trees. We have

$$\mathbb{E} m_k = k + \sum_{j=0}^{2k\Delta} \text{Binomial}(2k\Delta, q, j) \mathbb{E} m_j.$$

By comparing both recursions, we can couple the recursion of $\tilde{I}_d(k)$ with the branching process and conclude that $\tilde{I}_d(k) \leq \mathbb{E} m_k$. By linearity of expectation, $\mathbb{E} m_k = \frac{k}{1-2q\Delta}$. So from the definition of $\tilde{I}_d(k)$, we have

$$I_d(k) \leq k \left(\frac{2Cq}{1-2q\Delta}\right)^d. \quad \square$$

Corollary 8.1. *Let μ_G be the distribution in Definition 5. For any two triangles t and t' in graph G such that $d(t, t') \geq k$,*

$$\mathbf{I}_{t \rightarrow t'} = O(p'^k).$$

Proof. We only need to verify that μ_G satisfies all the properties in Theorem 8.1. In the graphical model of μ_G , any two triangles are connected if and only if they share an edge. So the maximum degree of the graphical model of μ_G is at most $3(n-2)$. From Lemma 6.1, μ_G is $O(p')$ -marginally small. Because $p' \ll 1/n$, we have that $p' < 1/2n$.

By induction one can verify that two triangles have distance k if and only if they share $3-k$ common vertices. So for two triangles t, t' that $d(t, t') = 2$, t and t' share 1 vertex. $N_1(t) \cap N_1(t')$ are the vertices that share 2 vertices with both t and t' . Therefore, $|N_1(t) \cap N_1(t')| = 2$.

Suppose two triangles t, t' satisfy $d(t, t') = 3$, i.e., t and t' do not share vertex. $N_2(t) \cap N_1(t')$ are the triangles that share 1 vertex with t and 2 vertices with t' . There are 3 of them in total.

We have

$$\sup_k \sum_{t, t': d(t, t')=k} |N_{k-1}(t) \cap N_1(t')| \leq 3. \quad \square$$

8.2 Influences in $\mathcal{L}_G(Y)$

In this section, we will apply the conclusions from the previous section about μ_G to bound the influence of $\mathcal{L}_G(Y)$. We start with a technical lemma on the marginals of μ_G and $\mathcal{L}_G(Y)$.

Lemma 8.2. *Suppose G is a c -uniformly 2-star dense graph. Let $A \subset E(G)$ be an edge set that $|A| \leq cn/4$ and $T \subset T_G$ be a triangle set that $|T| \leq cn/4$. Let $y_A \in \{0, 1\}^A$, $x_T \in \{0, 1\}^T$ be arbitrary consistent configurations on A and T . For any edge $e \notin A$,*

$$\mu_G(Y_e = 1 | Y_A = y_A, X_T = x_T) = \Omega(np'),$$

If further $T_e \cap T = \emptyset$ or $x_{T_e \cap T} = \vec{0}$, then

$$\mu_G(Y_e = 1 | Y_A = y_A, X_T = x_T) = \Theta(np').$$

For any triangle $t \notin T$ such that $(y_A)_{E^} = \vec{1}$ where $E^* = E(t) \cap A$, we have*

$$\mu_G(X_t = 1 | Y_A = y_A, X_T = x_T) = O\left(\frac{1}{n^3 p'^2}\right).$$

Proof. For the first statement, consider all triangles in G that contains e , denoted by T_e . If $T_e \cap T \neq \emptyset$ and $x_{T_e \cap T} \neq \vec{0}$, then one of the triangles in T that contain e is already in x_T , which means

$$\mu_G(Y_e = 1 | Y_A = y_A, X_T = x_T) = 1.$$

Now we only consider the case when $T_e \cap T = \emptyset$ or $x_{T_e \cap T} = \vec{0}$. As G is c -uniformly 2-star dense, there are at least $c(n-2)$ such triangles. Let $T^* \subset T_e$ be the set of triangles that does not contain any edge from A or any triangles in T . We have $|T^*| \geq c(n-2) - cn/4 - cn/4 = \Theta(n)$. Suppose $T^* = \{t_1, \dots, t_k\}$, $k = |T^*|$. We have

$$\begin{aligned} \mu_G(Y_e = 1 | Y_A = y_A, X_T = x_T) &= \sum_{i=1}^k \mu_G(X_{t_i} = 1 | Y_A = y_A, X_T = x_T, X_{t_j} = 0, \forall j < i) \\ &= \Theta(kp') = \Theta(np'). \end{aligned}$$

The first equality is because μ_G is $\Omega(p')$ -marginally large, and the conditioning is an event on triangles other than X_{t_i} .

Now let's look at the second statement.

$$\mu_G(X_t = 1 | Y_A = y_A, X_T = x_T) = \frac{\mu_G(X_t = 1 | Y_{A \setminus E^*} = (y_A)_{A \setminus E^*}, X_T = x_T)}{\mu_G(Y_{E^*} = \vec{1} | Y_{A \setminus E^*} = (y_A)_{A \setminus E^*}, X_T = x_T)}.$$

Let $E^* = \{e_1, \dots, e_k\}$, $k \leq 3$. From the first statement, the denominator can be bounded as follows,

$$\begin{aligned} \mu_G(Y_{E^*} = \vec{1} | Y_{A \setminus E^*} = (y_A)_{A \setminus E^*}, X_T = x_T) \\ = \prod_{i=1}^k \mu_G(Y_{e_i} = 1 | Y_{A \setminus E^*} = (y_A)_{A \setminus E^*}, X_T = x_T, Y_j = 1, \forall j < i) \\ = \Omega((np')^k) = \Omega(n^3 p'^3). \end{aligned}$$

Since μ_G is $O(p')$ -marginally small, the nominator is bounded by $O(p')$. We have

$$\mu_G(X_t = 1 | Y_A = y_A, X_T = x_T) = O(1/(n^3 p'^2)). \quad \square$$

Lemma 6.3. *Suppose G is a c -uniformly 2-star dense graph for a constant $c > 0$. Consider distribution $\mathcal{L}_G(Y)$ as in Definition 5. For $A \subset E(G)$ with $|A| \leq cn/4$ and any edge $e \notin A$,*

$$\mathbf{I}_{A \rightarrow e}^M = O(|A|/(n^3 p'^2)) = \tilde{O}(|A|/n).$$

For any pair of edges $e, e' \in E(G)$ that do not share common nodes,

$$\mathbf{I}_{e' \rightarrow e}^M = O(1/(n^3 p')) = \tilde{O}(1/n^2).$$

Proof. The proof idea is to consider edge variable Y as function of the triangle variable X , and use Lemma C.2 to bound influence between one pair edges as the sum of influences between many pairs of triangles.

Let T_A be the set of triangles that contain edges in A , T_e be the set of triangles that contain edges in e . Let $T_o = T_A \cap T_e$. Since any triangle in T_o have at least one edge in A , $|T_o| \leq |A|$.

Suppose $X \sim \mu_G$, Y be the corresponding edge distribution, $\mathcal{L}(Y) = \mu_G$. Suppose $\mathbf{I}_{A \rightarrow e}^M$ is given by conditioning y_A^a and y_A^b , i.e.,

$$\mathbf{I}_{A \rightarrow e}^M = d_{\text{TV}}(\mu_G(Y_e | Y_A = y_A^a), \mu_G(Y_e | Y_A = y_A^b)).$$

From Lemma 4.4,

$$\begin{aligned} \mathbf{I}_{A \rightarrow e}^M &\leq 2(\mu_G(X_{T_o} \neq \vec{0} | Y_A = y_A^a) + \mu_G(X_{T_o} \neq \vec{0} | Y_A = y_A^b)) \\ &\quad + d_{\text{TV}}(\mu_G(Y_e | Y_A = y_A^a, X_{T_o} = \vec{0}), \mu_G(Y_e | Y_A = y_A^b, X_{T_o} = \vec{0})). \end{aligned}$$

From Lemma 8.2, we can bound the first two terms.

$$\mu_G(X_{T_o} \neq \vec{0} | Y_A = y_A^a) \leq \sum_{t \in T_o} \mu_G(X_t = 1 | Y_A = y_A^a) = O(|A|/(n^3 p'^2)),$$

and similarly, $\mu_G(X_{T_o} \neq \vec{0} | Y_A = y_A^b) = O(|A|/(n^3 p'^2))$.

Let $X^a \sim X|Y_A = y_A^a, X_{T_o} = \vec{0}$, $X^b \sim X|Y_A = y_A^b, X_{T_o} = \vec{0}$ be independent, $T_C = \{t|t \in T_A, X_t^a = 1 \text{ or } X_t^b = 1\}$ we can use Lemma 4.2,

$$\begin{aligned}
\mathbf{I}_{A \rightarrow e}^M &\leq \mathbb{E}_{X_{T_A}^a, X_{T_A}^b} d_{\text{TV}}(\mu_G(Y_e|X_{T_A}^a, X_{T_o} = \vec{0}), \mu_G(Y_e|X_{T_A}^b, X_{T_o} = \vec{0})) + O(|A|/(n^3 p'^2)) \\
&\leq \mathbb{E}_{X_{T_A}^a, X_{T_A}^b} d_{\text{TV}}(\mu_G(X_{T_e}|X_{T_A}^a, X_{T_o} = \vec{0}), \mu_G(X_{T_e}|X_{T_A}^b, X_{T_o} = \vec{0})) + O(|A|/(n^3 p'^2)) \\
&\leq \mathbb{E}_{X_{T_A}^a, X_{T_A}^b} \mathbf{I}_{\{t|X_t^a \neq X_t^b\} \rightarrow T_e} + O(|A|/(n^3 p'^2)) \\
&\leq \mathbb{E}_{X_{T_A}^a, X_{T_A}^b} \mathbf{I}_{T_C \rightarrow T_e} + O(|A|/(n^3 p'^2)).
\end{aligned}$$

By Lemma C.2, $\mathbf{I}_{T_C \rightarrow T_e} \leq \sum_{t \in T_e} \mathbf{I}_{T_C \rightarrow t}$.

Now we need to bound $\mathbf{I}_{T_C \rightarrow t}$ for $t \in T_e$. From Corollary 8.1, if T_C and t are adjacent then $\mathbf{I}_{T_C \rightarrow t} = O(p')$, otherwise, $\mathbf{I}_{T_C \rightarrow t} = O(|T_C|p'^2)$. So

$$\sum_{t \in T_e} \mathbf{I}_{T_C \rightarrow t} = O(|N_1(T_C) \cap T_e|p' + |T_e||T_C|p'^2).$$

By the definition of X^a, X^b and T_C , $T_C \subset T_A \setminus T_e$. So for any $t \in T_C$, t share at most 1 vertex with e . There are at most 2 triangles that contains edge e and is adjacent to t . Therefore, $|N_1(T_C) \cap T_e| \leq 2|T_C|$. We have

$$\sum_{t \in T_e} \mathbf{I}_{T_C \rightarrow t} = O(|T_C|p').$$

Here we used that $|T_e| \leq n - 2$, $np' = o(1)$.

Now we have $\mathbf{I}_{A \rightarrow e}^M = O(p' \mathbb{E}_{X_{T_A}^a, X_{T_A}^b} |T_C| + |A|/(n^3 p'^2))$, and it remains to bound $\mathbb{E}_{X_{T_A}^a, X_{T_A}^b} |T_C|$. By linearity of expectation,

$$\mathbb{E}_{X_{T_A}^a, X_{T_A}^b} |T_C| \leq \sum_{t \in T_A} (\mu_G(X_t^a = 1|Y_A = y_A^a, X_{T_o} = \vec{0}) + \mu_G(X_t^b = 1|Y_A = y_A^a, X_{T_o} = \vec{0})).$$

By Lemma 8.2, each term on the right hand side is $O(1/n^3 p'^2)$. So

$$\mathbb{E}_{X_{T_A}^a, X_{T_A}^b} |T_C| = O\left(\frac{|T_A|}{n^3 p'^2}\right) = O\left(\frac{|A|}{n^2 p'^2}\right)$$

and hence $\mathbf{I}_{A \rightarrow e}^M = O(|A|/(n^3 p'^2))$.

Let T_e and $T_{e'}$ be the set of triangles that contains e and e' respectively. Note that e and e' do not share common nodes, so T_e and $T_{e'}$ are disjoint. Similar to the previous argument, suppose

$$\mathbf{I}_{e' \rightarrow e}^M = d_{\text{TV}}(\mu_G(Y_e|Y_{e'} = 1), \mu_G(Y_e|Y_{e'} = 0)).$$

Let $X^a \sim X|Y_{e'} = 1$, $X^b \sim X|Y_{e'} = 0$ be independent, $T_C = \{t|t \in T_{e'}, X_t^a = 1\}$, we have

$$\mathbf{I}_{e' \rightarrow e}^M \leq \mathbb{E}_{X^a} \mathbf{I}_{T_C \rightarrow T_e}.$$

Again by Lemma C.2, $\mathbf{I}_{T_C \rightarrow T_e} \leq \sum_{t_c \in T_C, t_e \in T_e} \mathbf{I}_{t_c \rightarrow t_e}$. Now let's bound $\mathbf{I}_{t_c \rightarrow t_e}$ for possible pairs of t_e and t_c . From Corollary 8.1, this is bounded by $p'^{d(t_c, t_e)}$.

Let T_o be the set of triangles

$$T_o = \{t : t \in T_{e'}, d(t, T_e) = 1\}.$$

This is equivalent to triangles that contains e' and share a node with e . There are at most 2 such triangles. For $t_c \in T_o$, there is at most 1 triangle in T_e with distance 1, so

$$\sum_{t_e \in T_e} \mathbf{I}_{t_c \rightarrow t_e} \leq O(p' + np'^2) = O(p').$$

For $t_c \in T_{e'} \setminus T_o$, there is at most 1 triangle in T_e with distance 2, so

$$\sum_{t_e \in T_e} \mathbf{I}_{t_c \rightarrow t_e} \leq O(p'^2 + np'^3) = O(p'^2).$$

Combining the two cases, we get

$$\mathbf{I}_{T_C \rightarrow T_e} = O(\mathbf{1}\{T_C \cap T_o \neq \emptyset\}p' + |T_C|p'^2).$$

Now we have $\mathbf{I}_{e' \rightarrow e}^M = O(p'^2 \mathbb{E}|T_C| + p' \mathbb{P}(T_C \cap T_o \neq \emptyset))$, where the randomness is from X^a . From Lemma 8.2, we have

$$\mathbb{E}|T_C| = \sum_{t \in T_{e'}} \mu_G(X_t = 1 | Y_{e'} = 1) = O(1/(n^2 p'^2)),$$

and

$$\mathbb{P}(T_C \cap T_o \neq \emptyset) \leq \sum_{t \in T_o} \mu_G(X_t = 1 | Y_{e'} = 1) = O(1/(n^3 p'^2)).$$

Combining the above equations, we have $\mathbf{I}_{e' \rightarrow e}^M = O(1/(n^3 p'))$. □

Corollary 6.4. *For any c -uniformly 2-star dense graph G and any $e \neq e'$,*

$$|\mu_G(Y_e = 1) - \mu_{G+e'}(Y_e = 1)| = \tilde{O}(1/n).$$

Proof. If $e' \in G$, this is trivial. So assume $e' \notin G$. Because

$$\mu_{G+e'}(Y_e = 1) = \mu_{G+e'}(Y_{e'} = 1)\mu_{G+e'}(Y_e = 1 | Y_{e'} = 1) + \mu_{G+e'}(Y_{e'} = 0)\mu_{G+e'}(Y_e = 1 | Y_{e'} = 0),$$

and $\mu_G(Y_e = 1) = \mu_{G+e'}(Y_e = 1 | Y_{e'} = 0)$, we have

$$|\mu_G(Y_e = 1) - \mu_{G+e'}(Y_e = 1)| \leq |\mu_{G+e'}(Y_e = 1 | Y_{e'} = 1) - \mu_{G+e'}(Y_e = 1 | Y_{e'} = 0)| \leq \mathbf{I}_{e' \rightarrow e}^M. \quad \square$$

9 Concentration and Mixing for Marginally Small Distributions

In this section we continue our study of marginally small distributions, and show concentration of Lipschitz functions and fast mixing of Glauber dynamics.

9.1 Glauber Dynamics and Path Coupling

Glauber Dynamics is a simple and commonly used discrete-time Markov chain to sample from distributions with local structures. Consider a distribution μ over $\Omega = \mathcal{X}^{[N]}$. Starting with arbitrary initial configuration $X \in \Omega$, at each time step, the Glauber dynamics sample a uniformly random location $i \in [N]$ and resamples X_i according to the conditional distribution given other locations, $\mu(x_i | x_{\sim i} = X_{\sim i})$. Let $D_{gl}(X)$ denote the variable generated by applying one step of Glauber dynamics on X .

When applying Glauber dynamics, the crucial property we care about is the mixing time, i.e., the time it takes for the dynamics to converge to the stationary distribution. There are many approaches to prove the mixing speed of Glauber dynamics, in this paper we will be using the path coupling technique introduced in [BD97].

Definition 11 (Weighted Hamming Distance). A path distance over Ω is a weighted hamming distance where $d(X, Y) = \sum_{i: X_i \neq Y_i} c_i$.

Lemma 9.1 (Path Coupling, [BD97]). *Let $d(X, Y)$ denote the path distance between a pair of configurations $X, Y \in \Omega$. If for any pair of X, Y that differs at one coordinate, there is a coupling $(D_{gl}(X), D_{gl}(Y))$ such that*

$$\mathbb{E} d(D_{gl}(X), D_{gl}(Y)) \leq (1 - \alpha) d(X, Y),$$

then the coupling can be extended to any pair of configurations that same inequality holds.

We say the Glauber dynamics satisfies path coupling with constant α with respect to distance d if the above holds. Path coupling imply mixing of Glauber dynamics.

Lemma 9.2. *If the Glauber dynamics of μ satisfies path coupling with constant α with respect to distance d defined in Definition 11, then for any starting configuration X_0 , we have*

$$d_{TV}(D_{gl}^t(X_0), \mu) \leq \delta$$

when $t \geq \log_{1-\alpha}(\epsilon \min_i c_i / (\sum_i c_i))$.

Proof. Now imagine X_0 is the initial configuration chosen by the algorithm and $Y_0 \sim \mu$. Let $X_t = D_{gl}^t(X_0)$ and $Y_t = D_{gl}^t(Y_0)$ denote the random variables we get after applying t steps of Glauber dynamics on X_0 and Y_0 . Using the Lemma above for t times, we get after t time steps,

$$\mathbb{E} d(X_t, Y_t) \leq (1 - \alpha)^t \sum_i c_i.$$

Since the distance $d(X, Y)$ is at least $\min_i c_i$ when $X \neq Y$, we get an upper bound on the total variation distance between the law of X_t and μ :

$$d_{TV}(X_t, \mu) \leq \mathbb{P}(X_t \neq Y_t) \leq \frac{\mathbb{E} d(X_t, Y_t)}{\min_i c_i} \leq (1 - \alpha)^t \frac{\sum_i c_i}{\min_i c_i}.$$

It follows that to generate a sample that is ϵ close to μ in total variation, we can start with arbitrary initial configuration and run Glauber dynamics for $\log_{1-\alpha}(\epsilon \min_i c_i / (\sum_i c_i))$ time-steps. \square

9.2 Concentration from Mixing

In this section we show that if a function can be chosen to represent the distance in path coupling, the function must concentrate.

Lemma 9.3 (Concentration From Path Coupling). *Let μ be a distribution over $\Omega = \{0, 1\}^N$. Let D_{gl} be the Glauber dynamics of μ . Suppose function $f : \{0, 1\}^N \rightarrow \mathbb{R}$ satisfies*

$$\sup_{x_1, \dots, x_N, x'_i} |f(x_1, \dots, x_i, \dots, x_N) - f(x_1, \dots, x'_i, \dots, x_N)| \leq L_i$$

Let D be the path distance on Ω defined by

$$D(x, x') = \sum_{i: x_i \neq x'_i} L_i.$$

Suppose that D satisfies path coupling with constant α , i.e., suppose that for x, x' that differ at precisely one coordinate, there exists a coupling between $D_{gl}(x)$ and $D_{gl}(x')$ such that

$$\mathbb{E} D(D_{gl}(x), D_{gl}(x')) \leq (1 - \alpha)D(x, x').$$

Then for any $\delta > 0$, $t > 0$,

$$\mu(f - \mathbb{E}_\mu(f) > t) \leq \exp\left(\frac{-t^2/2}{4T\|L\|_2^2/N + \|L\|_\infty t/3}\right) + \delta,$$

where $T = \log_{1-\alpha} \frac{\delta \inf_i L_i}{\|L\|_1}$ is the mixing time of D_{gl} . Further, if μ is q -marginally small,

$$\mu(f - \mathbb{E}_\mu(f) > t) \leq \exp\left(\frac{-t^2/2}{8Tq\|L\|_2^2/N + \|L\|_\infty t/3}\right) + \delta.$$

Proof. Consider running Glauber dynamics with initial condition $X_0 = \vec{0}$. Let $X_t = D_{gl}^t(X_0)$, where $t = 0, 1, \dots, T$. Define a sequence of martingale

$$g_t = \mathbb{E}[f(X_T) | X_t].$$

The martingale difference $Y_t = g_t - g_{t-1}$ can be written as

$$Y_t = f(D_{gl}^{T-t}(X_t)) - \mathbb{E}_{X'_t | X_{t-1}} f(D_{gl}^{T-t}(X'_t)).$$

From the definition of D , for any $x, x' \in \Omega$, $D(x, x') \geq |f(x) - f(x')|$, hence

$$\begin{aligned} Y_t &\leq \mathbb{E}_{X'_t | X_{t-1}} D(D_{gl}^{T-t}(X_t), D_{gl}^{T-t}(X'_t)) \\ &\leq \mathbb{E}_{X'_t | X_{t-1}} D(X_t, X'_t) \\ &\leq D(X_t, X_{t-1}) + \mathbb{E}_{X'_t | X_{t-1}} D(X'_t, X_{t-1}). \end{aligned}$$

The second inequality is by the data processing inequality and the third inequality is by the triangle inequality. Since X_t and X_{t-1} differ in at most one coordinate, $Y_t \leq 2\|L\|_\infty$ almost surely. Also, by the upper bound above,

$$\mathbb{E}[Y_t^2|X_{t-1}] \leq \mathbb{E}[D(X_t, X_{t-1})^2] + 3(\mathbb{E}[D(X_t, X_{t-1})|X_{t-1}])^2 \leq 4\mathbb{E}[D(X_t, X_{t-1})^2|X_{t-1}].$$

The Glauber dynamics chooses a location $i \in [N]$ uniformly at random, when i is chosen, $D(X_t, X_{t-1}) \leq L_i$. So

$$\mathbb{E}[Y_t^2|X_{t-1}] \leq 4 \sum_i L_i^2/N = 4\|L\|_2^2/N.$$

When μ is q -marginally small, the probability that we change coordinate i in the Glauber dynamics is at most

$$\frac{\mathbb{P}((X_{t-1})_i = 0)q + \mathbb{P}((X_{t-1})_i = 1)}{N}.$$

Note that $\mathbb{P}((X_{t-1})_i = 1) \leq q$ is maintained when we run Glauber dynamics. So the above quantity is bounded by $2q/N$. We have

$$\mathbb{E}[Y_t^2|X_{t-1}] \leq 4 \sum_i 2qL_i^2/N = 8q\|L\|_2^2/N.$$

Therefore, we can apply McDiarmid's inequality (Lemma 4.9), yielding

$$\mathbb{P}(g_T - \mathbb{E} g_T > t) \leq \exp\left(\frac{-t^2/2}{8qT\|L\|_2^2/N + \|L\|_\infty t/3}\right).$$

By a standard path coupling argument, Lemma 9.2, $d_{\text{TV}}(\mathcal{L}(X_T), \mu) \leq \delta$. Since $g_T = f(X_T)$ by definition, we have

$$\mu(f - \mu(f) > t) \leq \exp\left(\frac{-t^2/2}{8qT\|L\|_2^2/N + \|L\|_\infty t/3}\right) + \delta. \quad \square$$

It turns out that we can apply path coupling to any q -marginally small distribution as long as q is much less than the inverse degree of its graphical model.

Theorem 9.1. *Let P be a distribution over $\{0, 1\}^N$ and let G be the graphical model of P with maximum degree Δ . If P is q -marginally small for $q \ll 1/\Delta$, then:*

1. *If $X \sim P$, then for any f that is L -Lipschitz on $\{0, 1\}^N$, $|f(X) - \mathbb{E} f(X)| = O(\sqrt{NqL^2 \log N})$ with high probability.*
2. *The Glauber dynamics for P mixes in $O(N \log N)$ time.*

Proof. We can use Lemma 9.3 by defining $D(x, x') = L$ for adjacent configurations. In the following, we show that D satisfies the path coupling condition with constant $\alpha = \Omega(1/N)$.

Suppose x and x' differ only at position i , $x_i = 1$, $x'_i = 0$. Consider the following coupling between $\mathcal{D}_{gl}(x)$ and $\mathcal{D}_{gl}(x')$. The Glauber dynamics chooses the same location in x and x' , say $j \in [N]$. If $j = i$ or j is not adjacent to i in G , since $x_j|x_{\sim i}$ and $x'_j|x'_{\sim i}$ have the same distribution, we can couple them by the same choice of x_j . If j is a neighbor of i , we couple them arbitrarily.

When the Glauber dynamics chooses i , the new distance is 0. When Glauber dynamics chooses a neighbor j of i , the new distance is either L or $2L$ depending on whether $\mathcal{D}_{gl}(x)$ and $\mathcal{D}_{gl}(x')$

chooses the same value on j . Since the distribution is q -small, they both choose $x_j = 0$ with at least $1 - 2q$ probability, from the union bound. When Glauber dynamics chooses any other coordinate, the new distance is still L . Therefore, under this coupling,

$$\begin{aligned}\mathbb{E} D(\mathbf{D}_{gl}(x), \mathbf{D}_{gl}(x')) &\leq \frac{N - \Delta - 1}{N} \cdot L + \frac{\Delta}{N} \cdot (1 - 2q) \cdot L + \frac{\Delta}{N} \cdot (2q) \cdot (2L) \\ &\leq \left(1 - \frac{1 - 6q\Delta}{N}\right) L \\ &= \left(1 - \frac{1 - 6q\Delta}{N}\right) D(x, x').\end{aligned}$$

Hence, D satisfies the contraction required by path coupling with constant $(1 - 6q\Delta)/N = \Omega(1/N)$.

Now we can apply Lemma 9.3. For any constant c and $\delta = N^{-c}$, the mixing time T is $\log_{(1-6q\Delta)/N}(N^{-c-1}) = O(N \log N)$. Choosing $t = \sqrt{cqTL^2 \log N} = O(\sqrt{NqL^2 \log^2 N})$, we have

$$\mathbb{P}(f(X) - \mathbb{E} f(X) > t) \leq \exp\left(\frac{cqTL^2 \log N/2}{8qTL^2 + L\sqrt{cqTL^2 \log N}/3}\right) + \delta \leq 2N^{-c/32}.$$

By considering $-f(X)$, the same bound applies to $\mathbb{P}(f(X) - \mathbb{E} f(X) < -t)$. \square

We have the following corollary as an immediate result, noting that μ_G is $O(p')$ -marginally small under arbitrary conditioning.

Corollary 6.5. *Let $f : \{0, 1\}^{\binom{[n]}{3}} \rightarrow \mathbb{R}$ be an L -Lipschitz function. Let T be a subset of $\binom{[n]}{3}$ with size $\binom{n}{3} - m$ and $x_T \in \{0, 1\}^T$ be a configuration on T . If $X \sim \mu_G(\cdot | X_T = x_T)$, then*

$$|f(X) - \mathbb{E} f(X)| = O(\sqrt{mp'L^2 \log m})$$

holds with high probability.

A similar result can be obtained for functions that are not everywhere Lipschitz, but are when restricted to a high probability set.

Corollary 9.4. *Under the assumptions of Theorem 9.1, if there exists a $\mathcal{Y} \subset \{0, 1\}^N$ such that f restricted to \mathcal{Y} is L -Lipschitz, X is in \mathcal{Y} with high probability, and $1/q$ is polynomially bounded by N , we still have $|f(X) - \mathbb{E}[f(X)]| = O(\sqrt{NqL^2 \log N})$ with high probability.*

Proof. By assumption, $\mathbb{P}(X \in \mathcal{Y}) \geq 1 - 1/N^c$ for some sufficiently large constant c . We take

$$\bar{f}(x) = \inf_{y \in \mathcal{Y}} f(y) + d_H(x, y) \cdot L,$$

where $d_H(x, y) = |\{i : x_i \neq y_i\}|$ is the Hamming distance. Note that (i) if $x \in \mathcal{Y}$, then $f(x) = \bar{f}(x)$, and (ii) \bar{f} is everywhere L -Lipschitz. Thus, we can apply Theorem 9.1 to \bar{f} . Since $\mathbb{E} f(X)$ and $\mathbb{E} \bar{f}(X)$ differ by at most $\sup_x |f(x) - \bar{f}(x)|/N^c \leq L/N^{c-1}$, we have that when $X \in \mathcal{Y}$ and $|\bar{f}(X) - \mathbb{E} \bar{f}(X)| = O(\sqrt{NqL^2 \log N})$, which happens with high probability,

$$|f(X) - \mathbb{E} f(X)| = |\bar{f}(X) - \mathbb{E} f(X)| \leq |\bar{f}(X) - \mathbb{E} \bar{f}(X)| + L/N^{c-1} = O(\sqrt{NqL^2 \log N}) + L/N^{c-1}.$$

This last quantity is $O(\sqrt{NqL^2 \log N})$ whenever c is a large enough constant. \square

Lemma 2.3. *For any fixed graph G over n vertices, $p' \ll 1/n$ and constant p , the Glauber dynamics on μ_G mixes in $O(n^3 \log n)$ time.*

Proof. The graphical model of μ_G has $\binom{n}{3}$ variables. Two triangles are connected if and only if they share an edge, so each variable has degree at most $3(n-2)$. Now, μ_G is $O(p')$ -marginally small by Lemma 6.1, so Theorem 9.1 implies that the Glauber dynamics mixes in $O(n^3 \log n)$ time. \square

References

- [ALG21] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *SIAM Journal on Computing*, (0):FOCS20–1, 2021.
- [Bar17] Boaz Barak. *The Complexity of Public-Key Cryptography*, pages 45–77. Springer International Publishing, Cham, 2017.
- [BB19] Matthew Brennan and Guy Bresler. Optimal average-case reductions to sparse pca: From weak assumptions to strong hardness. *arXiv preprint arXiv:1902.07380*, 2019.
- [BB20] Matthew Brennan and Guy Bresler. Reducibility and statistical-computational gaps from secret leakage. In *Conference on Learning Theory*, pages 648–847. PMLR, 2020.
- [BBH18] Matthew Brennan, Guy Bresler, and Wasim Huleihel. Reducibility and computational lower bounds for problems with planted sparse structure. In *COLT*, pages 48–166, 2018.
- [BBH19] Matthew Brennan, Guy Bresler, and Wasim Huleihel. Universality of computational lower bounds for submatrix detection. *arXiv preprint arXiv:1902.06916*, 2019.
- [BBH⁺20] Matthew Brennan, Guy Bresler, Samuel B Hopkins, Jerry Li, and Tselil Schramm. Statistical query algorithms and low-degree tests are almost equivalent. *arXiv preprint arXiv:2009.06107*, 2020.
- [BBN19] Matthew Brennan, Guy Bresler, and Dheeraj Nagaraj. Phase transitions for detecting latent geometry in random graphs. *arXiv preprint arXiv:1910.14167*, 2019.
- [BBN20] Matthew Brennan, Guy Bresler, and Dheeraj Nagaraj. Phase transitions for detecting latent geometry in random graphs. *Probability Theory and Related Fields*, 178(3):1215–1289, 2020.
- [BD97] Russ Bubley and Martin Dyer. Path coupling: A technique for proving rapid mixing in markov chains. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 223–231. IEEE, 1997.
- [BDER16] Sébastien Bubeck, Jian Ding, Ronen Eldan, and Miklós Z Rácz. Testing for high-dimensional geometry in random graphs. *Random Structures and Algorithms*, 2016.
- [BHK⁺16] Boaz Barak, Samuel B Hopkins, Jonathan Kelner, Pravesh Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 428–437. IEEE, 2016.

- [BR13a] Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *COLT*, pages 1046–1066, 2013.
- [BR13b] Quentin Berthet and Philippe Rigollet. Optimal detection of sparse principal components in high dimension. *The Annals of Statistics*, 41(4):1780–1815, 2013.
- [BT06] Andrej Bogdanov and Luca Trevisan. Average-case complexity. *Foundations and Trends® in Theoretical Computer Science*, 2(1):1–106, 2006.
- [CGKM22] Sitan Chen, Aravind Gollakota, Adam Klivans, and Raghu Meka. Hardness of noise-free learning for two-hidden-layer neural networks. *Advances in Neural Information Processing Systems*, 35:10709–10724, 2022.
- [CGŠV21] Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Rapid mixing for colorings via spectral independence. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1548–1557. SIAM, 2021.
- [CLM16] T Tony Cai, Xiaodong Li, and Zongming Ma. Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow. *The Annals of Statistics*, 44(5):2221–2251, 2016.
- [CMW20] Michael Celentano, Andrea Montanari, and Yuchen Wu. The estimation error of general first order methods. In *Conference on Learning Theory*, pages 1078–1141. PMLR, 2020.
- [Com15] Richard Combes. An extension of mcdiarmid’s inequality. *arXiv preprint arXiv:1511.05240*, 2015.
- [DKS17] Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 73–84. IEEE, 2017.
- [Dob70] R. Dobrushin. Prescribing a system of random variables by conditional distributions. *Theory of Probability & Its Applications*, 15(3):458–486, 1970.
- [DV21] Amit Daniely and Gal Vardi. From local pseudorandom generators to hardness of learning. In *Conference on Learning Theory*, pages 1358–1394. PMLR, 2021.
- [FGR⁺13] Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 655–664. ACM, 2013.
- [FPV15] Vitaly Feldman, Will Perkins, and Santosh Vempala. On the complexity of random satisfiability problems with planted solutions. pages 77–86, 2015.
- [GMZ17] Chao Gao, Zongming Ma, and Harrison H Zhou. Sparse cca: Adaptive estimation and computational barriers. *The Annals of Statistics*, 45(5):2074–2101, 2017.

- [Gol11] Oded Goldreich. Notes on Levin’s theory of average-case complexity. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 233–247. Springer, 2011.
- [GS14] David Gamarnik and Madhu Sudan. Limits of local algorithms over sparse random graphs. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 369–376. ACM, 2014.
- [GVV22] Aparna Gupte, Neekon Vafa, and Vinod Vaikuntanathan. Continuous lwe is as hard as lwe & applications to learning gaussian mixtures. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1162–1173, 2022.
- [HKP⁺17] Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detecting hidden structures. *Proceedings of the fifty-eighth IEEE Foundations of Computer Science*, pages 720–731, 2017.
- [Hop18] Samuel B Hopkins. *Statistical Inference and the Sum of Squares Method*. PhD thesis, Cornell University, 2018.
- [HS22] Brice Huang and Mark Sellke. Tight lipschitz hardness for optimizing mean field spin glasses. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 312–322. IEEE, 2022.
- [HWX15] Bruce E Hajek, Yihong Wu, and Jiaming Xu. Computational lower bounds for community detection on random graphs. In *COLT*, pages 899–928, 2015.
- [IS03] Yuri Ingster and IA Suslina. *Nonparametric goodness-of-fit testing under Gaussian models*, volume 169. Springer Science & Business Media, 2003.
- [KLN18] Jeong Han Kim, Sang June Lee, and Joohan Na. On the total variation distance between the binomial random graph and the random intersection graph. *Random Structures & Algorithms*, 52(4):662–679, 2018.
- [KM22] Frederic Koehler and Elchanan Mossel. Reconstruction on trees and low-degree polynomials. *Advances in Neural Information Processing Systems*, 35:18942–18954, 2022.
- [LC12] Lucien Le Cam. *Asymptotic methods in statistical decision theory*. Springer Science & Business Media, 2012.
- [Lev86] Leonid A Levin. Average case complete problems. *SIAM Journal on Computing*, 15(1):285–286, 1986.
- [LZ22] Yuetian Luo and Anru R Zhang. Tensor clustering with planted structures: Statistical optimality and computational limits. *The Annals of Statistics*, 50(1):584–613, 2022.
- [McD98] Colin McDiarmid. Concentration. In *Probabilistic methods for algorithmic discrete mathematics*, pages 195–248. Springer, 1998.
- [MW15] Zongming Ma and Yihong Wu. Computational barriers in minimax submatrix detection. *The Annals of Statistics*, 43(3):1089–1116, 2015.

- [MW22] Andrea Montanari and Alexander S Wein. Equivalence of approximate message passing and low-degree polynomials in rank-one matrix estimation. *arXiv preprint arXiv:2212.06996*, 2022.
- [Ros08] Benjamin Rossman. On the constant-depth complexity of k-clique. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 721–730. ACM, 2008.
- [Ros14] Benjamin Rossman. The monotone complexity of k-clique on random graphs. *SIAM Journal on Computing*, 43(1):256–279, 2014.
- [SW22] Tselil Schramm and Alexander S Wein. Computational barriers to estimation from low-degree polynomials. *The Annals of Statistics*, 50(3):1833–1858, 2022.
- [SZB21] Min Jae Song, Ilias Zadik, and Joan Bruna. On the cryptographic hardness of learning single periodic neurons. *Advances in neural information processing systems*, 34:29602–29615, 2021.
- [WBS16] Tengyao Wang, Quentin Berthet, and Richard J Samworth. Statistical and computational trade-offs in estimation of sparse principal components. *The Annals of Statistics*, 44(5):1896–1930, 2016.
- [WEAM19] Alexander S Wein, Ahmed El Alaoui, and Cristopher Moore. The kikuchi hierarchy and tensor pca. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1446–1468. IEEE, 2019.
- [Wei22] Alexander S Wein. Average-case complexity of tensor decomposition for low-degree polynomials. *arXiv preprint arXiv:2211.05274*, 2022.
- [ZSWB22] Ilias Zadik, Min Jae Song, Alexander S Wein, and Joan Bruna. Lattice-based methods surpass sum-of-squares in clustering. In *Conference on Learning Theory*, pages 1247–1248. PMLR, 2022.
- [ZX17] Anru Zhang and Dong Xia. Tensor svd: Statistical and computational limits. *arXiv preprint arXiv:1703.02724*, 2017.

A Relation Between Models

We will analyze the parameter regimes where each of the three models, Erdős-Rényi, RIG and RGT are close to each other. The RGT can be thought of as partway between RIG and Erdős-Rényi.

A.1 RGT vs Erdős-Rényi

RGT is close to Erdős-Rényi with equal edge density when $p' \ll n^{-3/2}$.

Theorem A.1. *If p is a positive constant, then*

- *When $p' \gg n^{-3/2}$,*

$$d_{TV}(\text{RGT}(n, p, p'), G(n, p + (1 - p)(1 - (1 - p')^{n-2}))) = 1 - o_n(1)$$

and there is an efficient algorithm to distinguish the two distributions with vanishing error.

- When $p' \ll n^{-3/2}$,

$$d_{TV}(\text{RGT}(n, p, p'), G(n, p + (1 - p)(1 - (1 - p')^{n-2}))) = o_n(1).$$

The lower bound on total variation is established through statistical testing using a signed version of triangle count. This test has also been employed to distinguish between Erdős-Rényi and random intersection graphs [BBN20], as well as between Erdős-Rényi and random geometric graphs [BDER16].

The ordinary triangle count with adjacency matrix A is given by

$$\sum_{\{i,j,k\} \subset [n]} A_{ij}A_{jk}A_{ik}.$$

Let $q = p + (1 - p)(1 - (1 - p')^{n-2}) = 1 - (1 - p)(1 - p')^{n-2}$ be the edge density. The *signed triangle count* is given by

$$\tau(A) = \sum_{\{i,j,k\} \subset [n]} (A_{ij} - q)(A_{jk} - q)(A_{ik} - q)$$

It is not hard to show that for Erdős-Rényi,

Lemma A.1. *For A being the adjacency matrix of an Erdős-Rényi random graph $G(n, q)$,*

$$\mathbb{E}[\tau(A)] = 0 \quad \text{and} \quad \text{Var}[\tau(A)] = \binom{n}{3} q^3 (1 - q)^3,$$

The calculation is straightforward and can be found in Section 3 of [BDER16]. For RGT, we will prove

Lemma A.2. *For $\text{RGT}(n, p, p')$,*

$$\mathbb{E}[\tau(A)] = \binom{n}{3} p'(1 - q)^3 + O(1) \quad \text{and} \quad \text{Var}[\tau(A)] = O(n^3).$$

Proof. Let $\tau_{ijk} = (A_{ij} - q)(A_{jk} - q)(A_{ik} - q)$. Let T be the set of triangles added in the construction of $\text{RGT}(n, p, p')$. Each triple is in T independently with probability p' . Let X_{ijk} be the indicator of whether the triple is included in T . Note that Given X_{ijk} , edges (i, j) , (j, k) and (k, i) become independent since all other triangles do not contain two of them. We have

$$\begin{aligned} \mathbb{E} \tau_{ijk} &= \mathbb{P}(X_{ijk} = 1)(1 - q)^3 + \mathbb{P}(X_{ijk} = 0)(\mathbb{E}[A_{ij} - q | X_{ijk} = 0])^3 \\ &= p'(1 - q)^3 + (1 - p')(p + (1 - p)(1 - (1 - p')^{n-3}))^3 \\ &= p'[(1 - q)^3 - p'^2(1 - p)^3(1 - p')^{3n-8}] \end{aligned}$$

Since $p' \leq 1/n$, we have $\mathbb{E} \tau_{ijk} = p'(1 - q)^3 + O(1/n^3)$. So $\mathbb{E} \tau(A) = \binom{n}{3} p'(1 - q)^3 + O(1)$.

Now let's bound the variance. One can attempt to calculate the exact value, but here we only upper bound it by $O(n^3)$ for simplicity. Note that if two triples i, j, k and i', j', k' do not overlap, τ_{ijk} and $\tau_{i',j',k'}$ are independent.

$$\begin{aligned} \text{Var}[\tau(A)] &= \sum_{\{i,j,k\} \subset [n]} \sum_{\{i',j',k'\} \subset [n]} \text{Cov}[\tau_{ijk}, \tau_{i'j'k'}] \\ &= \binom{n}{3} \text{Var}[\tau_{123}] + \binom{n}{4} \binom{4}{2} \text{Cov}[\tau_{123}, \tau_{124}] + \binom{n}{5} \binom{5}{2} \binom{3}{1} \text{Cov}[\tau_{123}, \tau_{145}] \end{aligned}$$

The first term is trivially bounded by $O(n^3)$ since $\text{Var}[\tau_{123}]$ is at most constant. It remains to bound the second and third term.

Let T_0 be a set of triangles of constant size that contains i, j , we have

$$\mathbb{E}[(A_{ij} - q)^2 | X_{T_0} = \vec{0}] = O(1) \text{ and } \mathbb{E}[q - A_{ij} | X_{T_0} = \vec{0}] = O(p') \quad (25)$$

The first equation is trivial as all the variables are bounded by constant, the second equation is equal to $(1-p)(1-p')^{n-2-|T_0|} - (1-p)(1-p')^{n-2} = O(|T_0|p') = O(p')$.

For the second term, $\text{Cov}[\tau_{123}, \tau_{124}]$ Let E_1 be the set of edges included in triangle $(1, 2, 3)$ or triangle $(1, 2, 4)$, let T_1 be the set of all triangles that contains at least 2 edges from E_1 , $|T_1| = 4$. Now $\mathbb{E}[\tau_{123}\tau_{124}] = \mathbb{E}_{X_{T_1}}[\tau_{123}\tau_{124} | X_{T_1}]$. Notice that given X_{T_1} , all the edges becomes independent. So

$$\begin{aligned} \mathbb{E}[\tau_{123}\tau_{124}] &= \mathbb{E}_{X_{T_1}}[\tau_{123}\tau_{124} | X_{T_1}] \\ &= \mathbb{E}_{X_{T_1}} \left[\mathbb{E}[(\mathbb{1}\{(1, 2) \in E(X_{T_1})\})(1-q)^2 + \mathbb{1}\{(1, 2) \notin E(X_{T_1})\})(A_{12} - q)^2) | X_{T_1}] \right. \\ &\quad \left. \prod_{e \in E(X_{T_1}) \setminus \{1, 2\}} (1-q) \prod_{e \in E_1 \setminus E(X_{T_1}) \setminus \{1, 2\}} \mathbb{E}[(A_e - q) | X_{T_1}] \right] \\ &= O(1) \cdot \sum_{X_{T_1}} \mathbb{P}(X_{T_1}) \prod_{e \in E(X_{T_1}) \setminus \{1, 2\}} (1-q) \prod_{e \in E_1 \setminus E(X_{T_1}) \setminus \{1, 2\}} \mathbb{E}[(A_e - q) | X_{T_1}] \end{aligned}$$

The third equality is because the term on edge $(1, 2)$ is always bounded by $O(1)$. When $|X_{T_1}| \geq 2$, $\mathbb{P}(X_{T_1}) \leq p'^2$, the term inside summation is $O(p')$. When $|X_{T_1}| \leq 1$, there are at least two edges in $E_1 \setminus E(X_{T_1}) \setminus \{1, 2\}$, so by (25), the term inside summation is still bounded by $O(p')$. There are constant number of terms, so we have $\mathbb{E}[\tau_{123}\tau_{124}] = O(p'^2)$.

Similar analysis holds for $\mathbb{E}[\tau_{123}\tau_{145}]$. Let E_2 be the set of edges included in one of the triangles and T_2 be the set of triangles that contains at least two edges in E_2 . $|T_2| = 6$.

$$\begin{aligned} \mathbb{E}[\tau_{123}\tau_{145}] &= \mathbb{E}_{X_{T_2}}[\tau_{123}\tau_{145} | X_{T_2}] \\ &= \mathbb{E}_{X_{T_2}} \left[\prod_{e \in E(X_{T_2})} (1-q) \prod_{e \in E_2 \setminus E(X_{T_2})} \mathbb{E}[(A_e - q) | X_{T_2}] \right] \\ &= O(1) \cdot \sum_{X_{T_2}} \mathbb{P}(X_{T_2}) \prod_{e \in E(X_{T_2})} (1-q) \prod_{e \in E_2 \setminus E(X_{T_2})} \mathbb{E}[(A_e - q) | X_{T_2}] \end{aligned}$$

When $|X_{T_2}| \geq 2$, $\mathbb{P}(X_{T_2}) \leq p'^2$, the term inside summation is $O(p')$. When $|X_{T_2}| \leq 1$, there are at least three edges in $E_2 \setminus E(X_{T_2})$, so by (25), the term inside summation is bounded by $O(p'^3)$. There are constant number of terms, so we have $\mathbb{E}[\tau_{123}\tau_{145}] = O(p'^2)$.

Therefore, $\text{Var}[\tau(A)] = O(n^3 + n^4 p'^2 + n^5 p'^2)$. Since $p' = O(1/n)$, we have $\text{Var}[\tau(A)] = O(n^3)$. \square

So we can construct a statistical test between

$$H_0 \sim G(n, p) \quad \text{and} \quad H_1 \sim G(n, p, p')$$

as follows: Calculate $\tau(A)$, if it's larger than $\binom{n}{3} p'(1-q)^3/2$, output alternative hypothesis, otherwise output null hypothesis. By Lemma A.1 and Lemma A.2, when $p' \gg n^{-3/2}$, the test success

with probability $1 - o_n(1)$ by Chebyshev Inequality. So the two problem have total variation distance $1 - o_n(1)$.

Now we prove the upper bound assuming $p' \ll n^{-3/2}$. The main lemma is given in [BBN20] that bounds the deviation we have from Erdős-Rényi when one triangle is added at a uniformly random place.

Lemma A.3 (Lemma 3.1, [BBN20]). *Let $G(n, t, p)$ denote the graph generated by an Erdős-Rényi graph $G(n, p)$ with a clique of size t planted at a uniformly random place. When t and p are constants,*

$$d_{\text{TV}}\left(G(n, t, p), G\left(n, p + (1 - p)\binom{t}{2}\binom{n}{2}^{-1}\right)\right) = O_n(n^{-3/2}).$$

We can use data processing inequality to show that adding i cliques instead of one leads to a TV distance of $O_n(in^{-3/2})$ when compared with Erdős-Rényi with same edge density.

Corollary A.4. *Let $G(n, i, t, p)$ denote the graph generated by an Erdős-Rényi graph $G(n, p)$ with i cliques of size t planted at a random place. When t and p are constants,*

$$d_{\text{TV}}(G(n, i, t, p), G(n, q)) = O_n(in^{-3/2})$$

where $q = p + (1 - p)[1 - (1 - \binom{t}{2}\binom{n}{2}^{-1})^i]$.

Intuitively, when we add $o(n^{3/2})$ number of triangles, the resulting graph is still close to Erdős-Rényi. This corresponds to adding each triangle with probability $o(n^{-3/2})$. Before formally proving this, let's introduce a technical lemma.

Lemma A.5 (Corollary 5.1 in [KLN18]). *For a positive integer N and $0 < p < q < 1$,*

$$d_{\text{TV}}(\text{Binomial}(N, p), \text{Binomial}(N, q)) \leq \gamma + 3\gamma^2.$$

Here $\gamma = (q - p)\sqrt{\frac{N}{p(1-p)}}$.

To simulate the process of adding each triangle independently with probability p' , we can add $M \sim \text{Poisson}(\binom{n}{3} \log(\frac{1}{1-p'}))$ number of triangles independently at uniformly random location. By Poisson splitting, each triangle is added $\text{Poisson}(\log(\frac{1}{1-p'}))$ times, which is equivalent to adding it with probability p' . By Corollary A.4 and convexity of total variation,

$$d_{\text{TV}}\left(\text{RGT}(n, p, p'), \mathbb{E}_M G(n, q(M))\right) = O(n^{-3/2} \mathbb{E}_M M) = O(n^{3/2} p') = o(1),$$

where $q(M) = 1 - (1 - p)\left(1 - \binom{n-2}{3}\binom{n}{3}^{-1}\right)^M$ is the edge density after adding M triangles to $G(n, p)$.

By Lemma A.5 and convexity of total variation,

$$d_{\text{TV}}(G(n, q), \mathbb{E}_M G(n, q(M))) \leq \mathbb{E}_M[\gamma(M)] = O(n \mathbb{E}_M[|q - q_M|]),$$

where

$$\gamma(M) = |q - q(M)| \sqrt{\frac{\binom{n}{2}}{\min\{q(1-q), q(M)(1-q(M))\}}} = O(n|q - q(M)|).$$

Now notice that $\binom{n}{3} \log(\frac{1}{1-p'}) = O(n^3 p')$. By Chebyshev's inequality,

$$\mathbb{E} M - \sqrt{n^3 p' \log n} \leq M \leq \mathbb{E} M + \sqrt{n^3 p' \log n}$$

with probability $O(\log^{-1} n)$. For any M within the above range,

$$\begin{aligned} |q - q(M)| &= \left| (1-p)(1-p')^{n-2} - (1-p) \left(1 - (n-2) \binom{n}{3}^{-1} \right)^M \right| \\ &= (1-p) \left| (n-2)p' - (n-2) \log\left(\frac{1}{1-p'}\right) + O(n^2 p'^2) + O(\sqrt{\log n n^{-1} p'}) \right| = o(n^{-1}). \end{aligned}$$

Here we used that $(1-a)^b = 1 - ab + O(a^2 b^2)$ when $ab \ll 1$. So we have

$$d_{\text{TV}}(G(n, q), \mathbb{E}_M G(n, q(M))) \leq \mathbb{E}_M [\gamma(M)] = o(1) + \log^{-1} n = o(1).$$

Combining this with the fact that $\text{RGT}(n, p, p')$ is close to $\mathbb{E}_M G(n, q(M))$, we have

$$d_{\text{TV}}(G(n, q), \text{RGT}(n, p, p')) = o(1)$$

for $p' \ll n^{3/2}$.

A.2 RGT vs RIG

We now consider random intersection graphs. Using Corollary A.4, we can prove that RIG is close to RGT when $d \gg n^{2.5}$, and is close to Erdős-Rényi when $d \gg n^3$.

Theorem 1.3. *For an RIG with constant edge density, i.e., $d = \Theta(1/\delta^2)$, if $d \gg n^{2.5}$, then*

$$d_{\text{TV}}(\text{RIG}(n, d, \delta), \text{RGT}(n, 1 - e^{-d\delta^2 + (n-2)d\delta^3(1-\delta)^{n-3}}, 1 - e^{-d\delta^3(1-\delta)^{n-3}})) = o(1).$$

Theorem A.2 ([BBN20], Theorem 3.1). *When $d \gg n^3$, we have*

$$d_{\text{TV}}(\text{RIG}(d, \delta), G(n, 1 - (1 - \delta^2)^d)) = o(1).$$

The idea is to provide another viewpoint of generating RIG. We define an $n \times d$ matrix M such that $M_{ij} = 1$ if S_i contains j , and $M_{ij} = 0$ otherwise. An edge i, j is added to the graph if they appear on the same column.

Equivalently, we can add d cliques defined by the columns of M . The size of each clique is determined by a multinomial distribution with size k having probability $\binom{n}{k} \delta^k (1-\delta)^{n-k}$. It can be shown that this distribution is closely approximated by adding a Poisson number of cliques of each size k . Therefore, we can view RIG as adding $\text{Poisson}(\Theta(dn^2 \delta^2))$ edges (2-cliques), $\text{Poisson}(\Theta(dn^3 \delta^3))$ triangles (3-cliques), $\text{Poisson}(\Theta(dn^4 \delta^4))$ 4-cliques, and so on.

Formally, we can define the following equivalent generation of RIG. Let $p_j := \mathbb{P}(|S_i| = j) = \binom{n}{j} \delta^j (1-\delta)^{n-j}$.

1. Sample $(M_0, M_1, \dots, M_n) \sim \text{Multinomial}(d, p_0)$.
2. For all $2 \leq j \leq n$, independently uniformly sample a random subset of size j from $[n]$ M_j times and plant a j -clique at those M_j places.

We can use a series of Poisson distribution to simulate the distribution of M_2, \dots, M_n . Define $\text{RIG}_{\mathbf{P}}(n, d, \delta)$ to be the model generated by the following process.

1. Sample $X \sim \text{Poisson}(d(1 - p_0 - p_1))$.
2. Sample $M_2, \dots, M_n \sim \text{Multinomial}(X, \gamma p_2, \dots, \gamma p_n)$ where $\gamma = (1 - p_0 - p_1)^{-1}$.
3. For all $2 \leq j \leq n$, independently uniformly sample a random subset of size j from $[n]$ M_j times and plant a j -clique at those M_j places.

Note that by Poisson splitting, the process of generating M_j is equivalent to independent sampling $M_j \sim \text{Poisson}(dp_j)$ for $2 \leq j \leq n$. It is proven in [BBN20] that $\text{RIG}_{\mathbf{P}}$ is close to RIG in total variation distance.

Lemma A.6 ([BBN20], Proposition 3.1). *For $d \gg n^2$ and $\delta = \Theta(1/\sqrt{d})$, it holds that*

$$d_{\text{TV}}(\text{RIG}(n, d, \delta), \text{RIG}_{\mathbf{P}}(n, d, \delta)) = o_n(1).$$

To compare it with RIG , we separate the process of adding 3-cliques and define $\text{RIG}_{\mathbf{P}_4}(n, d, \delta)$ by the distribution of $\text{RIG}_{\mathbf{P}}$ conditioned on $M_3 = 0$. Similar to the argument of Theorem 3.1 in [BBN20], we can prove it is close to Erdős-Rényi random graph when $d \gg n^{2.5}$.

Lemma A.7. *For $d \gg n^{2.5}$ and constant edge density, i.e., $\delta = \Theta(1/\sqrt{d})$, it holds that*

$$d_{\text{TV}}(\text{RIG}_{\mathbf{P}_4}(n, d, \delta), G(n, 1 - e^{-d\delta^2 + (n-2)d\delta^3(1-\delta)^{n-3}})) = o_n(1).$$

The proof is deferred to Appendix B as it is similar to the proof of Theorem A.1.

Notice that to transform from $\text{RIG}_{\mathbf{P}_4}$ to $\text{RIG}_{\mathbf{P}}$, we only need to sample $M_3 \sim \text{Poisson}(dp_3)$ and add a uniformly random triangle to $\text{RIG}_{\mathbf{P}_4}$ independently for M_3 times. By Poisson splitting, if we let X_{ijk} be the number of times that triple i, j, k is sampled, we have $X_{ijk} \sim \text{Poisson}(\frac{dp_3}{\binom{n}{3}}) = \text{Poisson}(d\delta^3(1-\delta)^{n-3})$. This means each triple is added to $\text{RIG}_{\mathbf{P}_4}$ independently with probability $1 - e^{-d\delta^3(1-\delta)^{n-3}}$. Combine this observation with Lemma A.7, we immediately have the following corollary.

Corollary A.8. *For $d \gg n^{2.5}$ and constant edge density, i.e., $\delta = \Theta(1/\sqrt{d})$, we have*

$$d_{\text{TV}}(\text{RIG}_{\mathbf{P}}(n, d, \delta), \text{RGT}(n, 1 - e^{-d\delta^2 + (n-2)d\delta^3(1-\delta)^{n-3}}, 1 - e^{-d\delta^3(1-\delta)^{n-3}}))$$

Combining this and Lemma A.6, we have proved Theorem 1.3.

B Deferred Proofs

B.1 Proof of Lemma 7.13 on the distance between $K_{uv}(X)$ and X

Proof. Recall that we define a partition of T_G as follows.

$$\begin{aligned} T_{G0} &\triangleq \{t \in T_G : t \text{ do not contain } u \text{ or } v\} \\ T_{Gv} &\triangleq \{t \in T_G : t \text{ contains } v\} \\ T_{Gu} &\triangleq \{t \in T_G : t \text{ contains } u\} \end{aligned}$$

Recall the following lemma.

Lemma 7.14. *Let G be c -uniformly 2-star dense. With high probability over the randomness of $X_{T_G^0}$, for any $w \notin \{u, v\}$, there are $\Theta(n^2 p')$ different w' such that $(w', v) \in G$ and $(w', w) \in E(X_{T_G^0})$.*

The same holds by exchanging u and v , so we have the the following by combining both conditions.

Lemma B.1. *With high probability over the randomness of $X_{T_{G^0}}$, for any $w \neq u, v$, there are $\Theta(n^2 p')$ different w' that $(w', v) \in G$ and $(w', w) \in E(X_{T_{G^0}})$, and $\Omega(n)$ different w' that $(w', u) \in G$ and $(w', w) \in E(X_{T_{G^0}})$*

We will prove the following claim. For any $X_{T_{G^0}}$ satisfying the condition in Lemma B.1,

$$d_{TV}(K_{uv}(X)|X_{T_{G^0}}, X|X_{T_{G^0}}) = O\left(\frac{\log^3 n}{n^6 p^{11/2}}\right) = \tilde{O}\left(\frac{1}{\sqrt{n}}\right). \quad (26)$$

This would then imply Lemma 7.13 as the condition in Lemma B.1 happens with high probability.

To use Lemma 4.5, we will show that the likelihood ratio between the two distributions concentrates. Note that the distribution of $X|X_{T_{G^0}}$ is just $\mu_G(X|X_{T_{G^0}})$. We can write the density of both distributions as

$$\begin{aligned} & \mathbb{P}(K_{uv}(X) = x|X_{T_{G^0}}) \\ &= \sum_{\substack{t_1=(v,w,w_1), t_2=(u,w,w_2): t_1, t_2 \in x \\ (w,w_1), (v,w_1), (w,w_2), (u,w_2) \in E(x^-) \\ (w,u), (w,v) \notin E(x^-)}} \frac{1}{|V_{uv}(x^-)| \cdot |W_{E(x^-)}(v, w)| \cdot |W_{E(x^-)}(u, w)|} \mu_G(x^-|X_{T_{G^0}}), \end{aligned}$$

where x^- stands for $x - t_1 - t_2$, and

$$\mathbb{P}(X = x|X_{T_{G^0}}) = \mu_G(X = x|X_{T_{G^0}}).$$

So the derivative is

$$\begin{aligned} & \frac{\mathbb{P}(K_{uv}(X) = x|X_{T_{G^0}})}{\mathbb{P}(X = x|X_{T_{G^0}})} \\ &= \sum_{\substack{t_1=(v,w,w_1), t_2=(u,w,w_2): t_1, t_2 \in x \\ (w,w_1), (v,w_1), (w,w_2), (u,w_2) \in E(x^-) \\ (w,u), (w,v) \notin E(x^-)}} \frac{1}{|V_{uv}(x^-)| \cdot |W_{E(x^-)}(v, w)| \cdot |W_{E(x^-)}(u, w)|} \frac{\mu_G(x^-)}{\mu_G(x)} \end{aligned} \quad (27)$$

Similarly to the proof of Lemma 7.12, we will show each term on the right hand side is close to constant with high probability, and the number of terms also concentrates.

Note that if t is chosen in K_{uv} to be added to X , it always increases the number of edges by 2, i.e., $e(x) - e(x - t_1 - t_2) = 2$. So

$$\frac{\mu_G(x - t_1 - t_2)}{\mu_G(x)} = \frac{p^{|x-t_1-t_2|} p^{-e(x-t_1-t_2)}}{p^{|x|} p^{-e(x)}} = \frac{p^2}{p'^2}.$$

From definition, $|V_{uv}(X)|$, $|W_{E(X)}(v, w)|$ and $|W_{E(X)}(u, w)|$ are $O(1)$ -Lipschitz, so we have the following lemma

Lemma B.2. *With high probability,*

$$\left| |W_{E(X)}(u, w)| - \mathbb{E} |W_{E(X)}(u, w)| \right| = O(\sqrt{n^2 p'} \log n),$$

$$\left| |W_{E(X)}(v, w)| - \mathbb{E} |W_{E(X)}(v, w)| \right| = O(\sqrt{n^2 p'} \log n),$$

and

$$\left| |V_{uv}(X)| - \mathbb{E} |V_{uv}(X)| \right| = O(\sqrt{n^2 p'} \log n).$$

Proof. $\mathcal{L}(X|X_{T_{G_0}})$ is a distribution over $O(n^2)$ triangles that is $\Theta(p')$ -marginally small. So we can apply Corollary 6.5. With high probability, any $O(1)$ Lipschitz function is within distance $O(\sqrt{n^2 p'} \log n)$ from its mean with high probability. \square

Lemma B.3. *For any $X_{T_{G_0}}$ satisfying condition in Lemma 7.14,*

$$\mathbb{E}[V_{uv}(X)|X_{T_{G_0}}] = \Theta(n),$$

and for any $w \neq u, v$

$$\mathbb{E}[W_{E(X)}(u, w)|X_{T_{G_0}}] = \Theta(n^3 p'^2) \text{ and } \mathbb{E}[W_{E(X)}(v, w)|X_{T_{G_0}}] = \Theta(n^3 p'^2).$$

Proof. Recall that $V_{uv}(X)$ consists of vertices w that $(w, u), (w, v) \notin E(X)$. Fix any $w \in W_{u,v}(G)$. Both (w, u) and (w, v) are included in at most $n - 2$ triangles in G , each being one on X with probability $O(p')$ as shown in Lemma 6.1. So

$$\mathbb{P}[(w, u), (w, v) \notin E(X)|X_{T_{G_0}}] \geq 1 - 2(n - 2)O(p') = \Omega(1),$$

as $p' \ll 1/n$. Because $|W_{u,v}(G)| \geq cn$, we have

$$\mathbb{E}[W_{E(X)}(u, w)|X_{T_{G_0}}] = \sum_{w \in W_{u,v}(G)} \mathbb{P}[(w, u), (w, v) \notin E(X)|X_{T_{G_0}}] = \Omega(n).$$

Fix any $w \neq u, v$. Consider a vertex w' such that $(w, w') \in E(X_{T_{G_0}}), (v, w') \in G$, since G is c -uniformly 2-star dense, (v, w') is included in $\Theta(n)$ triangles, each having probability $\Theta(p')$ of being in X . So

$$\mathbb{P}[(v, w') \in E(X)|X_{T_{G_0}}] = \Theta(np').$$

From the assumption on $X_{T_{G_0}}$, there are $\Theta(n^2 p')$ such w' , so

$$\mathbb{E}[W_{E(X)}(v, w)|X_{T_{G_0}}] = \Theta(n^2 p') \cdot \Theta(np') = \Theta(n^3 p'^2).$$

The same holds for $\mathbb{E}[W_{E(X)}(u, w)|X_{T_{G_0}}]$ \square

Note that $|W_{E(X)}(v, w)| \geq |W_{E(X_{T_{G_0}})}(v, w)| = \Theta(n)$, so $\mathbb{E}[W_{E(X)}(v, w)|X_{T_{G_0}}] = \Theta(n)$. By Lemma B.2 and Lemma B.3, for any v and w , conditioned on any $X_{T_{G_0}}$,

$$\left| \frac{|V_{uv}(X)| - \mathbb{E} |V_{uv}(X)|}{\mathbb{E} |V_{uv}(X)|} \right| = O(\sqrt{p'} \log n), \quad \left| \frac{|W_{E(X)}(v, w)| - \mathbb{E} |W_{E(X)}(v, w)|}{\mathbb{E} |W_{E(X)}(v, w)|} \right| = O\left(\frac{\log n}{n^2 p'^{3/2}}\right),$$

$$\text{and} \quad \left| \frac{|W_{E(X)}(u, w)| - \mathbb{E} |W_{E(X)}(u, w)|}{\mathbb{E} |W_{E(X)}(u, w)|} \right| = O\left(\frac{\log n}{n^2 p'^{3/2}}\right)$$

holds with high probability.

As $|V_{uv}(X^-)| = |V_{uv}(X - t_1 - t_2)|$ and $|V_{uv}(X)|$ differ by at most 2, so are $|W_{E(X^-)}(v, w)|$ and $|W_{E(X)}(u, w)|$. So with high probability, each term in (27) concentrates,

$$\begin{aligned} & \frac{1}{|V_{uv}(X^-)| \cdot |W_{E(X^-)}(v, w)| \cdot |W_{E(X^-)}(u, w)|} \frac{\mu_G(X^-)}{\mu_G(X)} \\ &= \frac{1}{\mathbb{E}|V_{uv}(X^-)| \cdot |W_{E(X^-)}(v, w)| \cdot \mathbb{E}|W_{E(X^-)}(u, w)|} \frac{p^2}{p'^2} \left(1 + O\left(\frac{\log n}{n^2 p'^{3/2}}\right)\right). \end{aligned} \quad (28)$$

It remains to show that the number of terms also concentrates. Let S be a set of ordered vertices

$$S(X) = \{(w; w_1; w_2) : t_1 = (v, w, w_1) \in X, t_2 = (u, w, w_2) \in X, (w, w_1), (v, w_1), (u, w_2), (w, w_2) \in E(X - t_1 - t_2), (u, w), (v, w) \notin E(X - t_1 - t_2)\}.$$

Lemma B.4. $|S(X)|$ is $O(\log^2 n)$ -Lipschitz with high probability. under $\mathcal{L}(X|X_{T_{g_0}})$,

$$\mathbb{E}[|S(X)||X_{T_{g_0}}] = \Theta(n^7 p'^6) = \tilde{\Theta}(n),$$

and with high probability,

$$||S(X)| - \mathbb{E}[|S(X)||X_{T_{g_0}}]| = O(n\sqrt{p'} \log^3 n) = \tilde{O}(\sqrt{n}).$$

Proof. We first prove that $|S(X)|$ is Lipschitz. Suppose X and X' differ at one triangle, $X' = X \vee e_{t_0}$. Let $T(u, w)$ denote the set of triangles that contains u, w , and similarly $T(v, w)$ for triangles that contain v, w . Let $S_w(X) := \{(w; w_1; w_2) : (w; w_1; w_2) \in S(X)\}$. Since t_1 and t_2 have to be in X , we have $|S_w(X)| \leq |X_{T(u, w)}||X_{T(v, w)}|$. We will show $S(X')$ and $S(X)$ differ by at most $O(\sup_w |X_{T(u, w)}||X_{T(v, w)}|)$ elements, and $|X_{T(u, w)}||X_{T(v, w)}| = \tilde{O}(1)$ with high probability.

First consider $S(X) \setminus S(X')$. Suppose $(w; w_1; w_2) \in S(X)$ but $(w; w_1; w_2) \notin S(X + t_0)$. Then we must have either $(w, v) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)$ or $(w, u) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)$. With loss of generality, assume it is $(w, v) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)$. Since $E(X + t_0 - t_1 - t_2)$ and $E(X - t_1 - t_2)$ differ by at most 3 edges, we have $|\{w : \exists w', (w; w_1; w_2) \in S(X) \setminus S(X')\}| \leq 3$. Notice that each for each w , at most $|S_w(X)|$ of pairs w_1, w_2 satisfy $(w; w_1; w_2) \in S(X)$, otherwise $(w, v) \in E(X - t_1 - t_2)$. So $|S(X) \setminus S(X')| \leq 3 \sup_w |X_{T(u, w)}||X_{T(v, w)}|$.

Then consider $S(X') \setminus S(X)$. If $(w; w_1; w_2) \in S(X') \setminus S(X)$, either $(w_1, v) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)$, $(w, w_1) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)$, $(w_2, u) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)$, or $(w, w_2) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)$. Let

$$S_1 = \{(w; w_1; w_2) : (w; w_1; w_2) \in S(X'), (w_1, v) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)\},$$

$$S_2 = \{(w; w_1; w_2) : (w; w_1; w_2) \in S(X'), (w, w_1) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)\},$$

$$S_3 = \{(w; w_1; w_2) : (w; w_1; w_2) \in S(X'), (w_2, u) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)\},$$

and

$$S_4 = \{(w; w_1; w_2) : (w; w_1; w_2) \in S(X'), (w, w_2) \in E(X + t_0 - t_1 - t_2) - E(X - t_1 - t_2)\}.$$

Then

$$S(X') \setminus S(X) \subset S_1 \cup S_2 \cup S_3 \cup S_4.$$

Still note that for each $|S_w| \leq |X_{T(u,w)}||X_{T(v,w)}|$, and $E(X+t_0-t_1-t_2)-E(X-t_1-t_2)$ has at most 3 edges, we have $|S_2|, |S_4| \leq 3 \sup_w |X_{T(u,w)}||X_{T(v,w)}|$. For S_1 and S_3 , notice that each w_1 has at most 2 different w such that $\exists w_2, (w; w_1; w_2) \in S(X')$ and $(v, w_1) \in E(X+t_0-t_1-t_2)-E(X-t_1-t_2)$. Otherwise by definition of $S(X')$, at least 3 triangles in X' contains edge (v, w_1) , so (v, w_1) would belong to $E(X-t_1-t_2)$ for any t_1, t_2 , contradicting that $(v, w_1) \in E(X+t_0-t_1-t_2)-E(X-t_1-t_2)$. Therefore, $|S_2| \leq 6 \sup_w |X_{T(u,w)}||X_{T(v,w)}|$, similarly, $|S_4| \leq 6 \sup_w |X_{T(u,w)}||X_{T(v,w)}|$. We have $|S(X') \setminus S(X)| \leq 18 \sup_w |X_{T(u,w)}||X_{T(v,w)}|$.

We have proven $|S(X)|$ changes by at most $18 \sup_w |X_{T(u,w)}||X_{T(v,w)}|$ when changing one triangle. Notice that by Lemma 6.1, $|X_{T(u,w)}|$ and $|X_{T(v,w)}|$ are stochastically dominated by $\text{Binomial}(n-2, O(p'))$, so both quantities are $O(\log n)$ with high probability. This means $|S(X)|$ is $O(\log^2 n)$ -Lipschitz with high probability. From Corollary 9.4, there are $O(n^2)$ triangles in $\mathcal{L}(X|X_{T_{G_0}})$, the distribution is $O(p')$ marginally small, so $|S(X)|$ concentrates around its mean with distance $O(n\sqrt{p'} \log^3 n)$ with high probability.

Next we show that $\mathbb{E}|S(X)| = \Theta(n^7 p'^6) = \tilde{\Theta}(n)$. It is sufficient to prove there are $\Theta(n^5 p'^2)$ pairs $(w; w_1; w_2)$ that $(w; w_1; w_2)$ is in $S(X)$ with probability $\tilde{\Theta}(1/n^2)$. Since G is a c -uniformly 2-star dense graph, there are $\Omega(n)$ vertices w such that $(w, u), (w, v) \in G$. From the assumption that for any w , $|W_{E(X_{T_{G_0}})}(v, w)| = \Theta(n^2 p')$ and $|W_{E(X_{T_{G_0}})}(u, w)| = \Theta(n^2 p')$, there are $\Theta(n^5 p'^2)$ pairs $(w; w_1; w_2)$ that $(w, u), (w, v) \in G$ and $(w, w_1), (w, w_2) \in E(X_{T_{G_0}})$.

Next we show that for such $(w; w_1; w_2)$, it is in $S(X)$ with probability $\tilde{\Omega}(1/n^2)$. Consider the event that $(w; w_1; w_2) \in S(X)$, i.e.,

$$\begin{aligned} t_1 &= (v, w, w_1), t_2 = (u, w, w_2), t_1 \in X, t_2 \in X, \\ (u, w_2), (w, w_2), (v, w_1), (w, w_1) &\in E(X-t), (w, u), (w, v) \notin E(X-t). \end{aligned}$$

Since $E(X-t_1-t_2) \supset E(X_{T_{G_0}})$, $(w, w_1), (w, w_2) \in E(X-t)$ always holds. First, $t_1, t_2 \in X$ with probability $\Theta(p'^2)$, next consider the event conditioning on this. (w, v) and (w, u) are included in at most $n-3$ triangles, each being included with probability $\Theta(p')$. So $(w, v), (w, u) \notin E(X-t)$ with probability at least

$$\mathbb{P}((w, v), (w, u) \notin E(X-t) | t_1, t_2 \in X) \geq (1 - \Theta(p'))^{2n-6} = \Theta(1).$$

Graph G is c -uniformly 2-star dense, so there are $\Theta(n)$ triangles contains (u, w_2) and (v, w_1) and do not contain (w, v) or w_u . By Lemma 6.1, conditioning on the all the events above, each triangle is included in X with probability $\Theta(p')$, so $(u, w_2), (v, w_1) \in E(X-t)$ with probability $\Theta(n^2 p'^2)$. Therefore, the event $(w; w_1; w_2) \in S(X)$ happens with probability $\Theta(n^2 p'^4)$. \square

As for any w ,

$$\frac{1}{\mathbb{E}|V_{uv}(X)| \cdot \mathbb{E}|W_{E(X)}(u, w)| \cdot \mathbb{E}|W_{E(X)}(v, w)|} \frac{p^2}{p'^2} = O\left(\frac{1}{n^7 p'^6}\right).$$

Lemma B.4 shows the following function is $O(\log^2 n / (n^7 p'^6))$ -Lipschitz with high probability.

$$f(X) \triangleq \sum_{(w; w_1; w_2) \in S(X)} \frac{1}{\mathbb{E}|V_{uv}(X)| \cdot \mathbb{E}|W_{E(X)}(u, w)| \cdot \mathbb{E}|W_{E(X)}(v, w)|} \frac{p^2}{p'^2}.$$

Again by Corollary 9.4, $f(X)$ concentrates around its mean with distance at most $O(\log^3 n / (n^6 p'^{11/2}))$ with high probability.

Now combine this with (28), we get for any $X_{T_{G_0}}$ satisfying condition in Lemma B.1, the likelihood ratio

$$L(x) := \frac{\mathbb{P}(K_{uv}(X) = x | X_{T_{G_0}})}{\mathbb{P}(X = x | X_{T_{G_0}})}$$

concentrates as

$$|L(X) - \mathbb{E} f(X)| = O\left(\frac{\log^3 n}{n^6 p'^{11/2}}\right) + O\left(|S(X)| \frac{\log n}{n^9 p'^{15/2}}\right),$$

which is $O\left(\frac{\log^3 n}{n^6 p'^{11/2}}\right)$ with high probability since $|S(X)| = \Theta(n^7 p'^6)$ with high probability.

We can use Lemma 4.7 (which states that concentration implies concentration about the mean for bounded random variables), as it is not hard to check that the likelihood ratio is always bounded by $\tilde{O}(n^3)$. Hence, with high probability,

$$|L(X) - \mathbb{E} L(X)| = O\left(\frac{\log^3 n}{n^6 p'^{11/2}}\right).$$

This completes the proof of (26), using that $\mathbb{E} L(X) = 1$ and Lemma 4.5. (26) then implies Lemma 7.13 as stated before. \square

B.2 Proof of Lemma A.7

Recall that $\text{RIG}_{P_4}(n, d, \delta)$ is generated by sampling $M_j \sim \text{Poisson}(dp_j)$ for $j = 2, 4, 5, \dots, n$ where $p_j = \binom{n}{j} \delta^j (1 - \delta)^{n-j}$. And then add a j -clique uniformly for M_j times. Denote (M_4, \dots, M_{12}) by \vec{M} . The event that $M_j > 0$ for $j > 12$ happens with probability $O(dn^{13} \delta^{13}) \ll 1$, so we can only consider the distribution conditioned on this event. By Poisson splitting, the addition of 2-cliques is equivalent to the following. For each $(i, j) \in \binom{[n]}{2}$, add edge (i, j) with probability $1 - e^{-dp_2/\binom{n}{2}} = 1 - e^{-d\delta^2(1-\delta)^{n-2}}$. This results in a Erdős-Rényi graph, $G(n, 1 - e^{-d\delta^2(1-\delta)^{n-2}})$.

By Corollary A.4, and convexity of total variation distance,

$$d_{\text{TV}}\left(\text{RIG}_{P_4}(n, d, \delta), \mathbb{E}_{\vec{M}} G(n, q(\vec{M}))\right) = O(n^{-3/2} \mathbb{E}_{\vec{M}} \|\vec{M}\|_1). \quad (29)$$

where

$$1 - q(\vec{M}) = e^{-d\delta^2(1-\delta)^{n-2}} \prod_{j=4}^{12} \left(1 - \binom{n-2}{j-2} \binom{n}{j}\right)^{M_j} = e^{-d\delta^2(1-\delta)^{n-2}} \prod_{j=4}^{12} \left(1 - \frac{j(j-1)}{n(n-1)}\right)^{M_j}$$

is the edge density of RIG_{P_4} conditioned on \vec{M} . The right hand side of (29) is $\mathbb{E}_{\vec{M}} \|\vec{M}\|_1 \leq \sum_{j \geq 4} dp_j = O(dn^4 \delta^4) \ll 1$ for $d \gg n^{5/2}$ and $\delta = \Theta(\sqrt{1/d})$.

Now it remains to show that $\mathbb{E}_{\vec{M}} G(n, q(\vec{M}))$ is close to Erdős-Rényi. Let $q = 1 - e^{-d\delta^2 + (n-2)d\delta^3(1-\delta)^{n-3}}$, by Lemma A.5,

$$d_{\text{TV}}(G(n, q), \mathbb{E}_{\vec{M}} G(n, q(\vec{M}))) = O(n \mathbb{E}_{\vec{M}} |q - q(\vec{M})|).$$

Since M_j follows Poisson distribution, we have that with probability $O(1/\log n)$, for any $4 \leq j \leq 12$,

$$\mathbb{E} M_j - \sqrt{dp_j \log n} \leq M_j \leq \mathbb{E} M_j + \sqrt{dp_j \log n}.$$

For any \vec{M} within the above range,

$$\begin{aligned} |q - q(\vec{M})| &= \left| e^{-d\delta^2 + (n-2)d\delta^3(1-\delta)^{n-3}} - e^{-d\delta^2(1-\delta)^{n-2}} \prod_{j=4}^{12} \left(1 - \frac{j(j-1)}{n(n-1)}\right)^{M_j} \right| \\ &= e^{-d\delta^2(1-\delta)^{n-2}} \left| \prod_{j=4}^n e^{-\binom{n-2}{j-2}d\delta^j(1-\delta)^{n-j}} \prod_{j=4}^{12} \left(1 - \frac{j(j-1)}{n(n-1)}\right)^{\mathbb{E} M_j} + O(\sqrt{dp_4 \log n/n^2}) \right| \\ &= o(n^{-1}). \end{aligned}$$

Combining this with (29), we have

$$d_{\text{TV}}(G(n, q), \text{RIGP}_4(n, d, \delta)) = o(1).$$

B.3 Proof of Lemma 7.2

Consider $\mathbb{E}_{G \sim \text{RGT}(n, p, p')} \mathbb{P}_{\mathcal{A}}[e \in \mathcal{A}(G + e)]$. Fix any edge e inside S . The probability that reverse transition \mathcal{A} resamples e is bounded by $O(np')$ by Corollary 6.2 when $G + e$ is c -uniformly 2-star dense. And by Lemma 5.2, $G \sim \text{RGT}(n, p, p')$ is $p^2/2$ uniformly 2-star dense with high probability. Overall, e gets resampled with probability $O(np')$, which means $p_e = 1 - O(np')$.

C Technical Lemmas

Lemma 5.4. *Let G be a c -uniformly two-star dense graph with n vertices, u, v are vertices in G and $(u, v) \notin G$. There exists $n_0 > 0$ such that for any $n > n_0$, we can choose for every $w \notin \{u, v\}$ a set $\bar{W}_G(u, w) \subset W_G(u, w)$ such that*

- For every $w \notin \{u, v\}$, $|\bar{W}_G(u, w)| = cn/3$, and
- For every pair $w_1, w_2 \notin \{u, v\}$, either $w_1 \notin \bar{W}_G(u, w_2)$ or $w_2 \notin \bar{W}_G(u, w_1)$.

Proof. We will use a random process to generate $\bar{W}_G(u, w)$ for all w and argue the desired properties is satisfied with non-zero probability. Consider the following process.

1. Start with $\bar{W}_G(u, w) = \emptyset$ for any $w \notin \{u, v\}$.
2. Loop through all pairs of different vertices $w_1, w_2 \notin \{u, v\}$ and do the following.
 - If $w_1 \in W_G(u, w_2)$ but $w_2 \notin W_G(u, w_1)$, put w_1 in $\bar{W}_G(u, w_2)$.
 - If $w_2 \in W_G(u, w_1)$ but $w_1 \notin W_G(u, w_2)$, put w_2 in $\bar{W}_G(u, w_2)$.
 - If $w_1 \in W_G(u, w_2)$ and $w_2 \in W_G(u, w_1)$, with $1/2$ probability put w_1 in $\bar{W}_G(u, w_2)$ and $1/2$ probability put w_2 in $\bar{W}_G(u, w_1)$.
3. If there exists w with $|\bar{W}_G(u, w)| < cn/3$, declare failure of the process. Otherwise, for all $w \in \{u, v\}$, delete an arbitrary set of vertices in $\bar{W}_G(u, w)$ so that it has size $cn/3$ after deletion.

It is clear that the second condition of the lemma is always satisfied through the process, and if the algorithm does not fail in step 3, the first condition is also satisfied. It remains to prove that the probability of failure is less than 1.

Consider a fixed vertex $w \notin \{u, v\}$ and the set $\bar{W}_G(u, w)$ after step 2 of the algorithm. Since G is c -uniformly 2-star dense, we know $W_G(u, w) \geq c(n-2)$. For each $w' \in W_G(u, w)$, it is included in $\bar{W}_G(u, w)$ in step 2 either with probability 1 or with probability $1/2$, independently. Therefore, the size of $\bar{W}_G(u, w)$ stochastically dominates $\text{Binomial}(c(n-2), 1/2)$. And by Chernoff bound, this is smaller than $cn/3$ with probability $e^{-\Omega(n)}$. So by union bounding over all w , the probability of failure of the algorithm is bounded by $ne^{-\Omega(n)}$, which is smaller than 1 for sufficiently large n . \square

Lemma C.1. *Sampling with replacement for $\text{Binomial}(n, p)$ times and sampling without replacement for $\text{Binomial}(n, p)$ times have total variation distance at most $p(1-p) + np^2$.*

Proof. For sampling with replacement k times, the probability of having a collision is

$$1 - \frac{n-1}{n} \cdot \frac{n-2}{n} \cdots \frac{n-k+1}{n} \leq \frac{\sum_{j=1}^k j}{n} \leq \frac{k^2}{n}.$$

Conditioned on not colliding, sampling with replacement has the same distribution as sampling without replacement, so the total variation distance is bounded by $\mathbb{E}[k^2/n]$ where $k \sim \text{Binomial}(n, p)$. The lemma follows from that the second moment of $\text{Binomial}(n, p)$ is $np(1-p) + np^2$. \square

Lemma C.2. *Consider a distribution on $\{0, 1\}^{\mathcal{X}}$. Suppose S and S' can be partitioned as $S = S_1 \cup S_2$ and $S' = S'_1 \cup S'_2$. Then*

- (i) $\mathbf{I}_{S \rightarrow S'} \leq \mathbf{I}_{S_1 \rightarrow S'} + \mathbf{I}_{S_2 \rightarrow S'}$,
- (ii) $\mathbf{I}_{S \rightarrow S'} \leq \mathbf{I}_{S \rightarrow S'_1} + \mathbf{I}_{S \rightarrow S'_2}$,
- (iii) $\mathbf{I}_{S \rightarrow S'} \leq \sum_{x \in S, x \in S'} \mathbf{I}_{x \rightarrow x'}$.

Proof. Suppose

$$\mathbf{I}_{S \rightarrow S'} = d_{\text{TV}}(P_{x_{S'}|x_S^a, \omega}, P_{x_{S'}|x_S^b, \omega}).$$

For the first statement, we can use the triangle inequality to upper bound the last display by

$$d_{\text{TV}}(P_{x_{S'}|x_{S_1}^a, x_{S_2}^a, \omega}, P_{x_{S'}|x_{S_1}^a, x_{S_2}^b, \omega}) + d_{\text{TV}}(P_{x_{S'}|x_{S_1}^a, x_{S_2}^b, \omega}, P_{x_{S'}|x_{S_1}^b, x_{S_2}^b, \omega}) \leq \mathbf{I}_{S_2 \rightarrow S'} + \mathbf{I}_{S_1 \rightarrow S'}.$$

For the second statement, we will design a coupling $(x_{S'}^a, x_{S'}^b)$ between $P_{x_{S'}|x_S^a, \omega}$ and $P_{x_{S'}|x_S^b, \omega}$ so that $\mathbb{P}[x_{S'}^a \neq x_{S'}^b]$ is bounded by $\mathbf{I}_{S \rightarrow S'_1} + \mathbf{I}_{S \rightarrow S'_2}$. First let $(x_{S'_1}^a, x_{S'_1}^b)$ be an optimal coupling between $P_{x_{S'_1}|x_S^a, \omega}$ and $P_{x_{S'_1}|x_S^b, \omega}$, so that

$$\mathbb{P}[x_{S'_1}^a \neq x_{S'_1}^b] \leq \mathbf{I}_{S \rightarrow S_1}.$$

When $x_{S'_1}^a = x_{S'_1}^b = x_{S'_1}$, let the conditional distribution of $(x_{S'_2}^a, x_{S'_2}^b)$ be the optimal coupling between $P_{x_{S'_2}|x_{S'_1}^a, x_S^a, \omega}$ and $P_{x_{S'_2}|x_{S'_1}^b, x_S^b, \omega}$, otherwise, couple $(x_{S'_2}^a, x_{S'_2}^b)$ arbitrarily, so

$$\mathbb{P}[x_{S'_2}^a \neq x_{S'_2}^b | x_{S'_1}^a = x_{S'_1}^b] \leq \mathbf{I}_{S \rightarrow S_2}.$$

Therefore, we have

$$\mathbb{P}[x_{S'}^a \neq x_{S'}^b] = \mathbb{P}[x_{S'_1}^a \neq x_{S'_1}^b] + \mathbb{P}[x_{S'_2}^a \neq x_{S'_2}^b | x_{S'_1}^a = x_{S'_1}^b] \leq \mathbf{I}_{S \rightarrow S_1} + \mathbf{I}_{S \rightarrow S_2}.$$

The proof for $\mathbf{I}_{S \rightarrow S'}^M \leq \mathbf{I}_{S \rightarrow S'_1}^M + \mathbf{I}_{S \rightarrow S'_2}^M$ is identical, except ω is fixed to be empty. Finally, item (iv) follows by repeatedly applying (i) and (ii). \square