

Counting Motifs with Graph Sampling

Jason M. Klusowski

Yihong Wu

*Department of Statistics and Data Science
Yale University*

JASON.KLUSOWSKI@YALE.EDU

YIHONG.WU@YALE.EDU

Editors: Sebastien Bubeck, Vianney Perchet and Philippe Rigollet

Abstract

Applied researchers often construct a network from data that has been collected from a random sample of nodes, with the goal to infer properties of the parent network from the sampled version. Two of the most widely used sampling schemes are *subgraph sampling*, where we sample each vertex independently with probability p and observe the subgraph induced by the sampled vertices, and *neighborhood sampling*, where we additionally observe the edges between the sampled vertices and their neighbors.

In this paper, we study the problem of estimating the number of motifs as induced subgraphs under both models from a statistical perspective. We show that: for parent graph G with maximal degree d , for any connected motif h on k vertices, to estimate the number of copies of h in G , denoted by $s = s(h, G)$, with a multiplicative error of ϵ ,

- For subgraph sampling, the optimal sampling ratio p is $\Theta_k(\max\{(s\epsilon^2)^{-\frac{1}{k}}, \frac{d^{k-1}}{s\epsilon^2}\})$, which only depends on the size of the motif but *not* its actual topology. Furthermore, we show that Horvitz-Thompson type estimators are universally optimal for any connected motifs.
- For neighborhood sampling, we propose a family of estimators, encompassing and outperforming the Horvitz-Thompson estimator and achieving the sampling ratio $O_k(\min\{(\frac{d}{s\epsilon^2})^{\frac{1}{k-1}}, \sqrt{\frac{d^{k-2}}{s\epsilon^2}}\})$, which again only depends on the size of h . This is shown to be optimal for all motifs with at most 4 vertices and cliques of all sizes.

The matching minimax lower bounds are established using certain algebraic properties of subgraph counts. These results allow us to quantify how much more informative neighborhood sampling is than subgraph sampling, as empirically verified by experiments on synthetic and real-world data. We also address the issue of adaptation to the unknown maximum degree, and study specific problems for parent graphs with additional structures, e.g., trees or planar graphs.

Keywords: Graph sampling, network motifs, graph homomorphism numbers, Horvitz-Thompson estimator, minimax lower bounds

1. Introduction

Counting the number of features in a graph is an important statistical and computational problem. These features are typically basic local structures like motifs Milo et al. (2002) or graphlets Pržulj et al. (2004) (e.g., patterns of small subgraphs). Seeking to capture the interactions and relationships between groups and individuals, applied researchers typically construct a network from data that has been collected from a random sample of nodes. This

scenario is sometimes due to resource constraints (e.g., massive social network, surveying a hidden population) or an inability to gain access the full population (e.g., historical data, corrupted data). Most of the problems encountered in practice are motivated by the need to infer global properties of the parent network (population) from the sampled version. For specific motivations and applications of statistical inference on sampled graphs, we refer the reader to the [Cormode and Duffield \(2014\)](#); [Leskovec and Faloutsos \(2006\)](#); [Kolaczyk \(2017\)](#) for comprehensive reviews as well as applications in computer networks and social networks.

From a computational and statistical perspective, it is desirable to design sublinear time (in the size of the graph) algorithms which typically involves random sampling as a primitive to reduce both time and sample complexities. Various sublinear-time algorithms based on edge and degree queries have been proposed to estimate graph properties such as the average degree [Goldreich and Ron \(2008\)](#); [Feige \(2006\)](#), triangles [Eden et al. \(2015\)](#), stars [Aliakbarpour et al.](#), and more general subgraph counts [Gonen et al. \(2011\)](#). In all of these works, however, some form of adaptive queries, e.g. breadth or depth first search, is performed, which can be impractical or unrealistic in the context of certain applications such as social network analysis [Apicella et al. \(2012\)](#) or econometrics [Chandrasekhar and Lewis \(2011\)](#), where a sampled graph is obtained and statistical analysis is to be conducted on the basis of this dataset alone. In this work, we focus on data arising from specific sampling models, in particular, subgraph sampling and neighborhood sampling [Lovász \(2012\)](#), two of the most popular and commonly used sampling models in part due to their simplicity and ease of implementation. In subgraph sampling, we sample each vertex independently with equal probability and observe the subgraph induced by these sampled vertices. In neighborhood sampling, we additionally observe the edges between the sampled vertices and their neighbors. Despite their ubiquity, theoretical understanding of these sampling models in the context of statistical inference and estimation has been lacking.

In this paper, we study the problem of estimating the counts of various classes of motifs, such as edges, triangles, cliques, and wedges, from a statistical perspective. Network motifs are important local properties of a graph. Detecting and counting motifs have diverse applications in a suite of scientific applications including gene regulation networks [Milo et al. \(2002\)](#), protein-protein interaction networks [Uetz et al. \(2000\)](#), and social networks [Wasserman and Faust \(1994\)](#). Throughout this paper, motifs will be viewed as *induced subgraphs* of the parent graph. For a subgraph H , the number of copies of H contained in G as induced subgraphs is denoted by $s(H, G)$. Many useful graph statistics can often be expressed in terms of induced subgraph counts, e.g., the global clustering coefficient, which is the density of induced open triangles. It is worth pointing out that in some literature motifs are also understood as (not necessarily induced) subgraphs [Milo et al. \(2002\)](#). In fact, it is well-known that the number of a given subgraph can be expressed as a linear combination of induced subgraph counts. For instance, if we denote the number of copies of H contained in G as subgraphs by $n(H, G)$, then for wedges, we have $n(\text{wedge}, G) = s(\text{open triangle}, G) + 3s(\text{closed triangle}, G)$.¹ For this reason, we focus on counting motifs as induced subgraphs.

1. More generally, we have (cf. [Lovász \(2012, Eq. \(5.15\)\)](#)):

$$n(H, G) = \sum_{H'} n(H, H')s(H', G), \tag{1}$$

Furthermore, while we make no assumption about the connectivity of the parent graph, we focus on motifs being *connected* subgraphs which is the most relevant case for applications. It is a classical result that subgraph count of disconnected subgraphs can be expressed as a fixed polynomials in terms of the connected ones; cf. [Whitney \(1932\)](#); [Lovász \(2012\)](#). Additionally, motifs in directed graphs have also been considered [Milo et al. \(2002\)](#); in this paper we focus on undirected simple graphs.

The purpose of this paper is to develop a statistical theory for estimating motif counts in sampled graph. We will be concerned with both methodologies as well as their statistical optimality, with focus on large graphs and the sublinear sample regime, where only a vanishing fraction of vertices are sampled. In particular, a few questions we want to address quantitatively are as follows:

- How does the sample complexity depend on the motif itself? For example, is estimating the count of open triangles as easy as estimating the closed triangles? How does the sample complexity of counting 4-cycles compare with that of counting 4-cliques?
- How much of the graph must be observed to ensure accurate estimation? For example, severe under-coverage issues have been observed in the study of protein-protein interaction networks [Han et al. \(2005\)](#).
- How much more informative is neighborhood sampling than subgraph sampling from the perspective of reducing the sample complexity?
- To what extent does additional structures of the parent graph, e.g., tree or planarity, impact the sample complexity?

Finally, let us also mention that motif counts e.g., triangles [Gao and Lafferty \(2017\)](#), wheels [Bickel et al. \(2011\)](#), and cycles [Mossel et al. \(2015\)](#) have been used as useful test statistics for generative network models such as the stochastic block models. Furthermore, edges counts of similarity and dependency graphs have been used in the context of testing and estimating change-point detection [Chen et al. \(2015\)](#); [Chu and Chen \(2017\)](#). In this paper we do not assume any generative network model, and the randomness of the problem comes solely from the sampling mechanism.

1.1. Sampling model

In this subsection, we formally describe the two graph sampling models we will study in the remainder of the paper.

Subgraph sampling. Fix a simple graph $G = (V, E)$ on $v(G)$ vertices. For $S \subset V$, we denote by $G[S]$ the vertex induced subgraph. If S represents a collection of vertices that are randomly sampled according to a sampling mechanism, we denote $G[S]$ by \tilde{G} . The first and simplest sampling model we consider is the subgraph sampling model, where each vertex is sampled with equal probability. In particular, we sample each vertex independently with probability p , where p is called the *sampling ratio* and can be thought of as the fraction of

where the summation ranges over all simple graphs H' (up to isomorphisms) obtained from H by adding edges.

the graph that is observed. Thus, the sample size $|S|$ is distributed as $\text{Bin}(v(G), p)$, and the probability of observing a subgraph isomorphic to H is equal to

$$\mathbb{P}[\tilde{G} \simeq H] = s(H, G)p^{v(H)}(1-p)^{v(G)-v(H)}.$$

There is also a variant of this model where exactly $n = pv(G)$ vertices are chosen uniformly at random without replacement from the vertex set V . In the sublinear sampling regime where $n \ll v(G)$, they are nearly equivalent.

Neighborhood sampling. In this model, in addition to observing $G[S]$, we also observe the labelled neighbors of all vertices in S , denoted by $G\{S\}$. That is, $G\{S\}$ is equal to $\tilde{G} = (V, \tilde{E})$, where $\tilde{E} = \cup_{v \in S} \cup_{u \in N_G(v)} \{u, v\}$ together with the colors $b_v \in \{0, 1\}$ for each $v \in V(\tilde{G})$, indicating which vertices were sampled. We refer to such bicolored graphs as *neighborhood subgraphs*, which is a union of stars with the root vertex of each star colored. This model is also known in the literature as ego-centric [Handcock and Gile \(2010\)](#) or star sampling [Capobianco \(1972\)](#); [Kolaczyk \(2009\)](#).

In other words, we randomly sample the rows of the adjacency matrix of G independently with probability p and then observe the rows together with the row indices. The graph then consists of unions of star graphs (not necessarily disjoint) together with colors indicating the root of the stars. Neighborhood sampling operates like subgraph sampling but neighborhood information is acquired for each sampled vertex. Hence neighborhood sampling is more informative in the sense that, upon sampling the same set of vertices, considerably more edges are observed. For an illustration and comparison of both subgraph and neighborhood sampling, see [Fig. 1](#). Thus it is reasonable to expect (and indeed we will prove in the sequel) that for the same statistical task, neighborhood sampling typically has significantly lower sample complexity than the subgraph sampling scheme. Note that in many cases, neighborhood sampling is more realistic than subgraph sampling (e.g., social network crawling), where sampling a vertex means that its immediate connections (e.g., friends list) are obtained for free.

A more general version of the neighborhood sampling model is described by Lovász in [Lovász \(2012, Section 1.7\)](#), where each sample consists of a radius- r (labeled) neighborhood rooted at a randomly chosen vertex. Since from a union of marked stars one can disassemble each star individually, our model is equivalent to this one with $r = 1$.

It turns out that the knowledge of the colors provides crucial information about the sampled graph and affects the quality of estimation (see [Appendix J](#)). In practice, the model with labels is more realistic since the experimenter would know which nodes were sampled. We henceforth assume that all sampled graphs obtained from neighborhood sampling are bicolored, with black and white vertices corresponding to sampled and non-sampled vertices, respectively. For a neighborhood subgraph h , let $V_b(h)$ (resp. $v_b(h)$) denote the collection (resp. number) of black vertices. Suppose H is a bicolored subgraph of G . Let $\mathbf{N}(H, G)$ be the number of ways that H can appear (isomorphic as a vertex-colored graph) in G from neighborhood sampling with $v_b(H)$ vertices. Thus

$$\mathbb{P}[\tilde{G} \cong H] = \mathbf{N}(H, G)p^{v_b(H)}(1-p)^{v(G)-v_b(H)}.$$

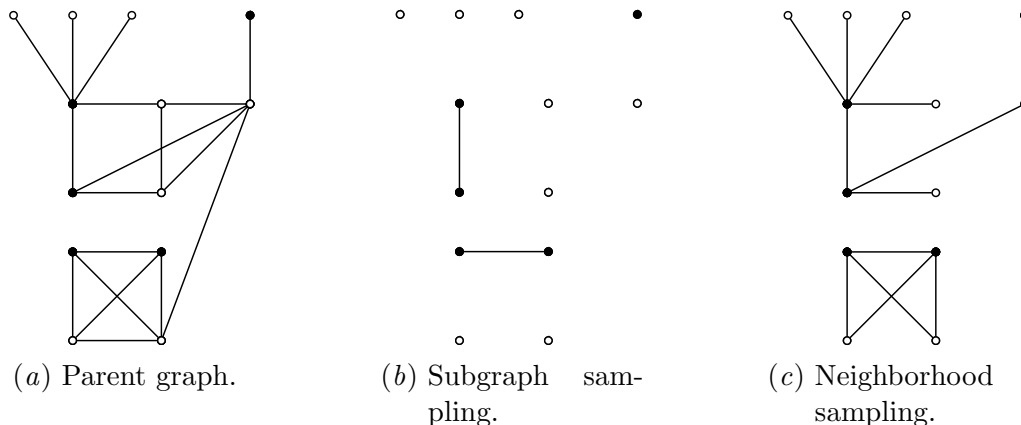


Figure 1: A comparison of subgraph and neighborhood sampling: Five vertices are sampled in the parent graph, and the observed graph is shown in the middle and right panel for the subgraph and neighborhood sampling, respectively.

1.2. Main results

Let h denote a motif, which is a connected graph on k vertices. As mentioned earlier, we do not assume any generative model or additional structures on the parent graph G , except that the maximal degree is at most d ; this parameter, however, need not be bounded, and one of the goals is to understand how the sample complexity depends on d . The goal is to estimate the motif count $s(h, G)$ based on the sampled graph \tilde{G} obtained from either subgraph or neighborhood sampling.

Methodologically speaking, Horvitz-Thompson (HT) estimator [Horvitz and Thompson \(1952\)](#) is perhaps the most natural idea to apply here. The HT estimator is an unbiased estimator of the population total by weighting the empirical count of a given item by the inverse of the probability of observing said item. To be precise, consider estimate the edge count in a graph with m edges and maximal degree d , the sampling ratio required by the HT estimator to achieve a relative error of ϵ scales as $\Theta(\max\{\frac{1}{\sqrt{m\epsilon}}, \frac{d}{m\epsilon^2}\})$, which turns out to be minimax optimal. For ϵ being a small constant, this yields a sublinear sample complexity when m is large and $m \gg d$.

For neighborhood sampling, which is more informative than subgraph sampling since more edges are observed, we show that the optimal sampling ratio can be improved to $\Theta(\min\{\frac{1}{\sqrt{m\epsilon}}, \frac{d}{m\epsilon^2}\})$, which, perhaps surprisingly, is not always achieved by the HT estimator. The main reason for its suboptimality in the high degree regime is the correlation between observed edges. To reduce correlation, we propose a family of linear estimators encompassing and outperforming the Horvitz-Thompson estimator. The key idea is to use the color information indicating which vertices are sampled. For example, in a neighborhood sampled graph it is possible to observe two types of edges: $\bullet \rightarrow \circ$ and $\bullet \rightarrow \bullet$. The estimator takes a linear combination of the count of these two types of edges with a negative weight on the latter, which, as counterintuitive as it sounds, significantly reduces the variance and achieves the optimal sample complexity.

For general motifs h on k vertices, for subgraph sampling, it turns out the simple HT scheme for estimating $s = s(h, G)$ achieves a multiplicative error of ϵ with the optimal

sampling fraction

$$\Theta_k \left(\max \left\{ \frac{1}{(s\epsilon^2)^{\frac{1}{k}}}, \frac{d^{k-1}}{s\epsilon^2} \right\} \right),$$

which only depends on the size of the motif but *not* its actual topology. For neighborhood sampling, the situation is more complicated and the picture is less complete. For general h , we propose a family of estimators that achieves the sample ratio:

$$\Theta_k \left(\min \left\{ \left(\frac{d}{s\epsilon^2} \right)^{\frac{1}{k-1}}, \sqrt{\frac{d^{k-2}}{s\epsilon^2}} \right\} \right)$$

which again only depends on the size of h . We conjecture that this is optimal for neighborhood sampling and we indeed prove this for (a) all motifs up to 4 vertices; (b) cliques of all sizes.

Let us conclude this part by providing some intuition on proving the impossibility results. The main apparatus is matching subgraph counts: If two graphs have matching subgraphs counts for all induced (resp. neighborhood) subgraphs up to size k , then the total variation of the sampled versions obtained from subgraph (resp. neighborhood) sampling are at $O(p^k)$. At a high level, this idea is akin to the method of moment matching, which have been widely used to prove statistical lower bound for functional estimation [Lepski et al. \(1999\)](#); [Cai and Low \(2011\)](#); [Wu and Yang \(2016\)](#); [Jiao et al. \(2015\)](#); in comparison, in the graph-theoretic context, moments correspond to graph homomorphism numbers which are indexed by subgraphs instead of integers [Lovász and Szegedy \(2010\)](#). To give a concrete example, consider the triangle motif and take

$$H = \triangle \text{---} \circ \quad H' = \square \text{---} \circ \quad (2)$$

which have matching subgraph counts up to size two (equal number of vertices and edges) but distinct number of triangles. Then with subgraph sampling, the sampled graph satisfies $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O(p^3)$. For neighborhood sampling, we can take

$$H = \triangle \text{---} \circ \text{---} \circ \quad H' = \square \text{---} \circ \text{---} \circ \quad (3)$$

which have *matching degree sequences* $(3, 2, 2, 2, 1)$ but distinct number of triangles. In general, these pairs of graphs can be either shown to exist by the strong independence of graph homomorphism numbers for connected subgraphs [Erdős et al. \(1979\)](#) or explicitly constructed by a linear algebra argument [Biggs \(1978\)](#); however, for neighborhood sampling it is significantly more involved as we need to relate the neighborhood subgraph counts to the injective graph homomorphism numbers. Based on these small pairs of graphs, the lower bound in general is constructed by using either H or H' as its connected components.

1.3. Notations

We use standard big- O notations, e.g., for any positive sequences $\{a_n\}$ and $\{b_n\}$, $a_n = O(b_n)$ or $a_n \lesssim b_n$ if $a_n \leq Cb_n$ for some absolute constant $C > 0$, $a_n = o(b_n)$ or $a_n \ll b_n$ or if $\lim a_n/b_n = 0$. Furthermore, the subscript in $a_n = O_r(b_n)$ indicates that $a_n \leq C_r b_n$ for some constant C_r depending on r only. For nonnegative integer k , let $[k] = \{1, \dots, k\}$.

Next we establish some graph-theoretic notations that will be used throughout the paper. Let $G = (V, E)$ be a simple, undirected graph. Let $e = e(G) = |E(G)|$ denote the number of edges, $v = v(G) = |V(G)|$ denote the number of vertices, and $cc = cc(G)$ be the number of connected components in G . The open neighborhood of a vertex u is denoted by $N_G(u) = \{v \in V(G) : \{u, v\} \in E(G)\}$. The degree of a vertex u , denoted by d_u , is the cardinality of $N_G(u)$. The closed neighborhood is defined by $N_G[u] = \{u\} \cup N_G(u)$. Two vertices u and v are said to be adjacent to each other, denoted by $u \sim v$, if $\{u, v\} \in E(G)$.

Two graphs G and G' are isomorphic, denoted by $G \simeq G'$, if there exists a bijection between the vertex sets of G and G' that preserves adjacency, i.e., if there exists a bijective function $g : V(G) \rightarrow V(G')$ such that $\{g(u), g(v)\} \in E(G')$ whenever $\{u, v\} \in E(G)$. If G and G' are vertex-colored graphs with colorings c and c' (i.e., a function that assigns a color to each vertex), then G and G' are isomorphic as vertex-colored graphs, denoted by $G \cong G'$, if there exist a bijection $g : V(G) \rightarrow V(G')$ such that $\{g(u), g(v)\} \in E(G')$ whenever $\{u, v\} \in E(G)$ and $c(v) = c'(g(v))$ for all vertices $v \in V(G)$.

We say that H is an (edge-induced) subgraph of G , denoted by $H \subset G$, if $V(H) \subset V(G)$ and $E(H) \subset E(G)$. For any $S \subset V(G)$, the subgraph of G induced by S is denoted by $G[S] \triangleq (S, E(G) \cap S \times S)$. Let $s(H, G)$ (resp. $n(H, G)$) be the number of vertex (resp. edge) induced subgraphs of G that are isomorphic to H ; in other words,

$$s(H, G) = \sum_{V \subset V(G)} \mathbb{1} \{G[V] \simeq H\} \tag{4}$$

$$n(H, G) = \sum_{g \subset G} \mathbb{1} \{g \simeq H\}. \tag{5}$$

For example, $s(\circ-\circ-\circ, \triangle) = 2$ and $n(\circ-\circ-\circ, \triangle) = 8$. Let $\omega(G)$ denote the clique number, i.e., the size of the largest clique in G . Let $e(G) = s(\circ-\circ, G)$, $t(G) = s(\triangle, G)$ and $w(G) = s(\text{wedge}, G)$ denote the number of edges, triangles and wedges of G , which are of particular interest.

1.4. Organization

The paper is organized as follows. In Section 2, we state our positive results in terms of minimax rates for squared error and design algorithms that achieve them for subgraph (Section 2.1) and neighborhood (Section 2.2) sampling. Appendix D discusses converse results and states counterpart minimax lower bounds for subgraph (Appendix D.2) and neighborhood (Appendix D.3) sampling. We further restrict the class of graphs to be acyclic or planar in Appendix E and explore whether such additional structure can be exploited to improve the quality of estimation. In Section 3, we perform a numerical study of the proposed estimators for counting edges, triangles, and wedges on both simulated and real-world data. We discuss a few open problems in Appendix F. Finally, in Appendix H, we prove some of the auxiliary lemmas and theorems that were stated in the main body of the paper.

2. Methodologies and performance guarantees

2.1. Subgraph sampling

The motivation for our estimation scheme is based on the observation that any motif count $s(h, G)$ can be written as a sum of indicator functions as in (4). Note that for a fixed subset of vertices $T \subset V(G)$, the probability it induces a subgraph in the sampled graph \tilde{G} that is isomorphic to h is

$$\mathbb{P}[\tilde{G}[T] \simeq h] = p^{v(h)} \mathbb{1}\{G[T] \simeq h\}.$$

In view of (4), this suggests the following unbiased estimator of $s(h, G)$:

$$\hat{s}_h \triangleq s(h, \tilde{G})/p^{v(h)}. \tag{6}$$

We refer to this estimator as the Horvitz-Thompson (HT) estimator [Horvitz and Thompson \(1952\)](#) since it also uses inverse probability weighting to achieve unbiasedness. The next theorem gives an upper bound on the mean-squared error for this simple scheme, which, somewhat surprisingly, turns out to be minimax optimal within a constant factor as long as the motif h is *connected*.

Theorem 1 (Subgraph sampling) *Let h be an arbitrary connected graph with k vertices. Let G be a graph with maximum degree at most d . Consider the subgraph sampling model with sampling ratio p . Then the estimator (6) satisfies*

$$\mathbb{E}_G |\hat{s}_h - s(h, G)|^2 \leq s(h, G) \cdot k2^k \left(\frac{1}{p^k} \vee \frac{d^{k-1}}{p} \right).$$

Furthermore,

$$\inf_{\tilde{s}} \sup_{\substack{G: d(G) \leq d \\ s(h, G) \leq s}} \mathbb{E}_G |\tilde{s} - s(h, G)|^2 = \Theta_k \left(\left(\frac{s}{p^k} \vee \frac{sd^{k-1}}{p} \right) \wedge s^2 \right).$$

The above result establishes the optimality of the HT estimator for classes of graphs with degree constraints. Since the lower bound construction actually uses instances of graphs containing many cycles, it is a priori unclear whether additional assumptions such as tree structures can help. Indeed, for the related problem of estimating the number of connected components with subgraph sampling, it has been shown that for parent graphs that are forests the sample complexity is strictly smaller [Klusowski and Wu \(2018\)](#). Nevertheless, for counting motifs such as edges or wedges, in [Theorem 13](#) and [Theorem 15](#) we show that the HT estimator (6) cannot be improved up to constant factors even if the parent graph is known to be a forest. The proof of the lower bound of [Theorem 1](#) is given in [Appendix D.2](#). The proof of [Theorem 1](#) is given in [Appendix I](#). To lighten the notation, throughout the rest of the paper, we set $q \triangleq 1 - p$.

2.2. Neighborhood sampling

Our methodology is again motivated by (4) which represents neighborhood subgraph counts as a sum of indicators. In contrast to subgraph sampling, a motif can be observed in the

sampled graph by sampling only some, but not all, of its vertices. For example, we only need to sample one vertex of an edge, or two vertices of a triangle to observe the full motif. More generally, for a subset T vertices in G , we can determine whether $H \simeq G[T]$ or not if at least $v(H) - 1$ vertices from T are sampled. This reduces the variance but introduces more correlation at the same time.

Throughout this subsection, the neighborhood sampled graph is again denoted by $\tilde{G} = G\{S\}$, and $b_v = \mathbb{1}\{v \in S\}$ indicates whether a given vertex v is sampled.

2.2.1. EDGES

We begin by discussing the Horvitz-Thompson type of estimator and why it falls short for the neighborhood sampling model. Extensions to general motifs are presented in Appendix B. Analogously to the estimator (6) designed for subgraph sampling, for neighborhood sampling, we can take the observed number of edges and re-weight it according to the probability of observing an edge. Note that with neighborhood sampling, a given edge is observed if and only if at least one of the end points is sampled. Thus, the corresponding Horvitz-Thompson type edge estimator is

$$\hat{e}_{\text{HT}} = \frac{e(\tilde{G})}{p^2 + 2pq}, \quad (7)$$

which is again an unbiased estimator for $e(G)$. To bound the variance, put $\tau = p^2 + 2pq \in [p, 2p]$ and write

$$e(\tilde{G}) = \sum_{A \in E(G)} r_A.$$

where $A = \{u, v\}$ and $r_A \triangleq \mathbb{1}\{b_u = 1 \text{ or } b_v = 1\} \sim \text{Bern}(\tau)$. For another edge $A' = \{v, w\}$ intersecting A , we have $\text{Cov}[r_A, r_{A'}] = \mathbb{P}[b_v = 1 \text{ or } b_u = b_w = 1] - \tau^2 \leq 3p$, by the union bound. Thus the number of non-zero covariance terms is determined by $n(\mathcal{J}_{\circ}, G)$, the number of \mathcal{J}_{\circ} contained in G as subgraphs, and we have

$$\text{Var}[e(\tilde{G})] \leq e(G)\tau + 2n(\mathcal{J}_{\circ}, G)(3p) \leq 2e(G)p(1 + 3d), \quad (8)$$

where we used the fact that $n(\mathcal{J}_{\circ}, G) \leq e(G)d$. Therefore, the variance of the Horvitz-Thompson estimator satisfies

$$\text{Var}[\hat{e}_{\text{HT}}] \lesssim \frac{e(G)d}{p}. \quad (9)$$

However, as we show next, this estimator is suboptimal when $p > \frac{1}{d}$, or equivalently, when the maximum degree exceeds $\frac{1}{p}$. In fact, the bound (9) itself is tight which can be seen by considering a star graph G with d leaves, and the suboptimality of the HT estimator is largely due to the *heavy correlation* between the observed edges. For example, for the star graph, the correlation is introduced through the root vertex, since with probability p we observe a full star, and with probability q a star with $\text{Bin}(d, p)$ number of black leaves. Thus, the key observation is to incorporate the colors of the vertices to reduce (or eliminate) correlation.

Next, we describe a class of estimators, encompassing and improving the Horvitz-Thompson estimator. Consider

$$\widehat{e} = \sum_{A \in E(\widetilde{G})} \mathcal{K}_A, \quad (10)$$

where \mathcal{K}_A has the form

$$\mathcal{K}_A = b_u(1 - b_v)f(d_u) + b_v(1 - b_u)f(d_v) + b_u b_v g(d_u, d_v); \quad (11)$$

here $A = \{u, v\}$ and f and g are functions of the degree of sampled vertices. For the neighborhood sampling model, this estimator is well-defined since the degree of any sampled vertex is observed without error. It is easy to see that





$$\mathbb{E}[\widehat{e}] = \sum_{\{u,v\} \in E(G)} [pq(f(d_u) + f(d_v)) + p^2g(d_u, d_v)]. \quad (12)$$

For simplicity, next we choose f and g to be constant; in other words, we do not use the degree information of the sampled vertices. This strategy works as long as the maximal degree d of the parent graph is known. To illustrate the main idea, we postpone the discussion on adapting to the unknown d to Appendix C. With $f \equiv \alpha$ and $g \equiv \beta$, the estimator (10) reduces to

$$\widehat{e} = \alpha N(\bullet \dashrightarrow, \widetilde{G}) + \beta N(\bullet \dashleftarrow, \widetilde{G}), \quad (13)$$

which is a linear combination of the counts of the two types of observed edges. In contrast to the HT estimator (7) which treats the two types of edges equally, the optimal choice will weigh them differently. Furthermore, somewhat counter-intuitively, the weights can be negative, which serves to reduce the correlation.

Table 1: Probability mass function of $\mathcal{K}_A \mathcal{K}_{A'}$ for two distinct intersecting edges (excluding zero values).

Graph				
Probability	pq^2	$2p^2q$	p^2q	p^3
Value	α^2	$\alpha\beta$	α^2	β^2

In view of (12), one way of making \widehat{e} unbiased is to set

$$pq(f(d_u) + f(d_v)) + p^2g(d_u, d_v) = 2pq\alpha + p^2\beta = 1. \quad (14)$$

Since the unbiased estimator is not unique, we set out to find the one with the minimum variance. Similar to (8), we have

$$\text{Var}[\widehat{e}] = e(G)\text{Var}[\mathcal{K}_A] + 2n(\text{graph with two intersecting edges}, G)\text{Cov}[\mathcal{K}_A, \mathcal{K}_{A'}] \leq e(G)(\text{Var}[\mathcal{K}_A] + 2d\text{Cov}[\mathcal{K}_A, \mathcal{K}_{A'}]), \quad (15)$$

where $A = \{u, v\}$ and $A' = \{v, w\}$ are distinct intersecting edges. Using Table 1, we find

$$\begin{aligned} \text{Var}[\mathcal{K}_A] &= 2pq\alpha^2 + p^2\beta^2 - 1 \\ \text{Cov}[\mathcal{K}_A, \mathcal{K}_{A'}] &= \alpha^2(pq^2 + p^2q) + p^3\beta^2 + 2p^2q\alpha\beta - 1. \end{aligned}$$

In fact, when the unbiased condition (14) is met, the covariance simplifies to $\text{Cov}[\mathcal{K}_A, \mathcal{K}_{A'}] = \frac{q}{p}(1 - p\alpha)^2 \geq 0$. Finally, optimizing the RHS of (15) over α, β subject to the constraint (14), we arrive the following performance guarantee for \hat{e} :

Theorem 2 *Set*

$$\alpha = \frac{1 + dp}{p(2 + (d - 1)p)}, \quad \beta = \frac{1 - d(1 - 2p)}{p(2 + (d - 1)p)}. \quad (16)$$

Then

$$\mathbb{E}[(\hat{e} - e(G))^2] = \text{Var}[\hat{e}] \leq \frac{e(G)(d + 1)q^2}{p(2 + (d - 1)p)} \lesssim e(G) \left(\frac{1}{p^2} \wedge \frac{d}{p} \right). \quad (17)$$

Furthermore, if p is bounded from one, then

$$\inf_{\hat{e}} \sup_{\substack{G: d(G) \leq d \\ e(h, G) \leq m}} \mathbb{E}_G |\hat{e} - e(G)|^2 = \Theta \left(\left(\frac{md}{p} \wedge \frac{m}{p^2} \right) \wedge m^2 \right)$$

The optimal weights in (16) appear somewhat mysterious. In fact, the following more transparent choice also achieves the optimal risk within constant factors:

- $p \leq 1/d$: we can set either $\alpha = \beta = \frac{1}{p^2 + 2pq}$ or $\alpha = \frac{1}{2pq}$ and $\beta = 0$, that is, we can use either the full HT estimator (7), or the HT estimator restricted to only edges of type $\bullet \circ$, which is the more probable one.
- $p > 1/d$: we choose $\alpha = \frac{1}{p}$ and $\beta = \frac{1 - 2q}{p^2}$. This is the unique weights that simultaneously kill all covariance terms and, at the same time, achieve zero bias. Note that although zero covariance is always possible, it is at a price of setting $\beta \approx -\frac{1}{p^2}$, which inflates the variance too much when p is small and hence suboptimal when $p \ll \frac{1}{d}$.

3. Numerical experiments


We perform our experiments on both synthetic and real-world data. For the synthetic data, we take as our parent graph G a realization of an Erdős-Rényi graph $\mathcal{G}(N, \delta)$ for various choices of parameters. For the real-world experiment, we study the collaboration network between jazz musicians in 198 bands that performed between 1912 and 1940 Gleiser and Danon (2003). Each node is a jazz musician and there is an edge between two musicians if and only if they played together in a band.

The error bars in Fig. 2 shows the variability of the relative error of edges over 10 independent experiments of subgraph and neighborhood sampling on a fixed parent graph G . Similar plots are provided in Fig. 3 and Fig. 4 in Appendix G for triangles and wedges, respectively. The solid black horizontal line shows the sample average and the whiskers show the mean \pm the standard deviation.


Specifically, for subgraph sampling, we always use the HT estimator (6). For neighborhood sampling, for counting triangles or wedges, we use the estimator (20) with choice of parameters given in Theorem 3 and for counting edges we use the adaptive estimator in Theorem 4 from Appendix C.

As predicted by the variance bounds, the estimators based on neighborhood sampling perform better than subgraph sampling. Furthermore, there is markedly less variability


across the 10 independent experiments in neighborhood sampling. However, this variability decreases as p grows. Furthermore, in accordance with our theory, counting bigger motifs (involving more vertices) is subject to more variability, which is evidenced in the plots for triangles and wedges by the wider spread in the whiskers.




(a) Jazz network (subgraph sampling).



(b) Jazz network (neighborhood sampling).



(c) Erdős-Rényi graph (subgraph sampling).



(d) Erdős-Rényi graph (neighborhood sampling).

Figure 2: Relative error of estimating the edge count. In Fig. 2a and Fig. 2b, the parent graph G is the jazz network with $d = 100$, $v(G) = 198$, $e(G) = 2742$. In Fig. 2c and Fig. 2d, G is a realization of the Erdős-Rényi graph $\mathcal{G}(1000, 0.005)$ with $d = 12$, and $e(G) = 2536$.

References

- Maryam Aliakbarpour, Amartya Shankha Biswas, Themis Gouleakis, John Peebles, Ronitt Rubinfeld, and Anak Yodpinyanee. Sublinear-time algorithms for counting star subgraphs via edge sampling. *Algorithmica*, pages 1–30.
- Coren L. Apicella, Frank W. Marlowe, James H. Fowler, and Nicholas A. Christakis. Social networks and cooperation in hunter-gatherers. *Nature*, 481(7382):497–501, 01 2012. URL <http://dx.doi.org/10.1038/nature10736>.
- Peter J. Bickel, Aiyou Chen, and Elizaveta Levina. The method of moments and degree distributions for network models. *Ann. Statist.*, 39(5):2280–2301, 2011. ISSN 0090-5364. doi: 10.1214/11-AOS904. URL <http://dx.doi.org/10.1214/11-AOS904>.
- Norman Biggs. On cluster expansions in graph theory and physics. *Quart. J. Math. Oxford Ser. (2)*, 29(114):159–173, 1978. ISSN 0033-5606. doi: 10.1093/qmath/29.2.159. URL <http://dx.doi.org/10.1093/qmath/29.2.159>.
- T.T. Cai and M. G. Low. Testing composite hypotheses, Hermite polynomials and optimal estimation of a nonsmooth functional. 39(2):1012–1041, 2011.
- Michael Capobianco. Estimating the connectivity of a graph. In Y. Alavi, D. R. Lick, and A. T. White, editors, *Graph Theory and Applications*, pages 65–74, Berlin, Heidelberg, 1972. Springer Berlin Heidelberg.
- Arun Chandrasekhar and Randall Lewis. Econometrics of sampled networks. 2011.
- Hao Chen, Nancy Zhang, et al. Graph-based change-point detection. *The Annals of Statistics*, 43(1):139–176, 2015.
- Lynna Chu and Hao Chen. Asymptotic distribution-free change-point detection for modern data. *arXiv preprint arXiv:1707.00167*, 2017.
- Graham Cormode and Nick Duffield. Sampling for big data: a tutorial. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1975–1975. ACM, 2014.
- Talya Eden, Amit Levi, Dana Ron, and C. Seshadhri. Approximately counting triangles in sublinear time. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science—FOCS 2015*, pages 614–633. IEEE Computer Soc., Los Alamitos, CA, 2015.
- Paul Erdős, László Lovász, and Joel Spencer. Strong independence of graphcopy functions. In *Graph theory and related topics (Proc. Conf., Univ. Waterloo, Waterloo, Ont., 1977)*, pages 165–172. Academic Press, New York-London, 1979.
- Uriel Feige. On sums of independent random variables with unbounded variance and estimating the average degree in a graph. *SIAM J. Comput.*, 35(4):964–984, 2006. ISSN 0097-5397. doi: 10.1137/S0097539704447304. URL <http://dx.doi.org/10.1137/S0097539704447304>.

- Chao Gao and John Lafferty. Testing network structure using relations between small subgraph probabilities. *arXiv preprint arXiv:1704.06742*, 2017.
- Pablo M Gleiser and Leon Danon. Community structure in jazz. *Advances in Complex Systems*, 6(04):565–573, 2003.
- Oded Goldreich and Dana Ron. Approximating average parameters of graphs. *Random Structures Algorithms*, 32(4):473–493, 2008. ISSN 1042-9832. doi: 10.1002/rsa.20203. URL <http://dx.doi.org/10.1002/rsa.20203>.
- Mira Gonen, Dana Ron, and Yuval Shavitt. Counting stars and other small subgraphs in sublinear-time. *SIAM J. Discrete Math.*, 25(3):1365–1411, 2011. ISSN 0895-4801. doi: 10.1137/100783066. URL <http://dx.doi.org/10.1137/100783066>.
- Jing-Dong J Han, Denis Dupuy, Nicolas Bertin, Michael E Cusick, and Marc Vidal. Effect of sampling on topology predictions of protein-protein interaction networks. *Nature Biotechnology*, 23(7):839, 2005.
- Mark S. Handcock and Krista J. Gile. Modeling social networks from sampled data. *Ann. Appl. Stat.*, 4(1):5–25, 2010. ISSN 1932-6157. URL <https://doi.org/10.1214/08-AOAS221>.
- Denis Hanson. On the product of the primes. *Canad. Math. Bull.*, 15:33–37, 1972. ISSN 0008-4395. doi: 10.4153/CMB-1972-007-7. URL <http://dx.doi.org/10.4153/CMB-1972-007-7>.
- Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952.
- Jiantao Jiao, Kartik Venkat, Yanjun Han, and Tsachy Weissman. Minimax estimation of functionals of discrete distributions. *IEEE Transactions on Information Theory*, 61(5):2835–2885, 2015.
- Jason M. Klusowski and Yihong Wu. Estimating the number of connected components in a graph via subgraph sampling. *arXiv preprint arXiv:1801.04339*, 2018.
- W. L. Kocay. Some new methods in reconstruction theory. In *Combinatorial mathematics, IX (Brisbane, 1981)*, volume 952 of *Lecture Notes in Math.*, pages 89–114. Springer, Berlin-New York, 1982.
- Eric D Kolaczyk. *Statistical Analysis of Network Data: Methods and Models*. Springer Science & Business Media, 2009.
- Eric D. Kolaczyk. *Topics at the Frontier of Statistics and Network Analysis: (Re)Visiting the Foundations*. SemStat Elements. Cambridge University Press, 2017. doi: 10.1017/9781108290159.
- Oleg Lepski, Arkady Nemirovski, and Vladimir Spokoiny. On estimation of the L_r norm of a regression function. 113(2):221–253, 1999.

- Jure Leskovec and Christos Faloutsos. Sampling from large graphs. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 631–636. ACM, 2006.
- László Lovász. *Large Networks and Graph Limits*, volume 60. American Mathematical Society, 2012.
- László Lovász and Balázs Szegedy. The graph theoretic moment problem. *arXiv preprint arXiv:1010.5159*, 2010.
- Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594): 824–827, 2002.
- Elchanan Mossel, Joe Neeman, and Allan Sly. Reconstruction and estimation in the planted partition model. *Probab. Theory Related Fields*, 162(3-4):431–461, 2015. ISSN 0178-8051. URL <https://doi.org/10.1007/s00440-014-0576-6>.
- M. Nair. On Chebyshev-type inequalities for primes. *Amer. Math. Monthly*, 89(2):126–129, 1982. ISSN 0002-9890. doi: 10.2307/2320934. URL <http://dx.doi.org/10.2307/2320934>.
- Natasa Pržulj, Derek G Corneil, and Igor Jurisica. Modeling interactome: scale-free or geometric? *Bioinformatics*, 20(18):3508–3515, 2004.
- Peter Uetz, Loic Giot, Gerard Cagney, Traci A Mansfield, Richard S Judson, James R Knight, Daniel Lockshon, Vaibhav Narayan, Maithreyan Srinivasan, Pascale Pochart, et al. A comprehensive analysis of protein–protein interactions in *saccharomyces cerevisiae*. *Nature*, 403(6770):623, 2000.
- Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*, volume 8. Cambridge university press, 1994.
- Douglas B. West. *Introduction to graph theory*. Prentice Hall, Inc., Upper Saddle River, NJ, 1996. ISBN 0-13-227828-6.
- Hassler Whitney. The coloring of graphs. *Ann. of Math. (2)*, 33(4):688–718, 1932. ISSN 0003-486X. doi: 10.2307/1968214. URL <http://dx.doi.org/10.2307/1968214>.
- Yihong Wu and Pengkun Yang. Minimax rates of entropy estimation on large alphabets via best polynomial approximation. *IEEE Transactions on Information Theory*, 62(6): 3702–3720, 2016.

Appendix A. Appendix notations

Let K_n , P_n , and C_n denote the complete graph (clique), path graph, and cycle graph on n vertices, respectively. Let $K_{n,n'}$ denote the complete bipartite graph (biclique) with nn' edges and $n + n'$ vertices. Let S_n denote the star graph $K_{1,n}$ on $n + 1$ vertices.

Define the following graph operations cf. West (1996): The *disjoint union* of graphs G and G' , denoted $G + G'$, is the graph whose vertex (resp. edge) set is the disjoint union of the vertex (resp. edge) sets of G and of G' . For brevity, we denote by kG to the disjoint union of k copies of G . The *join* of G and G' , denoted by $G \vee G'$, is obtained from the disjoint union $G + G'$ by connecting all $v \in V(G)$ and all $v' \in V(G')$, that is, $G \vee G' = (V(H) \cup V(H'), E(H) \cup E(H') \cup (V(H) \times V(H')))$, where $H \simeq G$ and $H' \simeq G'$ and $V(H)$ and $V(H')$ are disjoint. For example, $nK_1 \vee n'K_1 = K_{n,n'}$. For $S \subset V(G)$, let $G - S$ denote the resulting graph after deleting all vertices in S and all incident edges, and $G - v \triangleq G - \{v\}$.

Appendix B. Cliques and general motifs for neighborhood sampling

In this appendix, we extend the results in Section 2.2 for neighborhood sampling from edges to general motifs. For ease of exposition, we start by developing the methodology for estimating cliques counts. Both the procedure and the performance guarantee readily extend to general motifs.

We now generalize the techniques for counting edges to estimate the number of cliques of size ω in a given graph. Note that there are two types of colored cliques one observe: (a) K_ω° : all but one vertex are sampled; (b) K_ω^\bullet : all vertices are sampled, with the first one being more probable when the sampling ratio is small. In the case of triangles, we have $K_3^\circ = \triangleleft$ and $K_3^\bullet = \blacktriangleleft$. Analogous to the estimator (13), we take a linear combination of these two types of clique counts as the linear estimator:

$$\widehat{s} = \alpha N(K_\omega^\circ, \widetilde{G}) + \beta N(K_\omega^\bullet, \widetilde{G}). \quad (18)$$

Similar to the design principles for counting edges, in the low sampling ratio regime $p < \frac{1}{d}$, we implement the Horvitz-Thompson estimator, so that the coefficients scale like $p^{-\omega}$; in the high sampling ratio regime $p > \frac{1}{d}$, we choose a *negative* β , which scale as $-p^{-2\omega}$, to reduce the correlation between various observed cliques. However, unlike the case of counting edges, we cannot perfectly eliminate all covariance terms but will be able to remove the leading one.

The following result, which includes Theorem 2 as a special case ($\omega = 2$), gives the performance guarantee of the estimator (18) and establishes its optimality in the worst case:

Theorem 3 (Cliques) *Set*

$$\begin{cases} \alpha = \frac{1}{p^{\omega-1}}, & \beta = \frac{1-\omega p}{p^\omega} & \text{if } p > 1/d \\ \alpha = \frac{1}{\omega p^{\omega-1}}, & \beta = \frac{1}{p^\omega} & \text{if } p \leq 1/d. \end{cases} \quad (19)$$

Then

$$\mathbb{E}_G |\widehat{s} - s(K_\omega, G)|^2 = \text{Var}_G[\widehat{s}] \leq s(K_\omega, G) \cdot \omega^3 2^{\omega+1} \left(\frac{d(G)}{p^{\omega-1}} \wedge \frac{d(G)^{\omega-2}}{p^2} \right).$$

Furthermore,

$$\inf_{\widehat{s}} \sup_{\substack{G: d(G) \leq d \\ s(K_\omega, G) \leq s}} \mathbb{E}_G |\widehat{s} - s(K_\omega, G)|^2 = \Theta_\omega \left(\frac{sd}{p^{\omega-1}} \wedge \frac{sd^{\omega-2}}{p^2} \wedge s^2 \right)$$

For brevity, we defer its proof to Appendix I.

To extend to general motif h on k vertices, note that in the neighborhood sampled graph, again it is possible to observe fully sampled or partially sampled (with one unsampled vertices) motifs. Consider the following estimator analogous to (18):

$$\widehat{s}_h = \alpha N(h^\circ, \widetilde{G}) + \beta N(h^\bullet, \widetilde{G}), \quad (20)$$

where $N(h^\bullet, \widetilde{G})$ is the count of h with all vertices sampled and $N(h^\circ, \widetilde{G})$ is the total count of h with exactly one unsampled vertex. For instance, if $h = \text{triangle}$, then $N(h^\bullet, \widetilde{G}) = N(\text{triangle}, G)$ and $N(h^\circ, \widetilde{G}) = N(\text{triangle with one vertex unsampled}, G) + N(\text{triangle with two vertices unsampled}, G)$. This example shows that in general, for motifs with less symmetry, there exist multiple partially sampled motifs and in principle they can be weighted differently. However, in (20) we elect to treat them equally, which turns out to be optimal for a wide class of motifs. Let us point out that if the parent graph has more structures, e.g., forest, then distinguishing different partially sampled motifs can lead to strict improvement; see Theorem 16.

The estimator (20) turns out to satisfy the same bound as in the clique case. To see this, note that in (47), the covariance terms are given in (46) which do not depend on the actual motif h . Furthermore, the sum of the indicators satisfies the same bound in terms of maximal degree provided that h is connected. Using the same optimized coefficients as in (19), the guarantee in Theorem 3 holds verbatim:

$$\mathbb{E}_G |\widehat{s}_h - s(h, G)|^2 = \text{Var}_G[\widehat{s}_h] \leq s(h, G) \cdot k^3 2^{k+1} \left(\frac{d(G)}{p^{k-1}} \wedge \frac{d(G)^{k-2}}{p^2} \right). \quad (21)$$

We conjecture that, similar to Theorem 1, this rate is optimal as long as the motif h is connected. So far we are able to prove this for cliques of all sizes (Theorem 12) and motifs on at most 4 vertices (Appendix K).

Appendix C. Adaptation to the maximum degree

In practice, the bound on the maximum degree d is likely unknown to the observer and obtaining a consistent estimate might be difficult if the high-degree vertices are rare. For example, in a star, most of the vertices have degree one except for the root. Even if a consistent estimate is obtained, it is unclear how to avoid it correlating with the data used to form \widehat{e} . Because \widehat{e} has the form of a sum, such correlations increase the number of cross terms in its variance decomposition.

To overcome these difficulties, we weight each observed edge according to the size of the neighborhood of its incident vertices. Once a vertex is sampled, its degree is exactly determined and thus incorporating this information does not introduce any additional randomness. This observations leads to the following adaptive estimator which achieves a risk that is similar to the optimal risk in Theorem 2:

Theorem 4 *Let \widehat{e} be given in (10) with $f(x) = \frac{px+q}{p(px+2q)}$ and $g(x, y) = \frac{1-pq(f(x)+f(y))}{p^2}$. Then for any graph G on N vertices and maximum degree bounded by d , \widehat{e} is an unbiased estimator of $e(G)$ and*

$$\text{Var}[\widehat{e}] \lesssim \frac{Nd}{p^2} \wedge \frac{e(G)d}{p}.$$

Remark 5 *The variance bound from Theorem 4 is weaker than Theorem 2 in the $p > 1/d$ regime – $\frac{Nd}{p^2}$ versus $\frac{e(G)}{p^2}$. They have the largest disparity when G consists of $N/(d+1)$ copies of the star graph S_{d+1} , in which case $e(G) = Nd/(d+1)$. This is due to the fact that with high probability $1-p$, all sampled vertices from S_{d+1} have degree one. Ideally, we would like to know the degree of the root of the star; however this is impossible unless the root is sampled. Nonetheless, we can still find a good estimate. More generally, in addition to using the degree d_u from a sampled vertex u , we may modify the estimator to incorporate degree information from a non-sampled vertex via an unbiased estimate, i.e., $\widehat{d}_u = \frac{|N_{\tilde{G}}(u)|}{p} = \sum_{v \in N_G(u)} \frac{b_v}{p}$. For example, we can redefine \mathcal{K}_A from (10) as*

$$\mathcal{K}_A = b_u(1 - b_v)f(d_u \vee \widehat{d}_v) + b_v(1 - b_u)f(d_u \vee \widehat{d}_u) + b_u b_v g(d_u, d_v).$$

We choose not to analyze such modified estimators here and instead leave their theoretical properties as an open question for future investigation.

Appendix D. Lower bounds

Throughout this appendix we assume that the sampling ratio p is bounded away from one.

D.1. Auxiliary results

We start with a result which is the general strategy of proving all lower bounds in this paper. A variant of this result was proved in Klusowski and Wu (2018) for the Bernoulli sampling model, however, an examination of the proof reveals that the conclusions also hold for neighborhood sampling. In the context of estimating motif counts, the essential ingredients involve constructing a pair of random graphs whose motif counts have different average values, and the distributions of their sampled versions are close in total variation, which is ensured by matching lower-order subgraphs counts in terms of \mathbf{s} for subgraph sampling or \mathbf{N} for neighborhood sampling. The utility of this result is to use a pair of smaller graphs (which can be found in an ad hoc manner) to construct a bigger pair of graphs and produce a lower bound that scales with an arbitrary positive integer s .

Theorem 6 (Theorem 11 in Klusowski and Wu (2018)) *Let f be a graph parameter that is invariant under isomorphisms and additive under disjoint union, i.e., $f(G + H) = f(G) + f(H)$. Fix a subgraph h . Let d, s, m and $M = s/m$ be integers. Let H and H' be two graphs such that $\mathbf{s}(h, H) \vee \mathbf{s}(h, H') \leq m$ and $\mathbf{d}(H) \vee \mathbf{d}(H') \leq d$. Suppose $M \geq 300$ and $\text{TV}(P, P') \leq 1/300$, where P (resp. P') denote the distribution of the isomorphism class of the (subgraph or neighborhood) sampled graph \tilde{H} (resp. \tilde{H}'). Let \tilde{G} denote the sampled version of G under the Bernoulli or neighborhood sampling models with probability p . Then*

$$\inf_{\widehat{f}} \sup_{\substack{G: \mathbf{d}(G) \leq d \\ \mathbf{s}(h, G) \leq s}} \mathbb{P}_G \left[|\widehat{f}(\tilde{G}) - f(G)| \geq \Delta \right] \geq 0.01. \quad (22)$$

where

$$\Delta = \frac{|f(H) - f(H')|}{8} \left(\sqrt{\frac{s}{m \text{TV}(P, P')}} \wedge \frac{s}{m} \right).$$

Next we recall the well-known fact [Whitney \(1932\)](#); [Kocay \(1982\)](#) that disconnected subgraphs counts are determined by (fixed polynomials of) connected subgraph counts. The following version is from [Klusowski and Wu \(2018, Corollary 1 and Lemma 9\)](#):

Lemma 7 *Let H and H' be two graphs with m vertices and $v \leq m$. Suppose $\mathfrak{s}(h, H) = \mathfrak{s}(h, H')$ for all connected h with $\mathfrak{v}(h) \leq v$. Then $\mathfrak{s}(h, H) = \mathfrak{s}(h, H')$ for all h with $\mathfrak{v}(h) \leq v$ and, furthermore,*

$$\text{TV}(P, P') \leq \binom{m}{v+1} p^{v+1},$$

where P (resp. P') denote the distribution of the isomorphism class of the subgraph sampled graph \tilde{H} (resp. \tilde{H}') with sampling ratio p .

The following version is for neighborhood sampling, which will be used in the proof of [Theorem 12](#). We need to develop an analogous result that expresses disconnected neighborhood subgraph counts as polynomials of the connected cones. This is done in [Lemma 20](#) in [Appendix H](#).

Lemma 8 *Let H and H' be two graphs with m vertices and $v \leq m$. Suppose $\mathbf{N}(h, H) = \mathbf{N}(h, H')$ for all connected, bicolored h with $\mathfrak{v}_b(h) \leq v$. Then*

$$\mathbf{N}(h, H) = \mathbf{N}(h, H') \tag{23}$$

for all h with $\mathfrak{v}_b(h) \leq v$ and, furthermore,

$$\text{TV}(P, P') \leq \binom{m}{v+1} p^{v+1}, \tag{24}$$

where P (resp. P') denote the distribution of the isomorphism class of the sampled graph \tilde{H} (resp. \tilde{H}') generated from neighborhood sampling with sampling ratio p .

Proof The first conclusion [\(23\)](#) follows from [Lemma 20](#). For the second conclusion [\(24\)](#), we note that conditioned on ℓ vertices are sampled, \tilde{H} is uniformly distributed over the collection of all bicolored neighborhood subgraphs h with $\mathfrak{v}_b(h) = \ell$. Thus,

$$\mathbb{P} \left[\tilde{H} \cong h \mid \mathfrak{v}_b(h) = \ell \right] = \frac{\mathbf{N}(h, H)}{\binom{m}{\ell}}.$$

By [\(23\)](#), we conclude that the isomorphism class of \tilde{H} and \tilde{H}' have the same distribution provided that no more than v vertices are sampled. Thus, $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq \mathbb{P}[\text{Bin}(m, p) \leq v+1]$, and consequently, $\mathbb{P}[\text{Bin}(m, p) \leq v+1] \leq \binom{m}{v+1} p^{v+1}$ follows from a union bound. \blacksquare

Lemma 9 *For any connected graph h with k vertices, there exists a pair of (in fact, connected) graphs H and H' , such that $\mathfrak{s}(h, H) \neq \mathfrak{s}(h, H')$ and $\mathfrak{s}(g, H) = \mathfrak{s}(g, H')$ for all connected g with $\mathfrak{v}(g) \leq k-1$.*

Proof The existence of such a pair H and H' follows from the strong independence² of connected subgraph counts Erdős et al. (1979, Theorem 1). For example, for $h = \triangle$, we can take the ad hoc construction in (2), which have equal number of vertices and edges but distinct number of triangles. Alternatively, next we provide an explicit construction using a linear algebra argument which is similar to that of Erdős et al. (1979, Theorem 3) and Biggs (1978, Section 2). Let $\{h_1, \dots, h_m\}$ denote all distinct (up to isomorphism) induced *connected* subgraph of h , ordered in increasing number of edges (arbitrarily among graphs with the same number of edges) so that h_1 is an isolated vertex and $h_m = h$. Then the matrix $B = (b_{ij})$ with $b_{ij} = \mathfrak{s}(h_i, h_j)$ is upper triangular with strictly positive diagonals. Thus B is invertible and the entries of B^{-1} are rational. Let $x = B^{-1}e_m$, where $e_m = (0, \dots, 0, 1)$. Then $x_m = 1$ since $b_{mm} = \mathfrak{s}(h, h) = 1$. Let $w = \alpha x \in \mathbb{Z}^m$, where $\alpha \in \mathbb{N}$ is the lowest common denominator of the entries of x . Now define H and H' as the disjoint union with weights given by the vector w :

$$H = \sum_{i=1}^m \max\{w_i, 0\}h_i, \quad H' = \sum_{i=1}^m \max\{-w_i, 0\}h_i. \quad (25)$$

By design, any connected induced subgraph of H and H' with at most $k - 1$ vertices belongs to $\{h_1, \dots, h_{m-1}\}$. For any $1 \leq i \leq m - 1$, since h_i is connected, we have $\mathfrak{s}(h_i, H) - \mathfrak{s}(h_i, H') = \sum_{j=1}^m w_j \mathfrak{s}(h_i, h_j) = 0$, and $\mathfrak{s}(h, H') = 0$ and $\mathfrak{s}(h, H) = \alpha \geq 1$. For example, for $h = \square$, such a solution is given by:

$$H = \square + 6 \times \circ - \circ \quad H' = 4 \times \triangle + 4 \times \circ$$

which have matching subgraphs of order three (vertices, edges and triangles). For a construction for general cliques, see Klusowski and Wu (2018, Eq. (47)). ■

Next we present graph-theoretic results that are needed for proving lower bound under the neighborhood sampling model. First, we relate the neighborhood subgraph counts \mathbf{N} to the usual subgraph \mathbf{n} . Since \mathbf{N} is essentially subgraph counts with prescribed degree for the sampled vertices (cf. Lovász (2012, p. 62)), this can be done by inclusion-exclusion principle similar to (1) that expresses the induced subgraph counts \mathfrak{s} in terms of the subgraph counts \mathbf{n} ; however, the key difference here is that the size of the subgraphs that appear in the linear combination is not bounded a priori. For example,

$$\begin{aligned} \mathbf{N}(\overset{\bullet}{\circ} \circ, G) &= \text{number of degree-2 vertices in } G \\ &= \sum_{k \geq 2} (-1)^{k-2} \binom{k}{2} \mathbf{n}(S_{k+1}, G), \end{aligned}$$

where S_{k+1} is the star graph with k leaves. The following lemma is a general statement:

2. This means that the closure of the range of their normalized version (subgraph densities) has nonempty interior.

Lemma 10 *Let h be a bicolored connected neighborhood graph and h_0 denote the uncolored version. Then for any G ,*

$$\mathbf{N}(h, G) = \sum_g c(g, h) \mathbf{n}(g, G) \tag{26}$$

where the sum is over all (uncolored) g obtained from h by either adding edges incident to the black vertices in h or adding vertices connected to black vertices in h . In particular, the coefficients $c(g, h)$ do not depend on G .

Proof The proof is by the inclusion-exclusion principle and essentially similar to the argument in Section 5.2, in particular, the proof of Proposition 5.6(b) in Lovász (2012).

Recall the definition of the subgraph count $\mathbf{n}(H, G)$ in (5) in terms of counting distinct subsets. It will be convenient to work with the labeled version counting graph homomorphisms. The following definitions are largely from Lovász (2012, Chapter 5). We say ψ is an injective homomorphism from H to G , if $\psi : V(H) \rightarrow V(G)$ is injective, and $(u, v) \in E(H)$ if $(\psi(u), \psi(v)) \in E(G)$. Denote by $\text{inj}(H, G)$ the number of injective homomorphisms from H to G . Then $\text{inj}(H, G) = \mathbf{n}(H, G) \text{aut}(H)$, where $\text{aut}(H)$ denotes the number of automorphisms (i.e. isomorphisms to itself) for H . Furthermore, for neighborhood subgraph h , $\text{aut}(h)$ denotes the number of automorphisms for h that also preserve the colors. For example, $\text{aut}(\bullet \circ) = 2$ and $\text{aut}(\circ \circ) = 4$. Throughout the proof, ψ always denotes an injection.

We use the following version of the inclusion-exclusion principle Lovász (2012, Appendix A.1). Let S be a ground set and let $\{A_i : i \in S\}$ be a collection of sets. For each $I \subset S$, define $A_I \triangleq \bigcap_{i \in I} A_i$ and $B_I \triangleq A_I \setminus \bigcup_{i \notin I} A_i$; in words, B_I denotes those elements that belong to exactly those A_i for $i \in I$ and none other. Then we have

$$|A_I| = \sum_{J \subset I} |B_J| \tag{27}$$

$$|B_I| = \sum_{J \subset I} (-1)^{|J|-|I|} |A_J|. \tag{28}$$

Fix G . Let \mathcal{G} denote the collection of (uncolored) subgraphs that are “extensions” of h , obtained from h by either adding edges between the black vertices in h or adding vertices attached to black vertices in h . For example, for $h = \bullet \circ$, we have $\mathcal{G} = \{\circ \circ, \circ \circ, \circ \circ, \circ \circ, \dots\}$ is the collection of all stars. Let the g^* be the maximal subgraph of G that is in \mathcal{G} ; in other words, $\mathbf{n}(g, G) = 0$, for any other $g \in \mathcal{G}$ containing g^* as a subgraph.

Now we define the ground set to be the edge set of g^* . Let h_0 be the uncolored version of h , then $E(h_0) \subset E(g^*)$. For every $I \subset E(g^*)$, define $A_I \triangleq \{\psi : V(g^*) \rightarrow V(g) : (\psi(u), \psi(v)) \in E(G) \text{ if } (u, v) \in I\}$ and $B_I \triangleq \{\psi : V(g^*) \rightarrow V(g) : (\psi(u), \psi(v)) \in E(G) \text{ if and only if } (u, v) \in I\}$. The key observation is that $|B_{E(h_0)}| = \text{aut}(h) \mathbf{N}(h, G)$, and $|A_{E(g)}| = \text{inj}(g, G) = \text{aut}(g) \mathbf{n}(h, G)$. Applying the inclusion-exclusion principle (28) yields

$$\text{aut}(h) \mathbf{N}(h, G) = \sum_{g: g \supset h_0} (-1)^{|E(g)| - |E(h_0)|} \text{inj}(g, G).$$

proving the desired (26). ■

The next result is the counterpart of Lemma 9, which shows the existence of a pair of graphs with matching lower order neighborhood subgraph counts but contain distinct number of copies of a certain motif; however, unlike Lemma 9, so far we can only deal with the clique motifs. For example for $\omega = 3$, we can use the ad hoc construction in (2); both graphs have the same degree sequence but distinct number of triangles. For $\omega = 4$, we can choose

$$\begin{aligned} H &= \text{K}_4 + 3 \times \text{C}_3 + 12 \times \text{C}_2 + 12 \times \text{C}_1 \\ H' &= 6 \times \text{C}_3 + 12 \times \text{C}_2 + 4 \times \text{C}_1 + 4 \times \text{C}_0 \end{aligned} \tag{29}$$

It is straightforward (although extremely tedious!) to verify that $\mathbf{N}(h, H) = \mathbf{N}(h, H')$ for all neighborhood subgraphs h with at most 2 black vertices. The general result is as follows:

Lemma 11 *There exists two graphs H and H' such that $s(K_\omega, H) - s(K_\omega, H') \geq 1$ and $\mathbf{N}(h, H) = \mathbf{N}(h, H')$ for all neighborhood subgraphs h such that $v_b(h) \leq \omega - 2$.*

Proof First we show that there exist a pair of graphs H and H' such that $n(g, H) = n(g, H')$ for all connected graphs g with at most ω vertices except for the clique K_ω , and $n(g, H) = n(g, H') = 0$ for all connected graphs g with more than ω vertices. Analogous to the proof of Lemma 9, this either follows from the strong independence of injective graph homomorphism numbers Erdős et al. (1979), or from the following linear algebra argument. Let $\{h_1, \dots, h_m\}$ denote all distinct (up to isomorphism) *connected* graphs of at most ω vertices. Order the graphs in increasing number of edges (arbitrarily among graphs with the same number of edges) so that h_1 is an isolated vertex and $h_m = K_\omega$. Then the matrix $B = (b_{ij})$ with $b_{ij} = n(h_i, h_j)$ is upper triangular with strictly positive diagonals. Then H and H' can be constructed from the vector $x = B^{-1}e_m$ similar to (25); see (29) for a concrete example for K_4 . By design, each connected component of H and H' has at most ω vertices, we have $n(g, H) = n(g, H') = 0$ for all connected g with $v(g) > \omega$.

Next we show that the neighborhood subgraph counts are matched up to order $\omega - 2$. For each neighborhood subgraph h with $v_b(h) \leq \omega - 2$, by Lemma 10, we have $\mathbf{N}(h, H) = \sum_{g \in \mathcal{G}} c(g, h)n(g, H)$, where the coefficients $c(g, h)$ are independent of H , and \mathcal{G} contains all subgraphs obtained from h by adding edges incident to black vertices in h or attaching vertices to black vertices in h . The crucial observation is two-fold: (a) since $v_b(h) \leq \omega - 2$, there exists at least a pair of white vertices in h , which are not connected. Since no edges are added between white vertices, the collection \mathcal{G} excludes the full clique K_ω ; (b) for each $g \in \mathcal{G}$, if g contains more than ω vertices, then $n(g, H) = n(g, H) = 0$; if g contains at most ω vertices (and not K_ω by the previous point), then $n(g, H) = n(g, H)$ by design. Therefore we conclude that $\mathbf{N}(h, H) = \mathbf{N}(h, H')$ for all neighborhood subgraphs h with $v_b(h) \leq \omega - 2$. ■

D.2. Subgraph sampling

Next we prove the lower bound part of Theorem 1:

Proof Throughout the proof, we assume that both d and s are at least some sufficiently large constant that only depends on $k = v(h)$ and we use c, c', c_0, c_1, \dots to denote constants that possibly depend on k only. We consider two cases separately.

Case I: $p \leq 1/d$. Let H and H' be the pair of graphs from Lemma 9, such that $s(h, H) - s(h, H') \geq 1$ and $s(g, H) = s(g, H')$ for all induced subgraphs g with $v(g) \leq k-1$. Therefore, by Lemma 8, we have $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O_k(p^{k-1})$. Let $r = s(h, H)$ which is a constant only depending on k . Applying Theorem 6 with $M = \lfloor s/r \rfloor$ yields the lower bound

$$\inf_{\tilde{s}} \sup_{\substack{G: d(G) \leq d \\ s(h, G) \leq s}} \mathbb{E}_G |\tilde{s} - s(h, G)|^2 = \Omega_k \left(\frac{s}{p^k} \wedge s^2 \right). \quad (30)$$

Case II: $p > 1/d$. We construct a pair of graphs H and H' with maximum degree d such that $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq c'$, where c' is a constant depending only on k , $s(h, H') = 0$ and $c_1 \ell^{k-1}/p \leq s(h, H) \leq c_2 \ell^{k-1}/p$. Choosing $\ell = c_3((sp)^{\frac{1}{k-1}} \wedge d)$ for some small constant c_3 and applying Theorem 6, we obtain

$$\inf_{\tilde{s}} \sup_{\substack{G: d(G) \leq d \\ s(h, G) \leq s}} \mathbb{E}_G |\tilde{s}(H) - s(h, G)|^2 = \Omega \left(\frac{\ell^{k-1}s}{p} \right) = \Theta_k \left(\frac{sd^{k-1}}{p} \wedge s^2 \right). \quad (31)$$

Combining (30) and (31) completes the proof of the lower bound of Theorem 1.

It remains to construct H and H' . The idea of the construction is to expand each vertex in h into an independent set, which was used in the proof of Erdős et al. (1979, Lemma 5). Here, we also need to consider the possibility of expanding into a clique. Next consider two cases:

Suppose h satisfies the ‘‘distinct neighborhood’’ property, that is, for each $v \in V(h)$, $N_h(v)$ is a distinct subset of $v(h)$. Such h includes cliques, paths, cycles, etc. Pick an arbitrary vertex $u \in V(h)$. Let $\{S_v : v \in V(h)\}$ be a collection of disjoint subsets, so that $|S_u| = \lceil c/p \rceil$ and $|S_v| = \lceil cd \rceil$, where c is a constant that only depends on $v(h) = k$ such that $ck \leq 1$. Define a graph H with vertex set $\cup_{v \in v(h)} S_v$ by connecting each pair of $a \in S_u$ and $b \in S_v$ whenever $(u, v) \in E(h)$. In other words, H is obtained by blowing up each vertex in h into an independent set and each edge into a complete bipartite graph. Repeating the same construction with h replaced by $h - u$ yields H' , in which case S_u consists of isolated vertices. By construction, the maximum degree of both graphs is at most d . Note that $H - S_u = H' - S_u$. Thus the sampled graph of H and H' have the same law provided that none of the vertices in S_u are sampled. Applying Lemma 23, we conclude that $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq (1-p)^{c/p} \leq c'$ for all $p \leq 1/2$, where c' is a constant depending only on k .

Furthermore,

$$\begin{aligned} s(h, H') &= \sum_{T \cap S_u = \emptyset} \mathbb{1} \{H'[T] \simeq h\} + \sum_{T \cap S_u \neq \emptyset} \mathbb{1} \{H'[T] \simeq h\} \\ &\stackrel{(a)}{=} \sum_{T \cap S_u = \emptyset} \mathbb{1} \{H'[T] \simeq h\} = \sum_{T \cap S_u = \emptyset} \mathbb{1} \{H[T] \simeq h\}, \end{aligned}$$

where (a) follows from the fact that $H'[T]$ contains isolated vertices whenever $T \cap S_u \neq \emptyset$ while h is connected by assumption. Note that since $|T| = k$, if $T \cap S_u = \emptyset$, then there exists $t, t' \in T$ such that t, t' belong to the same independent set S_v for some v . By construction,

t and t' have the same neighborhood, contradicting $H[T] \simeq h$. Thus, we conclude that $\mathfrak{s}(h, H') = 0$. For H , we have

$$\mathfrak{s}(h, H) = \sum_{T \cap S_u \neq \emptyset} \mathbb{1}\{H[T] \simeq h\} \geq |S_u| \prod_{v \neq u} |S_v| \geq c^k \ell^{k-1}/p,$$

and, similarly, $\mathfrak{s}(h, H) \leq |S_u| (\sum_{v \neq u} |S_v|)^{k-1} \leq (2ck)^k \ell^{k-1}/p$.

Next suppose that h does have distinct neighborhoods, thus there exist $\{u_1, \dots, u_\ell\} \subset V(h)$ with $\ell \geq 2$ such that the neighborhood $N_{u_i}(h)$ are identical, denoted by T . Let $g \triangleq h[T] \vee \ell K_1$ is an induced (by $T \cup \{u_1, \dots, u_\ell\}$) subgraph of h . We define H with vertex set $\cup_{v \in v(h)} S_v$ by the same procedure as above, except now all vertices are expanded into a clique, with $|S_{u_1}| = \lceil c/p \rceil$ and $|S_v| = \lceil cd \rceil$ for $v \neq u$. Finally, as before, we connect each pair of $a \in S_u$ and $b \in S_v$ whenever $(u, v) \in E(h)$. Define H' by repeating the same construction with h replaced by $h - u_1$. Analogous to the above we have $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq c$ and it remains to show that $\mathfrak{s}(h, H') = 0$. Indeed, for any set T of k vertices that does not include any vertex from S_{u_1} , since S_{u_i} forms a clique and u_1, \dots, u_ℓ form an independent set in h , the number of induced g in $H'[T]$ is strictly less than that in h . Thus, there exists no $T \subset \cup_{v \neq u_1} S_v$ such that $H'[T]$ is isomorphic to h , and hence $\mathfrak{s}(h, H') = 0$. Entirely analogously, we have $\mathfrak{s}(h, H) = \Theta_k(\ell^{k-1}/p)$. \blacksquare

D.3. Neighborhood sampling

To illustrate the main idea, we only prove the lower bound for cliques. The proof for other motifs (of size up to four) is similar but involves several ad hoc constructions; see Appendix K.

Theorem 12 (Cliques) *For neighborhood sampling with sampling ratio p ,*

$$\inf_{\tilde{\mathfrak{s}}} \sup_{\substack{G: d(G) \leq d \\ \mathfrak{s}(h, G) \leq s}} \mathbb{E}_G |\tilde{\mathfrak{s}} - \mathfrak{s}(K_\omega, G)|^2 = \Theta_\omega \left(\left(\frac{sd}{p^{\omega-1}} \wedge \frac{sd^{\omega-2}}{p^2} \right) \wedge s^2 \right)$$

Proof For the lower bound, consider two cases. For simplicity, denote the minimax risk on the left-hand side by R .

Case I: $p > 1/d$. Applying Lemma 23 with G being the complete $(\omega - 2)$ -partite graph of $(\omega - 2)\ell$ vertices, $H_1 = K_{1/p, 1/p}$, and $H_2 = (2/p)K_1$, we obtain two graphs H and H' with $\mathfrak{s}(K_\omega, H) \asymp \frac{\ell^{\omega-2}}{p^2}$ and $\mathfrak{s}(K_\omega, H') = 0$, and $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq c < 1$ for all $p \leq 1/2$. By Theorem 6 with $M = s/(\ell^{\omega-2}/p^2)$, we obtain the lower bound $R \gtrsim \frac{s\ell^{\omega-2}}{p^2}$. Let $\ell = cd$ if $\frac{d^{\omega-2}}{p^2} \leq s$ and $\ell = c(p^2 s)^{\frac{1}{\omega-2}}$ if $\frac{d^{\omega-2}}{p^2} > s$, for some small constant c . In either case, we find that $\mathfrak{s}(K_\omega, H) \leq s$, $\mathfrak{s}(K_\omega, H') \leq s$, and $R \asymp \frac{s\ell^{\omega-2}}{p^2} \asymp \frac{sd^{\omega-2}}{p^2} \wedge s^2$.

Case II: $p \leq 1/d$. We use a different construction. Let $\ell = c(d \wedge s^{1/\omega})$ for some small constant c . Let H and H' be the two graphs from Lemma 11 such that $\mathfrak{s}(K_\omega, H) - \mathfrak{s}(K_\omega, H') \geq 1$ and $\mathbf{N}(h, H) = \mathbf{N}(h, H')$ for all neighborhood subgraphs h with $\mathfrak{v}_b(h) \leq \omega - 2$. By

Lemma 8, we have $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O_\omega(p^{\omega-1}) < 1$. Next we amplify the gap $|\mathfrak{s}(K_\omega, H) - \mathfrak{s}(K_\omega, H')| = \Omega(\ell^\omega)$ by expanding each vertex into an independent set, similar in what is done in the proof of Theorem 1. For each vertex in H , we associate ℓ distinct isolated vertices, and connect each pair of vertices by an edge if and only if they were connected in H . This defines a new graph F with $\ell v(H)$ vertices and similarly we construct F' from H' . In this way, the subgraph counts of F and F' also match up to order $\omega - 2$, and, in view of Lemma 8, $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O_\omega((\ell p)^{\omega-1})$. Furthermore, the number of cliques satisfies $\mathfrak{s}(K_\omega, F) = \mathfrak{s}(K_\omega, H)\ell^\omega$ and $\mathfrak{s}(K_\omega, F') = \mathfrak{s}(K_\omega, H')\ell^\omega$. Thus, $\mathfrak{s}(K_\omega, F) \asymp \mathfrak{s}(K_\omega, F') \asymp |\mathfrak{s}(K_\omega, H) - \mathfrak{s}(K_\omega, H')| = \ell^\omega$. Applying Theorem 6 with $M = s/\ell^\omega$ yields $R \gtrsim (\ell^\omega (\sqrt{\frac{s/\ell^\omega}{(p\ell)^{\omega-1}} \wedge \frac{s}{\ell^\omega}}))^2 \asymp \frac{s\ell}{p^{\omega-1}} \wedge s^2 \asymp \frac{sd}{p^{\omega-1}} \wedge \frac{s^{1+1/\omega}}{p^{\omega-1}} \wedge s^2 \asymp \frac{sd}{p^{\omega-1}} \wedge s^2$, where the last step follows from the assumption that $p \leq 1/d$. \blacksquare

Appendix E. Graphs with additional structures

In this appendix, we explore how estimation of motif counts can be improved by prior knowledge of the parent graph structure.

It is a priori unclear whether additional structures such as tree or planarity helps for estimating motif counts with neighborhood sampling. Nevertheless, for counting edges, in Theorem 14 we show that the Horvitz-Thompson estimator (6) can only be marginally improved, in the sense that the lower bound continues to hold up to a sub-polynomial factor $p^{o(1)}$ where $o(1)$ is uniformly vanishing as $p \rightarrow 0$. For planar graphs, Theorem 19 shows a similar statement.

In particular, for counting edges, we show that even if the parent graph is known to be a forest a priori, for neighborhood sampling, the bound in Theorem 2 remains optimal up to a subpolynomial factor in p . Similarly, for subgraph sampling, we cannot improve the rate in Theorem 1. We also discuss some results for planar graphs. In what follows, we let \mathcal{F} and \mathcal{P} denote the collection of all forests and planar graphs, respectively.

The next results shows that for estimating edge counts, even if it is known a priori that the parent graph is a forest, the risk in Theorem 1 and Theorem 2 cannot be improved in terms of the exponents on p . The proofs of all the following results are given in Appendix I.

Theorem 13 *For subgraph sampling with sampling ratio p ,*

$$\inf_{\hat{\mathbf{e}}} \sup_{G \in \mathcal{F}: \substack{d(G) \leq d \\ e(G) \leq m}} \mathbb{E}_G |\hat{\mathbf{e}} - \mathbf{e}(G)|^2 \asymp \left(\frac{m}{p^2} \vee \frac{md}{p} \right) \wedge m^2. \quad (32)$$

Theorem 14 *For neighborhood sampling with sampling ratio p ,*

$$\inf_{\hat{\mathbf{e}}} \sup_{G \in \mathcal{F}: \substack{d(G) \leq d \\ e(G) \leq m}} \mathbb{E}_G |\hat{\mathbf{e}} - \mathbf{e}(G)|^2 = \Omega \left(\frac{m}{p^{2+o(1)}} \wedge \frac{md}{p^{1+o(1)}} \wedge \frac{m^2}{p^{o(1)}} \right),$$

where $o(1) = 1/\sqrt{\log \frac{1}{p}}$ is with respect to $p \rightarrow 0$ and uniform in all other parameters.

For estimating the wedge count under subgraph sampling, the following result shows that the risk in Theorem 1 cannot be improved even if we know the parent graph is a forest.

Theorem 15 *For subgraph sampling with sampling ratio p ,*

$$\inf_{\widehat{w}} \sup_{\substack{G \in \mathcal{F}: d(G) \leq d \\ w(G) \leq w}} \mathbb{E}_G |\widehat{w} - w(G)|^2 \asymp \left(\frac{w}{p^3} \vee \frac{wd^2}{p} \right) \wedge w^2. \quad (33)$$

On the other hand, for neighborhood sampling, the tree structure can be exploited to improve the rate. Analogous to (13), we consider an estimator of the form

$$\widehat{w} = \lambda N(\triangleleft, \widetilde{G}) + \alpha N(\triangleleft_{\circ}, \widetilde{G}) + \beta N(\triangleleft_{\bullet}, \widetilde{G}), \quad (34)$$

If we weight \triangleleft and \triangleleft_{\circ} equally, i.e., $\alpha = \lambda$, this estimator reduces to (20) and hence inherits the same performance guarantee in (21), which by Theorem 31, is optimal. However, as will be seen in Theorem 17, there is added flexibility by this three-parameter family of estimators that produces improved bounds when the parent graphs satisfies certain additional structure. It should also be mentioned that the alternative choices $\lambda = \frac{5-8p}{p^2(4p-3)}$, $\alpha = \frac{1}{p^2}$, and $\beta = \frac{3p-2}{p^3(4p-3)}$ yield the same performance bound as in (21).

For this next result, we show that we can improve the performance of the wedge estimator (34) if the parent graph is a forest by choosing alternate values of the parameters: $\alpha = \frac{1}{2pq}$, $\lambda = \frac{1}{p^2}$, and $\beta = 0$. These choices eliminate the largest term in the variance of (34), which is proportional to $n(S_4, G)(4\alpha^2 p^3 q^2 + 4\lambda\beta p^4 q + \beta^2 p^5 + p^4 q \lambda^2 - 1)$. We immediately get the following variance bound:

$$\text{Var}[\widehat{w}] \lesssim \frac{w(G)}{p^2} \vee \frac{w(G)d}{p}. \quad (35)$$

Note also that $s(P_3, G) = \sum_u \binom{d_u}{2}$ whenever G is a forest. Hence another estimator we can use is $\sum_u \frac{b_u}{p} \binom{d_u}{2}$ which has variance of order $\frac{w(G)d^2}{p}$. Putting this all together, we obtain the following result.

Theorem 16 *For neighborhood sampling with sampling ratio p ,*

$$\inf_{\widehat{w}} \sup_{\substack{G \in \mathcal{F}: d(G) \leq d \\ w(G) \leq w}} \mathbb{E}_G |\widehat{w} - w(G)|^2 \lesssim \left(\frac{w}{p^2} \vee \frac{wd}{p} \right) \wedge \left(\frac{wd^2}{p} \right) \wedge w^2.$$

The next theorem shows that the minimax bound from Theorem 16 is optimal.

Theorem 17 *For neighborhood sampling with sampling ratio p and $w \geq d$,*

$$\inf_{\widehat{w}} \sup_{\substack{G \in \mathcal{F}: d(G) \leq d \\ w(G) \leq w}} \mathbb{E}_G |\widehat{w} - w(G)|^2 = \Omega \left(\left(\frac{w}{p^2} \vee \frac{wd}{p} \right) \wedge \left(\frac{wd^2}{p} \right) \wedge w^2 \right).$$

In the context of estimating triangles, the next set of results show that planarity improves the rates of estimation for both sampling models. Despite the smaller risk however, for subgraph sampling, the optimal estimator is still the Horvitz-Thompson type.

Theorem 18 For subgraph sampling with sampling ratio p ,


$$\inf_{\hat{t}} \sup_{G \in \mathcal{P}: \substack{d(G) \leq d \\ t(G) \leq t}} \mathbb{E}_G |\hat{t} - t(G)|^2 \asymp \left(\frac{t}{p^3} \vee \frac{td}{p^2} \right) \wedge t^2.$$

Theorem 19 For neighborhood sampling with sampling ratio p ,


$$\left(\left(\frac{t}{p^{7/3}} \wedge \frac{td}{p^2} \right) \vee \frac{td}{p} \right) \wedge t^2 \lesssim \inf_{\hat{t}} \sup_{G \in \mathcal{P}: \substack{d(G) \leq d \\ t(G) \leq t}} \mathbb{E}_G |\hat{t} - t(G)|^2 \lesssim \left(\left(\frac{t}{p^3} \wedge \frac{td}{p^2} \right) \vee \frac{td}{p} \right) \wedge t^2.$$

Appendix F. Open problems


Here we mention a number of interesting questions that remain open:

- As mentioned in the introduction, a more general (and powerful) version of the neighborhood sampling model is to observe a labeled radius- r ball rooted at a randomly chosen vertex [Lovász \(2012\)](#). The current paper focuses on the case of $r = 1$. For $r = 2$, we note for example that a triangle could be observed simply by sampling only one of its vertices, i.e., . Thus, a Horvitz-Thompson type of estimator is $\frac{1}{3p} \mathbf{N}(\triangle, \tilde{G})$ and the variance scales as $1/p$. When p is small, this outperforms the neighborhood sampling counterpart ($r = 1$) in [Theorem 3](#), where the variance scales as $1/p^2$. Understanding the statistical limits of r -hop neighborhood sampling is an interesting and challenging research direction. In particular, the lower bound will potentially involve more complicated graph statistics as opposed neighborhood subgraph counts. For example, this may involve developing a theory for graphs generated by unions of r -hop neighborhoods and algebraic properties of their subgraph counts.
- In this paper we have focused on counting motifs as induced subgraphs. As shown in [\(1\)](#), subgraph counts can be expressed as linear combinations of induced subgraph counts. However, this does not necessarily mean their sample complexity are the same. Although we do not have a systematic understanding so far, here is a concrete example that demonstrates this: consider estimating the number of (not necessarily) 4-cycles with neighborhood sampling. Note that to observe a C_4 one only needs to sample the two diagonal vertices. Thus, a simple unbiased estimator is $\frac{1}{2p^2} \mathbf{n}(\square, \tilde{G})$, whose variance scales as $O(1/p^2)$ and is much smaller than the best error rate for estimating induced C_4 's which scales as $1/p^3$, as given by [Theorem 12](#). The explanation for this phenomenon is that although we have the deterministic relationship $\mathbf{n}(\square, G) = \mathbf{s}(\square, G) + \mathbf{s}(\square, G) + \mathbf{s}(\square, G)$ and each of the three subgraph counts can be estimated at the rate of p^{-3} , the statistical errors cancel each other and result in a faster rate.


Appendix G. Additional experimental results




(a) Jazz network (subgraph sampling).



(b) Jazz network (neighborhood sampling).




(c) Erdős-Rényi graph (subgraph sampling).




(d) Erdős-Rényi graph (neighborhood sampling).


Figure 3: Relative error of counting triangles. In Fig. 3a and Fig. 3b, the parent graph is the jazz network used in Section 3 with $t(G) = 17899$. In Fig. 3c and Fig. 3d, the parent graph is a realization of the Erdős-Rényi graph $\mathcal{G}(1000, 0.02)$ with $d = 35$, and $t(G) = 1319$.




(a) Jazz network (subgraph sampling).



(b) Jazz network (neighborhood sampling).



(c) Erdős-Rényi graph (subgraph sampling).



(d) Erdős-Rényi graph (neighborhood sampling).

Figure 4: Relative error of counting wedges. In Fig. 4a and Fig. 4b, the parent graph is the jazz network used in Section 3 with $w(G) = 103212$. In Fig. 4c and Fig. 4d, the parent graph is a realization of the Erdős-Rényi graph $\mathcal{G}(1000, 0.0015)$ with $d = 7$, and $w(G) = 1187$.

Appendix H. Auxiliary lemmas

Lemma 20 (Kocay’s Edge Theorem for Colored Graphs) *Let h be a bicolored disconnected graph. Then $N(h, G)$ can be expressed as a polynomial, independent of G , in $N(g, G)$, where g is bicolored, connected, and $v_b(g) \leq v_b(h)$. Moreover, if $\prod_{g \in \mathcal{G}} N(g, G)$ is a term in the polynomial, then $\sum_{g \in \mathcal{G}} v_b(g) \leq v_b(h)$ and the corresponding coefficient is bounded by $3^{[v_b(h)]^2}$. The number of terms in the polynomial representation is bounded by*

the number of $\mathbf{v}_b(h)$ -tuples $(g_1, \dots, g_{\mathbf{v}_b(h)})$ of all bicolored neighborhood subgraphs such that $\sum_{i=1}^{\mathbf{v}_b(h)} \mathbf{v}_b(g_i) \leq \mathbf{v}_b(h)$ and $\mathbf{N}(g_i, h) \neq 0$.

Proof For a disconnected graph g' , note that g' can be decomposed into two graphs g'_1 and g'_2 , where g'_1 is connected and $\mathbf{v}_b(g'_2) \leq \mathbf{v}_b(g') - 1$. Then,

$$\mathbf{N}(g'_1, G)\mathbf{N}(g'_2, G) = \sum_g a_g \mathbf{N}(g, G), \quad (36)$$

where the sum runs over all graphs g with $\mathbf{v}_b(g) \leq \mathbf{v}_b(g'_1) + \mathbf{v}_b(g'_2) = \mathbf{v}_b(g')$ and a_g is the number of decompositions of $V(g)$ into $V(g'_1) \cup V(g'_2)$ and $V_b(g)$ into $V_b(g'_1) \cup V_b(g'_2)$ (not necessarily disjoint) such that $g\{V_b(g'_1)\} \cong g'_1$ and $g\{V_b(g'_2)\} \cong g'_2$.

The only disconnected graph satisfying the above decomposition property for $\mathbf{v}_b(g) = \mathbf{v}_b(g')$ is $g \cong g'$, and hence

$$\mathbf{N}(g', G) = \frac{1}{a_{g'}} \left[\mathbf{N}(g'_1, G)\mathbf{N}(g'_2, G) - \sum_g a_g \mathbf{N}(g, G) \right], \quad (37)$$

where $\mathbf{v}_b(g'_2) \leq \mathbf{v}_b(g') - 1$ and the sum ranges over all g that are either connected and $\mathbf{v}_b(g) \leq \mathbf{v}_b(g')$ or disconnected and $\mathbf{v}_b(g) \leq \mathbf{v}_b(g') - 1$. Furthermore, each a_g can be bounded by the number of ways of decomposing a set of size $\mathbf{v}_b(g')$ into two sets (with possible overlap), or $3^{\mathbf{v}_b(g')}$.

We will now prove the following claim using induction. Let h be a bicolored disconnected graph. For each $k < \mathbf{v}_b(h)$,

$$\mathbf{N}(h, G) = \sum_{\mathcal{G}} c_{\mathcal{G}} \prod_{g \in \mathcal{G}} \mathbf{N}(g, G), \quad (38)$$

where \mathcal{G} contains at least one disconnected g' for which $\mathbf{v}_b(g') \leq \mathbf{v}_b(h) - k$, $\sum_{g \in \mathcal{G}} \mathbf{v}_b(g) \leq \mathbf{v}_b(h)$, $|\mathcal{G}| \leq 3^{k\mathbf{v}_b(h)}$, and the number of terms is bounded by the number of k -tuples (g_1, \dots, g_k) of all bicolored neighborhood graphs such that $\sum_{i=1}^k \mathbf{v}_b(g_i) \leq \mathbf{v}_b(h)$ and $\mathbf{N}(g_i, h) \neq 0$.

The base case $k = 1$ is established by decomposing h into two graphs h_1 and h_2 with h_1 connected and $\mathbf{v}_b(h_2) \leq \mathbf{v}_b(h) - 1$ and applying (37) with $g' \cong h$, $g'_1 \cong h_1$, and $g'_2 \cong h_2$.

Next, suppose (38) holds. Then applying (37) to each disconnected g' , we have

$$\begin{aligned} \mathbf{N}(h, G) &= \sum_{\mathcal{G}} c_{\mathcal{G}} \mathbf{N}(g', G) \prod_g \mathbf{N}(g, G) \\ &= \sum_{\mathcal{G}} \frac{c_{\mathcal{G}}}{c_{g'}} [\mathbf{N}(g'_1, G)\mathbf{N}(g'_2, G) - \sum_{h'} c_{h'} \mathbf{N}(h', G)] \prod_g \mathbf{N}(g, G) \\ &= \sum_{\mathcal{G}} \frac{c_{\mathcal{G}}}{c_{g'}} \mathbf{N}(g'_1, G)\mathbf{N}(g'_2, G) \prod_g \mathbf{N}(g, G) - \sum_{\mathcal{G}} \sum_{h'} \frac{c_{\mathcal{G}} c_{h'}}{c_{g'}} \mathbf{N}(h', G) \prod_g \mathbf{N}(g, G). \end{aligned} \quad (39)$$

Note that $\mathbf{v}_b(g'_2) \leq \mathbf{v}_b(g') - 1 \leq \mathbf{v}_b(h) - (k + 1)$ and if h' is disconnected, then $\mathbf{v}_b(h') \leq \mathbf{v}_b(g') - 1 \leq \mathbf{v}_b(h) - (k + 1)$. Finally, we observe that (39) has the form

$$\sum_{\tilde{\mathcal{G}}} c_{\tilde{\mathcal{G}}} \prod_g \mathbf{N}(g, G), \quad (40)$$

where $\tilde{\mathcal{G}}$ contains at least one disconnected g' for which $\mathbf{v}_b(g') \leq \mathbf{v}_b(h) - (k+1)$, $\mathbf{v}_b(g') \leq \mathbf{v}_b(h) - (k+1)$, $\sum_{g \in \tilde{\mathcal{G}}} \mathbf{v}_b(g) \leq \mathbf{v}_b(h)$, and $|c_{\tilde{\mathcal{G}}}| \leq \left| \frac{c_{\mathcal{G}}}{c_{g'}} \right| \vee \left| \frac{c_{\mathcal{G}} c_{h'}}{c_{g'}} \right| \leq 3^{(k+1)\mathbf{v}_b(h)}$. The number of terms is bounded by the number of $(k+1)$ -tuples (g_1, \dots, g_{k+1}) of all bicolored neighborhood graphs such that $\sum_{i=1}^{k+1} \mathbf{v}_b(g_i) \leq \mathbf{v}_b(h)$ and $\mathbf{N}(g_i, h) \neq 0$. Repeat this until $k = \mathbf{v}_b(h)$ and so that the right hand side of (38) contains no disconnected g in its terms. \blacksquare

Lemma 21 *Let H and H' be two graphs on M vertices. Suppose there exists a constant $B > 0$ and positive integer k such that for each connected subgraph h ,*

$$|\mathbf{N}(h, H) - \mathbf{N}(h, H')| \leq BM^{\mathbf{v}_b(h)-k}.$$

Then for each subgraph h ,

$$|\mathbf{N}(h, H) - \mathbf{N}(h, H')| \leq BQ_h \mathbf{v}_b(h) 3^{[\mathbf{v}_b(h)]^2} M^{\mathbf{v}_b(h)-k},$$

where Q_h is the number of $\mathbf{v}_b(h)$ -tuples $(g_1, \dots, g_{\mathbf{v}_b(h)})$ of all bicolored neighborhood graphs such that $\sum_{i=1}^{\mathbf{v}_b(h)} \mathbf{v}_b(g_i) \leq \mathbf{v}_b(h)$ and $\mathbf{N}(g_i, H) \neq 0$ or $\mathbf{N}(g_i, H') \neq 0$.

Proof Let h be a disconnected subgraph. By Lemma 20, $\mathbf{N}(h, H) = \sum_{\mathcal{G}_h} c_{\mathcal{G}_h} \prod_{g \in \mathcal{G}_h} \mathbf{N}(g, H)$, where $\sum_{g \in \mathcal{G}_h} \mathbf{v}_b(g) \leq \mathbf{v}_b(h)$, g is connected, $|c_{\mathcal{G}_h}| \leq 3^{[\mathbf{v}_b(h)]^2}$, and the number of terms is bounded by Q_h . Thus, using the fact that if x_1, \dots, x_n and y_1, \dots, y_n are positive real numbers, $|x_1 \cdots x_n - y_1 \cdots y_n| \leq \sum_{i=1}^n |x_i - y_i| x_1 \cdots x_{i-1} y_{i+1} \cdots y_n$, we have

$$\begin{aligned} |\mathbf{N}(h, H) - \mathbf{N}(h, H')| &\leq \sum_{\mathcal{G}_h} |c_{\mathcal{G}_h}| \left| \prod_{g \in \mathcal{G}_h} \mathbf{N}(g, H) - \prod_{g \in \mathcal{G}_h} \mathbf{N}(g, H') \right| \\ &\leq \sum_{\mathcal{G}_h} |c_{\mathcal{G}_h}| \sum_i |\mathbf{N}(g_i, H) - \mathbf{N}(g_i, H')| \prod_{j \leq i-1} \mathbf{N}(g_j, H) \prod_{j \geq i+1} \mathbf{N}(g_j, H'), \end{aligned}$$

where $\{g_i\}$ is an ordering of $\{g\}_{g \in \mathcal{G}_h}$. Next, we use the fact that $\max\{\mathbf{N}(g, H), \mathbf{N}(g, H')\} \leq \binom{M}{\mathbf{v}_b(g)} \leq M^{\mathbf{v}_b(g)}$ to bound

$$\sum_i |\mathbf{N}(g_i, H) - \mathbf{N}(g_i, H')| \prod_{j \leq i-1} \mathbf{N}(g_j, H) \prod_{j \geq i+1} \mathbf{N}(g_j, H') \leq B|\mathcal{G}_h| M^{\sum_{g \in \mathcal{G}_h} \mathbf{v}_b(g)-k}$$

Since $|\mathcal{G}_h| \leq \sum_{g \in \mathcal{G}_h} \mathbf{v}_b(g) \leq \mathbf{v}_b(h)$, the above is further bounded by $B\mathbf{v}_b(h)M^{\mathbf{v}_b(h)-k}$. Thus,

$$\begin{aligned} |\mathbf{N}(h, H) - \mathbf{N}(h, H')| &\leq B\mathbf{v}_b(h)M^{\mathbf{v}_b(h)-k} \sum_{\mathcal{G}_h} |a_{\mathcal{G}_h}| \\ &\leq BQ_h \mathbf{v}_b(h) 3^{[\mathbf{v}_b(h)]^2} M^{\mathbf{v}_b(h)-k}. \end{aligned}$$

\blacksquare

Next we present two results on the total variation that will be used in the regime of $p > \frac{1}{d}$. The main idea is the following: if a subset T of vertices are not sampled, for subgraph

sampling, in the observed graph we delete all edges incident to T , i.e., the edge set of $G\{T\}$, and for neighborhood sampling, we delete all edges within T , that is, the edge set of $G[T]$. Therefore, for two parent graphs, if missing T leads to isomorphic graphs, then by a natural coupling, the total variation between the sampled graphs is at most the probability that T is not completely absent in the sample.

Lemma 22 *Let $G_\theta = K_{A,\Delta-\theta} + K_{B,\Delta+\theta}$ for integer θ between zero and Δ . Consider the neighborhood sampling model with sampling ratio p . Suppose $|\theta - \theta'| \asymp \sqrt{\frac{\Delta}{p}}$ and both A and B are at most $1/p$. For neighborhood sampling with sampling ratio p , there exists $0 < c < 1$ such that*

$$\text{TV}(P_{\tilde{G}_\theta}, P_{\tilde{G}_{\theta'}}) \leq c.$$

Proof Note that G_θ is the union of two complete bipartite graphs. Suppose that none of the $A + B$ “left” side vertices are sampled. Then G_θ can be described by $K_{A,X} + K_{B,Y} + (2\Delta - (X + Y))K_1$, where $(X, Y) \sim \text{Bin}(\Delta - \theta, p) \otimes \text{Bin}(\Delta + \theta, p)$. Thus, if $(X', Y') \sim \text{Bin}(\Delta - \theta', p) \otimes \text{Bin}(\Delta + \theta', p)$, then

$$\text{TV}(P_{\tilde{G}_\theta}, P_{\tilde{G}_{\theta'}}) \leq 1 - q^{A+B} + q^{A+B} \text{TV}(P_{(X,Y)}, P_{(X',Y')}).$$

Furthermore, observe that

$$\text{TV}(P_{(X,Y)}, P_{(X',Y')}) \leq \text{TV}(P_X, P_{X'}) + \text{TV}(P_Y, P_{Y'}),$$

where

$$\begin{aligned} \text{TV}(P_X, P_{X'}) &= \text{TV}(\text{Bin}(\Delta - \theta, p), \text{Bin}(\Delta - \theta', p)), \\ \text{TV}(P_Y, P_{Y'}) &= \text{TV}(\text{Bin}(\Delta + \theta, p), \text{Bin}(\Delta + \theta', p)). \end{aligned}$$

This shows that if $|\theta - \theta'| \asymp \sqrt{\frac{\Delta}{p}}$ and both A and B are $O(\frac{1}{p})$, then $\text{TV}(P_{\tilde{G}_\theta}, P_{\tilde{G}_{\theta'}})$ is less than a constant less than one. \blacksquare

Lemma 23 *Let G , H_1 , and H_2 be an arbitrary graphs and let $H = G \vee H_1$ and $H' = G \vee H_2$. If $v = v(H_1) = v(H_2) \leq 1/p$, then for neighborhood sampling with sampling ratio p ,*

$$\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq 1 - q^v \leq 1 - q^{1/p}, \quad (41)$$

More generally, for $H = (V, E)$ and $H' = (V, E')$ defined on the same set V of vertices, if $T \subset V$ is such that $(V \setminus T, E \setminus E(H[T]))$ and $(V \setminus T, E' \setminus E(H'[T]))$ are isomorphic, then (41) holds with $v = |T|$.

Proof Suppose that none of the v vertices in H_1 or H_2 are sampled. Then H_1 and H_2 are isomorphic to each other. Thus,

$$\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq \mathbb{P}[\text{at least one vertex in } H_1 \text{ or } H_2 \text{ is sampled}] = 1 - q^v.$$

The second claim follows from the same argument. \blacksquare

The following lemma, which was used in the proof of Theorems 14 and 16, relies on a number-theoretic fact:

Lemma 24 *There exist two sequences of integers $(\alpha_1, \dots, \alpha_{k+1})$ and $(\beta_1, \dots, \beta_{k+1})$ such that*

$$\sum_{x \in [k+1]} x^i \alpha_x = 0 \quad i = 0, 2, 3, \dots, k,$$

$$\sum_{x=1}^{k+1} x^i \beta_x = 0 \quad i = 0, 1, 3, \dots, k,$$

and

$$\sum_{x \in [k+1]} x \alpha_x = \text{lcm}(1, \dots, k+1),$$

$$\sum_{x \in [k+1]} x^2 \beta_x = \text{lcm}^2(1, \dots, k+1),$$

where lcm stands for the least common multiple. Moreover, there exists universal constants A and B such that

$$\sum_{x \in [k+1]} |\alpha_x| \leq A^k, \quad \sum_{x \in [k+1]} |\beta_x| \leq B^k. \quad (42)$$

Proof We first introduce the quantity

$$\gamma_i = \sum_{x=1}^k \frac{(-1)^{x+1}}{x^i} \binom{k}{x}.$$

The key observation is that $\sum_{x=0}^{k+1} (-1)^x \binom{k+1}{x} D(x) = 0$ for all polynomials D with degree less than or equal to k . Hence we can set

$$\alpha_x = \left(\gamma_1 - \frac{1}{x} \right) (-1)^x \binom{k+1}{x} \text{lcm}(1, \dots, k+1)$$

and

$$\beta_x = \left(\gamma_1^2 - \gamma_2 - \frac{\gamma_1}{x} + \frac{1}{x^2} \right) (-1)^x \binom{k+1}{x} \text{lcm}^2(1, \dots, k+1),$$

where $x = 1, 2, \dots, k+1$. A well-known number theoretic fact is that the least common multiple of the k integers is in fact significantly smaller than their product. In fact, we have the estimates [Nair \(1982\)](#), [Hanson \(1972\)](#)

$$2^{k-1} \leq \text{lcm}(1, \dots, k) \leq 3^k, \quad \text{for all } k \geq 1,$$

which shows (42). ■

Lemma 25 *For the two graphs H and H' from Theorem 14 constructed with $(\alpha_1, \dots, \alpha_{k+1})$ from Lemma 24, we have for neighborhood sampling with sampling ratio p ,*

$$\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O(pA^k + (p\ell A^k)^k),$$

provided $p\ell A^k < 1$.

Proof There are four types of connected subgraphs of H and H' : edge with one black vertex, edge with two black vertices, S_u , $u > 1$ with white center, S_u , $u > 1$ with black center. If g is an edge with one black vertex $\mathbf{N}(g, H) = 2\ell\alpha + \ell \sum_{x=1}^{k+1} xw_x$ and $\mathbf{N}(g, H') = 2\ell\alpha' + \ell \sum_{x=1}^{k+1} xw'_x$. If g is an edge with two black vertices $\mathbf{N}(g, H) = \ell\alpha$ and $\mathbf{N}(g, H') = \ell\alpha'$. If $g \cong S_u$ with white center, then $\mathbf{N}(g, H) = \sum_{x=1}^{k+1} w_x \binom{\ell x}{v_b(g)}$ and $\mathbf{N}(g, H') = \sum_{x=1}^{k+1} w'_x \binom{\ell x}{v_b(g)}$ and furthermore,

$$\begin{aligned} |\mathbf{N}(g, H) - \mathbf{N}(g, H')| &= \left| \sum_{x=1}^{k+1} w_x \binom{\ell x}{v_b(g)} - \sum_{x=1}^{k+1} w'_x \binom{\ell x}{v_b(g)} \right| \\ &= \frac{\ell}{v_b(g)} \left| \sum_{x=1}^{k+1} xw_x - \sum_{x=1}^{k+1} xw'_x \right|. \end{aligned}$$

If $g \cong S_u$ with black center, then

$$\begin{aligned} \mathbf{N}(g, H) &= \sum_{x=1}^{k+1} w_x \binom{\ell x}{v_b(g) - 1} \mathbb{1}\{\ell x = u\} \\ \mathbf{N}(g, H') &= \sum_{x=1}^{k+1} w'_x \binom{\ell x}{v_b(g) - 1} \mathbb{1}\{\ell x = u\} \end{aligned}$$

We find that $|\mathbf{N}(g, H) - \mathbf{N}(g, H')| \leq 2a^k(\ell(k+1))^{v_b(g)-1}$ and $|\mathbf{N}(g, H)| \leq 2a^k(\ell(k+1))^{v_b(g)}$.

Let $v = v(H) = v(H') \leq (\ell(k+1) + 1)a^k$. Then

$$\mathrm{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq \frac{1}{2} \sum_{h: v_b(h) \leq k} |\mathbf{N}(h, H) - \mathbf{N}(h, H')| p^{v_b(h)} q^{v-v_b(h)} + \mathbb{P}[\mathrm{Bin}(v, p) \geq k+1],$$

where the sum runs over all bicolored graphs with at most k black vertices. By Lemma 21, for each subgraph h ,

$$|\mathbf{N}(h, H) - \mathbf{N}(h, H')| \leq v_b(h) 3^{\lfloor v_b(h) \rfloor^2} (2v_b(h)a^k(k+3))^{v_b(h)} (\ell(k+1))^{v_b(h)-1},$$

where we used the bound $Q_h \leq \lfloor v_b(h)(k+3) \rfloor^{v_b(h)}$. Hence,

$$\begin{aligned} \mathrm{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) &\leq \frac{1}{2} \sum_{h: 1 \leq v_b(h) \leq k} |\mathbf{N}(h, H) - \mathbf{N}(h, H')| p^{v_b(h)} q^{v-v_b(h)} + \mathbb{P}[\mathrm{Bin}(v, p) \geq k+1] \\ &\leq \frac{1}{2} \sum_{h: 1 \leq v_b(h) \leq k} v_b(h) 3^{\lfloor v_b(h) \rfloor^2} (2v_b(h)a^k(k+3))^{v_b(h)} (\ell(k+1))^{v_b(h)-1} p^{v_b(h)} q^{v-v_b(h)} + \\ &\quad \mathbb{P}[\mathrm{Bin}((\ell(k+1) + 1)a^k, p) \geq k+1] \\ &\leq (pA^k) \sum_{v=0}^k (p\ell A^k)^v + \sum_{v=k+1}^{\infty} (p\ell A^k)^v \\ &= O(pA^k + (p\ell A^k)^{k+1}), \end{aligned}$$

for some constant $A > 0$ and provided $p\ell A^k < 1$. ■

Lemma 26 For the two graphs H and H' from Theorem 17 constructed with $(\beta_1, \dots, \beta_{k+1})$ from Lemma 24, we have for neighborhood sampling with sampling ratio p ,

$$\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O(pA^k + (p\ell A^k)^2 + (p\ell A^k)^k),$$

provided $p\ell A^k < 1$.

Proof There are two types of connected subgraphs of H and H' : S_u , $u > 1$ with white center and S_u , $u > 1$ with black center. If $g \cong S_u$ with white center, then $\mathbf{N}(g, H) = \sum_{x=1}^{k+1} w_x \binom{\ell x}{\mathbf{v}_b(g)}$ and $\mathbf{N}(g, H') = \sum_{x=1}^{k+1} w'_x \binom{\ell x}{\mathbf{v}_b(g)}$ and furthermore, since $\sum_{x=1}^{k+1} x^i w_x = \sum_{x=1}^{k+1} x^i w'_x$ for $i = 0, 1, 3, \dots, \mathbf{v}_b(g)$,

$$\begin{aligned} |\mathbf{N}(g, H) - \mathbf{N}(g, H')| &= \left| \sum_{x=1}^{k+1} w_x \binom{\ell x}{\mathbf{v}_b(g)} - \sum_{x=1}^{k+1} w'_x \binom{\ell x}{\mathbf{v}_b(g)} \right| \\ &= \frac{\ell^2}{\mathbf{v}_b(g)(\mathbf{v}_b(g) - 1)} \left| \sum_{x=1}^{k+1} x^2 w_x - \sum_{x=1}^{k+1} x^2 w'_x \right|. \end{aligned}$$

If $g \cong S_u$ with black center, then

$$\begin{aligned} \mathbf{N}(g, H) &= \sum_{x=1}^{k+1} w_x \binom{\ell x}{\mathbf{v}_b(g) - 1} \mathbb{1}\{\ell x = u\} \\ \mathbf{N}(g, H') &= \sum_{x=1}^{k+1} w'_x \binom{\ell x}{\mathbf{v}_b(g) - 1} \mathbb{1}\{\ell x = u\} \end{aligned}$$

We find that $|\mathbf{N}(g, H) - \mathbf{N}(g, H')| \leq 2a^k(\ell(k+1))^{\mathbf{v}_b(g)-1}$ and $|\mathbf{N}(g, H)| \leq a^k(\ell(k+1))^{\mathbf{v}_b(g)}$.

Let $v = \mathbf{v}(H) = \mathbf{v}(H') \leq (\ell(k+1) + 1)a^k$. Then

$$\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq \frac{1}{2} \sum_{h: \mathbf{v}_b(h) \leq k} |\mathbf{N}(h, H) - \mathbf{N}(h, H')| p^{\mathbf{v}_b(h)} q^{v - \mathbf{v}_b(h)} + \mathbb{P}[\text{Bin}(v, p) \geq k + 1],$$

where the sum runs over all bicolored graphs with at most k black vertices. By Lemma 21, for each subgraph h with $\mathbf{v}_b(h) \neq 2$,

$$|\mathbf{N}(h, H) - \mathbf{N}(h, H')| \leq \mathbf{v}_b(h) 3^{[\mathbf{v}_b(h)]^2} (2\mathbf{v}_b(h)a^k(k+3))^{\mathbf{v}_b(h)} (\ell(k+1))^{\mathbf{v}_b(h)-1},$$

where we used the bound $Q_h \leq [\mathbf{v}_b(h)(k+3)]^{\mathbf{v}_b(h)}$. Hence,

$$\begin{aligned} \text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) &\leq \frac{1}{2} \sum_{h: 1 \leq \mathbf{v}_b(h) \leq k} |\mathbf{N}(h, H) - \mathbf{N}(h, H')| p^{\mathbf{v}_b(h)} q^{v - \mathbf{v}_b(h)} + \mathbb{P}[\text{Bin}(v, p) \geq k + 1] \\ &\leq \frac{1}{2} \sum_{h: \mathbf{v}_b(h) \neq 2, \mathbf{v}_b(h) \leq k} \mathbf{v}_b(h) 3^{[\mathbf{v}_b(h)]^2} (2\mathbf{v}_b(h)a^k(k+3))^{\mathbf{v}_b(h)} (\ell(k+1))^{\mathbf{v}_b(h)-1} p^{\mathbf{v}_b(h)} q^{v - \mathbf{v}_b(h)} + \\ &\quad a^k \ell^2 p^2 + \mathbb{P}[\text{Bin}((\ell(k+1) + 1)a^k, p) \geq k + 1] \\ &\leq (pA^k) \sum_{v=0}^k (p\ell A^k)^v + (p\ell A^k)^2 + \sum_{v=k+1}^{\infty} (p\ell A^k)^v \\ &= O(pA^k + (p\ell A^k)^2 + (p\ell A^k)^{k+1}), \end{aligned}$$

for some constant $A > 0$ and provided $p\ell A^k < 1$. ■

Lemma 27 *There exists two planar graphs H and H' on order ℓ vertices with matching degree sequences and maximum degree equal to $\ell+1$ such that for neighborhood sampling with sampling ratio p , $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O(p^2 + p^3\ell^3)$ and $|\mathbf{w}(H) - \mathbf{w}(H')| = 3|\mathbf{t}(H) - \mathbf{t}(H')| \asymp \ell$ provided $p = O(1/\ell)$. Furthermore, there exists two planar graphs H and H' on order ℓ vertices such that for neighborhood sampling with sampling ratio p , $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O(p)$ and $|\mathbf{t}(H) - \mathbf{t}(H')| \asymp \ell$.*

Proof The proof follows from an examination of the two graphs below. Note that $\mathbf{N}(h, H) = \mathbf{N}(h, H')$ for all connected h with $\mathbf{v}_b(h) = 1$ and since $|\mathbf{N}(h, H) - \mathbf{N}(h, H')| = O(1)$ for all connected h with $\mathbf{v}_b(h) = 2$, it follows from Lemma 21 with $k = 2$ that $|\mathbf{N}(h, H) - \mathbf{N}(h, H')| = O(1)$ for all h with $\mathbf{v}_b(h) = 2$. Thus,

$$\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = \sum_h |\mathbf{N}(h, H) - \mathbf{N}(h, H')| p^{\mathbf{v}_b(h)} q^{v - \mathbf{v}_b(h)} = O(p^2 + \sum_{k=3}^{\infty} \ell^k p^k) = O(p^2 + p^3 \ell^3),$$

provided $p = O(1/\ell)$. The identity $|\mathbf{w}(H) - \mathbf{w}(H')| = 3|\mathbf{t}(H) - \mathbf{t}(H')| = \ell - 2$ follows from the fact that H and H' have matching degree sequences (corresponding to matching subgraphs from neighborhood sampling with one vertex).

For the second statement, consider two planar graphs H and H' on $\ell + 2$ vertices, where H consists of ℓ triangles sharing a common edge, and H' consists of ℓ wedges sharing a pair of non-adjacent vertices; see Fig. 5 for an illustration for $\ell = 5$. Note that if neither of




Figure 5: Example of H and H' with $\ell = 5$

the two highest-degree vertices in each graph (degree $\ell + 1$ in H and degree ℓ in H') are sampled and all incident edges removed, the two graphs are isomorphic. This shows that $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq 1 - q^2 = O(p)$. Also, note that $\mathbf{t}(H) = \ell$ and $\mathbf{t}(H') = 0$. ■

Table 2: The graph H with $\ell = 5$ and $d(H) = \ell + 1 = 6$

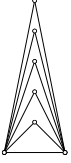
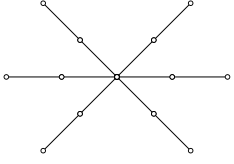
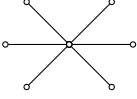

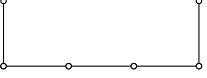
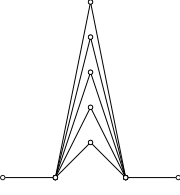
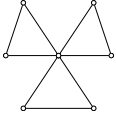
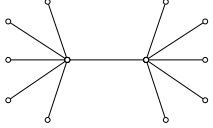
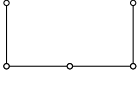

Copies	Components
1	
2	
2	
$\frac{\ell+1}{3}$	
$2(\ell + 1)$	

 Table 3: The graph H' with $\ell = 5$ and $d(H') = \ell + 1 = 6$

Copies	Components
1	
2	
1	
$\ell + 1$	
$2(\ell + 1)$	

Appendix I. Additional proofs

Proof [Proof of Theorem 3] Let $b_v \triangleq \mathbb{1}\{v \in S\} \stackrel{\text{i.i.d.}}{\sim} \text{Bern}(p)$. For any $T \subset V(G)$, let $b_T \triangleq \prod_{v \in T} b_v$. Write

$$\begin{aligned} \hat{s} &= \sum_{T \subset V(G)} \alpha \mathbb{1}\{\tilde{G}\{T\} \simeq K_\omega^\circ\} + \beta \mathbb{1}\{\tilde{G}\{T\} \simeq K_\omega^\bullet\} \\ &= \sum_{T \subset V(G)} f(T) \mathbb{1}\{G[T] \simeq K_\omega\}, \end{aligned} \quad (43)$$

where

$$f(T) \triangleq \alpha \sum_{v \in T} b_{T \setminus \{v\}} (1 - b_v) + \beta b_T.$$


Similar to (14), enforcing unbiasedness, we have the constraint $\mathbb{E}[f(T)] = 1$, i.e.,

$$\omega p^{\omega-1} q \alpha + p^\omega \beta = 1 \quad (44)$$

Furthermore, whenever $|T \cap T'| = t \in [\omega]$, we have

$$\mathbb{E}[f(T)f(T')] = \alpha^2 (tqp^{2\omega-t-1} + (\omega-t)^2 q^2 p^{2\omega-t-2}) + 2(\omega-t)qp^{2\omega-t-1}\alpha\beta + \beta^2 p^{2\omega-t} \quad (45)$$

$$\begin{aligned} &= p^{-t} \left[\alpha^2 tqp^{2\omega-1} + (\alpha(\omega-t)qp^{\omega-1} + \beta p^\omega)^2 \right] \\ &\stackrel{(44)}{=} p^{-t} \left[\alpha^2 tqp^{2\omega-1} + (1-tq\alpha p^{\omega-1})^2 \right] \end{aligned} \quad (46)$$

This follows from evaluating the probability of observing a pair of intersecting cliques with two, one, or zero unsampled vertices. For example, the four summands in (45), in the case of $\omega = 4$ and $t = 2$, correspond to , respectively.

Let $c_t \triangleq \text{Cov}[f(T), f(T')] = p^{-t} \left[\alpha^2 tqp^{2\omega-1} + (\alpha(\omega-t)qp^{\omega-1} + \beta p^\omega)^2 \right] - 1$ for $|T \cap T'| = t$. Denote by $T_{\omega,t}$ the subgraph correspond to two intersecting ω -cliques sharing t vertices. Then

$$\begin{aligned} \text{Var}[\widehat{s}] &= \sum_{T \cap T' \neq \emptyset} \text{Cov}(f(T), f(T')) \mathbb{1} \{G[T] \simeq K_\omega, G[T'] \simeq K_\omega\} \quad (47) \\ &= \sum_{t=1}^{\omega} c_t \sum_{|T \cap T'|=t} \mathbf{n}(T_{\omega,t}, G) \leq \mathbf{s}(K_\omega, G) d^\omega \sum_{t=1}^{\omega} c_t \binom{\omega}{t} d^{-t}. \end{aligned}$$

Next consider two cases separately.

Case I: $p \leq \frac{1}{d}$. In this case we choose $\alpha = \frac{1}{\omega p^{\omega-1}}$ and $\beta = \frac{1}{p^\omega}$. Then $c_t = p^{-t} \left(\frac{tpq}{\omega^2} + (1 - \frac{tpq}{\omega})^2 \right) \leq 2p^{-t}$. Furthermore, for the special case of $t = \omega$, we have $c_\omega \leq p^{-(\omega-1)}$. Thus,

$$\text{Var}[\widehat{s}] \leq \mathbf{s}(K_\omega, G) \left(d^\omega \sum_{t=1}^{\omega} \binom{\omega}{t} (pd)^{-t} + p^{-(\omega-1)} \right) \leq \mathbf{s}(K_\omega, G) \omega 2^{\omega+1} dp^{-(\omega-1)}. \quad (48)$$

Case II: $p \leq \frac{1}{d}$. In this high-degree regime, the pairs of cliques sharing one vertex ($t = 1$) dominate (i.e., open triangle for counting edge and bowties for counting triangles). Thus our strategy is to choose the coefficients to eliminate these covariance terms. In fact, (46) for $t = 1$ simplifies wonderfully to

$$c_1 = \frac{q}{p} (1 - \alpha p^{\omega-1})^2.$$

Thus we choose $\alpha = \frac{1}{p^{\omega-1}}$ and $\beta = \frac{1-\omega q}{p^\omega}$. Hence $c_t \leq 2\omega^2 p^{-t}$ for all $t \geq 2$, and

$$\text{Var}[\widehat{s}] \leq \mathbf{s}(K_\omega, G) d^\omega 2\omega^2 \sum_{t=2}^{\omega} \binom{\omega}{t} (pd)^{-t} \leq \mathbf{s}(K_\omega, G) 2^{\omega+1} \omega^3 d^{\omega-2} p^{-2}. \quad (49)$$

Combining (48) and (49) completes the proof. ■

Proof [Proof of Theorem 1] Since \widehat{s} is unbiased, it remains to bound its variance. Let $b_v \triangleq \mathbb{1} \{v \in S\}$, which are iid as $\text{Bern}(p)$. For any $T \subset V(G)$, let $b_T \triangleq \prod_{v \in T} b_v$. Then

$$\widehat{s} = p^{-k} \sum_{T \subset V(G)} b_T \mathbb{1} \{G[T] \simeq h\}.$$

Hence

$$\begin{aligned}
 \text{Var}[\widehat{s}] &= p^{-2k} \sum_{T \cap T' \neq \emptyset} \text{Cov}(b_T, b_{T'}) \mathbb{1} \{G[T] \simeq h, G[T'] \simeq h\} \\
 &\leq p^{-2k} \sum_{T \cap T' \neq \emptyset} \mathbb{E}[b_{T \cup T'}] \mathbb{1} \{G[T] \simeq h, G[T'] \simeq h\} \\
 &= \sum_{t=1}^k p^{-t} \sum_{|T \cap T'|=t} \mathbb{1} \{G[T] \simeq h, G[T'] \simeq h\} \\
 &\leq \sum_{t=1}^k p^{-t} s(h, G) \binom{k}{t} d^{k-t} \leq s(h, G) (2d)^k \cdot k \max\{(pd)^{-k}, (pd)^{-1}\},
 \end{aligned}$$

where the penultimate step follows from the fact that the maximum degree of G is d and, crucially, h is connected. \blacksquare

Proof [Proof of Theorems 13, 15, and 18] The upper bounds are achieved by Horvitz-Thompson estimation as in Theorem 1. However, for Theorem 15, we are able to achieve a smaller variance because $n(\text{triangle}, G)$ is of order td for planar G instead of td^2 and hence

$$\text{Var}[\widehat{t}_{HT}] \lesssim \frac{n(\text{triangle}, G)}{p^3} + \frac{n(\text{square}, G)}{p^2} + \frac{n(\text{pentagon}, G)}{p} \lesssim \frac{t}{p^3} + \frac{td}{p^2} + \frac{td}{p} \asymp \frac{t}{p^3} \vee \frac{td}{p^2}.$$

For the lower bound, the proof follows the same lines as Appendix D.2 in that we use two different constructions depending on whether $p \leq 1/d$ or $p > 1/d$.

For edges, let $H = S_\ell$ and $H' = (\ell + 1)S_1$ with $\ell = c(d \wedge m)$ for some small constant $c > 0$. Then $\text{TV}(P_{\widetilde{H}}, P_{\widetilde{H}'}) \leq p(1 - q^\ell) \leq p \wedge (\ell p^2)$.

For wedges, when $p \leq 1/d$, let $H = P_4 + K_1$ and $H' = P_3 + P_2$. Then $\text{TV}(P_{\widetilde{H}}, P_{\widetilde{H}'}) \leq O(p^3)$. When $p > 1/d$, let $H = S_\ell$ and $H' = (\ell + 1)K_1$. Then $\text{TV}(P_{\widetilde{H}}, P_{\widetilde{H}'}) \leq p$. Finally set $\ell = c(d \wedge w)$ for some universal constant $c > 0$.

Finally, for triangles, let H be the graph which consists of ℓ triangles that share the same edge plus ℓ isolated vertices. Let H' be the graph which consists of two S_ℓ star graphs with an edge between their roots. Choose $\ell = c(d \wedge t)$ for some small universal constant $c > 0$. Then $\text{TV}(P_{\widetilde{H}}, P_{\widetilde{H}'}) \leq p^2(1 - q^\ell) \leq p^2 \wedge (p^3 \ell)$. \blacksquare

Proof [Proof of Theorem 14] Let (w_1, \dots, w_{k+1}) and (w'_1, \dots, w'_{k+1}) be two sequences of integers defined by $w_x = \max\{\alpha_x, 0\}$ and $w'_x = \max\{-\alpha_x, 0\}$, where $(\alpha_1, \dots, \alpha_{k+1})$ is as in Lemma 24. Consider the disjoint union of stars

$$H \simeq \sum_{x=1}^{k+1} w_x S_{\ell_x} + \ell \alpha S_1 \quad \text{and} \quad H' \simeq \sum_{x=1}^{k+1} w'_x S_{\ell_x} + \ell \alpha' S_1,$$

for integer $\ell > 1$.

Note, for example, that $\mathbf{e}(H) = \ell(\sum_{x=1}^{k+1} xw_x + \alpha)$ and $\mathbf{v}(H) = \mathbf{e}(H) + \sum_{x=1}^{k+1} w_x + \ell\alpha = \sum_{x=1}^{k+1} (\ell x + 1)w_x + 2\alpha\ell$. Thus, $\mathbf{e}(H) \vee \mathbf{e}(H') \leq \ell a^k$ for some universal $a > 0$. Note that

$$\mathbf{e}(H) - \mathbf{e}(H') = \ell(\alpha - \alpha') = \frac{\ell}{2} \left(\sum_{x=1}^{k+1} xw_x - \sum_{x=1}^{k+1} xw'_x \right) \geq \frac{\ell}{2},$$

and by Lemma 25 there exists universal $A > 0$ such that

$$\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = \frac{1}{2} \sum_h |\mathbf{N}(h, H) - \mathbf{N}(h, H')| p^{\mathbf{v}_b(h)} q^{v - \mathbf{v}_b(h)} = O(pA^k + (p\ell A^k)^k),$$

provided $p\ell A^k < 1$.

By Theorem 6, we have

$$\inf_{\hat{\mathbf{e}}} \sup_{G \in \mathcal{F}: \mathbf{d}(G) \leq d, \mathbf{e}(G) \leq m} \mathbb{P} [|\hat{\mathbf{e}} - \mathbf{e}(G)| \geq \Delta_\ell] \geq c.$$

where

$$\begin{aligned} \Delta_\ell &\gtrsim |\mathbf{e}(H) - \mathbf{e}(H')| \left(\sqrt{\frac{m}{\mathbf{e}(H) \vee \mathbf{e}(H') \text{TV}(P_{\tilde{H}}, P_{\tilde{H}'})}} \wedge \frac{m}{\mathbf{e}(H) \vee \mathbf{e}(H')} \right) \\ &\gtrsim \sqrt{\frac{m\ell}{pc^k + (p\ell c^k)^k}} \wedge \frac{m}{c^k}, \end{aligned}$$

for some universal constants $c > 0$ provided $p\ell c^k < 1$. Next, choose

$$\ell = \begin{cases} \left(\frac{1}{pc^k} \right)^{1-1/k} \wedge \frac{m}{a^k} & \text{if } p > \left(\frac{1}{dc^k} \right)^{k/(k-1)} \\ d \wedge \frac{m}{a^k} & \text{if } p \leq \left(\frac{1}{dc^k} \right)^{k/(k-1)}. \end{cases} \quad (50)$$

Taking $k = \sqrt{\log \frac{1}{p}}$ yields the desired lower bound. \blacksquare

Proof [Proof of Theorem 17] Let (w_1, \dots, w_{k+1}) and (w'_1, \dots, w'_{k+1}) be two sequences of integers defined by $w_x = \max\{\beta_x, 0\}$ and $w'_x = \max\{-\beta_x, 0\}$, where $(\beta_1, \dots, \beta_{k+1})$ is as in Lemma 24. Let

$$H \simeq \sum_{x=1}^{k+1} w_x S_{\ell x} \quad \text{and} \quad H' \simeq \sum_{x=1}^{k+1} w'_x S_{\ell x},$$

for integer $\ell > 1$. Note, for example, that $\mathbf{e}(H) = \ell \sum_{x=1}^{k+1} xw_x$, $\mathbf{v}(H) = \sum_{x=1}^{k+1} (\ell x + 1)w_x$, and $\mathbf{w}(H) = \sum_{x=1}^{k+1} \binom{\ell x}{2} w_x$. This means that $\mathbf{w}(H) \vee \mathbf{w}(H') \leq \ell^2 a^{2k}$ for some universal $a > 0$. Note that

$$\mathbf{w}(H) - \mathbf{w}(H') = \frac{\ell^2}{2} \left(\sum_{x=1}^{k+1} x^2 w_x - \sum_{x=1}^{k+1} x^2 w'_x \right) \geq \frac{\ell^2}{2}.$$

By Lemma 26, we have that $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O(pA^k + (p\ell A^k)^2 + (p\ell A^k)^k)$ for some universal $A > 0$. By Theorem 6, we have

$$\inf_{\hat{w}} \sup_{G \in \mathcal{F}: \substack{d(G) \leq d \\ w(G) \leq w}} \mathbb{P}[|\hat{w} - w(G)| \geq \Delta_\ell] \geq c.$$

where

$$\begin{aligned} \Delta_\ell &\gtrsim |w(H) - w(H')| \left(\sqrt{\frac{w}{w(H) \vee w(H') \text{TV}(P_{\tilde{H}}, P_{\tilde{H}'})}} \wedge \frac{w}{w(H) \vee w(H')} \right) \\ &\gtrsim \sqrt{\frac{w\ell^2}{pc^k + (p\ell c^k)^2 + (p\ell c^k)^k}} \wedge \frac{w}{c^k}, \end{aligned}$$

for some universal constant $c > 0$. Next, choose $k = 2$ and $\ell = c(d \wedge w^{1/2})$ when $p \leq 1/d$ for some universal constant $c > 0$. For $p > 1/d$ and $w \geq d$, we use Lemma 22 with $A = B = 1$ and $\Delta = cd$. Then $w(H) \asymp w(H') \asymp d^2$ and $|w(H) - w(H')| \asymp d\sqrt{\frac{d}{p}}$, and $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) < c < 1$. By Theorem 6, we have $\inf_{\hat{w}} \sup_{G \in \mathcal{F}: \substack{d(G) \leq d \\ w(G) \leq w}} \mathbb{E}_G |\hat{w} - w(G)|^2 \gtrsim \frac{wd}{p}$. ■

Proof [Proof of Theorem 19] Let R denote the minimax risk. The bound $R \lesssim \frac{td}{p^2}$ follows immediately from Theorem 12 with $\omega = 3$. For the other regimes, we modify the estimator (43) from Theorem 3. To accomplish this, observe that $n(\triangle, G)$ is of order td for planar G , since the number of triangles that share a common vertex is at most d . Choosing $\alpha = \frac{1}{2qp^2}$ so that, in the notation of the proof of Theorem 3, $c_1 - 1 \asymp \frac{1}{p}$ and $c_2 - 1 = p^{-2} \left[2\alpha^2 qp^5 + (1 - 2q\alpha p^2)^2 \right] \asymp \frac{1}{p}$, we have $\text{Var}[\hat{s}] \lesssim \frac{t}{p^3} \vee \frac{td}{p}$. This yields the bound $R \lesssim \frac{t}{p^3} \vee \frac{td}{p}$. Thus, $R \lesssim \left(\frac{t}{p^3} \vee \frac{td}{p} \right) \wedge \frac{td}{p^2} = \left(\frac{t}{p^3} \wedge \frac{td}{p^2} \right) \vee \frac{td}{p}$. For the lower bound, consider two cases:

Case I: $p \leq 1/d$. By Lemma 27, there exists two planar graphs H and H' on order ℓ vertices such that $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O(p^2 + p^3\ell^3)$ and $t(H) \asymp t(H') \asymp |t(H) - t(H')| \asymp \ell$ provided $p = O(1/\ell)$. We choose $\ell = p^{-1/3} \wedge t$ if $p > 1/d^3$. Otherwise, if $p \leq 1/d^3$, we choose $\ell = d \wedge t$. By Theorem 6, this produces a lower bound of $R \gtrsim \left(\frac{t}{p^{7/3}} \wedge \frac{td}{p^2} \right) \wedge t^2$.

Case II: $p > 1/d$. We use the second statement of Lemma 27 which guarantees the existence of two planar graphs H and H' on order ℓ vertices such that $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) = O(p)$ and $t(H) \asymp |t(H) - t(H')| \asymp \ell$. Choosing $\ell = d \wedge t$ yields the lower bound $R \gtrsim \frac{td}{p} \wedge t^2$. ■

Proof [Proof of Theorem 4] To make \hat{e} unbiased, in view of (12), we set

$$1 = \mathbb{E}[\mathcal{K}_A] = pq(f(d_u) + f(d_v)) + p^2 g(d_u, d_v).$$

This determines

$$g(d_u, d_v) = \frac{1 - pq(f(d_u) + f(d_v))}{p^2}.$$

An easy calculation shows that

$$\text{Var}[\mathcal{K}_A] = \frac{(1 - pq(f(d_u) + f(d_v)))^2}{p^2} + pq(f^2(d_u) + f^2(d_v)) - 1$$

and if $A = \{u, w\}$ and $A' = \{w, v\}$ in G , then

$$\text{Cov}[\mathcal{K}_A, \mathcal{K}_{A'}] = \frac{q}{p}(1 - pf(d_u))(1 - pf(d_v)).$$

Otherwise, $\text{Cov}[\mathcal{K}_A, \mathcal{K}_{A'}] = 0$ if A and A' do not intersect. Thus,

$$\begin{aligned} \text{Var}[\widehat{\mathbf{e}}] &= \frac{q}{p} \sum_{u \neq v} d_{uv}(1 - pf(d_u))(1 - pf(d_v)) \\ &\quad + \sum_{\{u, v\} \in E(G)} \left[\frac{(1 - pq(f(d_u) + f(d_v)))^2}{p^2} + pq(f^2(d_u) + f^2(d_v)) - 1 \right], \end{aligned} \quad (51)$$

where d_{uv} denotes the cardinality of $N_G(u) \cap N_G(v)$. To gain a better idea for how to choose f , we first suppose that $f \equiv \alpha$. Thus, (51) reduces to the mean square error of (10) or

$$\text{Var}[\widehat{\mathbf{e}}] = \frac{2q}{p} \mathfrak{n}(P_3, G)(1 - p\alpha)^2 + \mathfrak{e}(G) \frac{q}{p^2} (1 + p(1 - 2\alpha((p - 2)p\alpha + 2)))$$

Next, let us minimize the above expression over all α . Doing so with

$$\alpha' = \left(\frac{1}{p} \right) \frac{p\mathfrak{e}(G) + p\mathfrak{n}(P_3, G)}{p\mathfrak{n}(P_3, G) + \mathfrak{e}(G)(2 - p)} + \left(\frac{1}{2p} \right) \frac{2q\mathfrak{e}(G)}{p\mathfrak{n}(P_3, G) + \mathfrak{e}(G)(2 - p)}.$$

yields

$$\text{Var}[\widehat{\mathbf{e}}] = \frac{q^2}{p} \frac{\mathfrak{e}(G)(\mathfrak{e}(G) + \mathfrak{n}(P_3, G))}{(2 - p)\mathfrak{e}(G) + p\mathfrak{n}(P_3, G)}. \quad (52)$$

Note that α' is a convex combination of $\frac{1}{p}$ and $\frac{1}{2p}$. These are the values that yield the risk bound for the non-adaptive estimator (10) in Theorem 2, viz.,

$$\alpha = \left(\frac{1}{p} \right) \mathbb{1} \left\{ d > \frac{1}{p} \right\} + \left(\frac{1}{2p} \right) \mathbb{1} \left\{ d \leq \frac{1}{p} \right\}.$$

Of course, this choice of α' is not feasible since it depends on the unknown quantities $\mathfrak{e}(G)$ and $\mathfrak{n}(P_3, G)$. However, noting that $\mathfrak{e}(G) = \sum_u d_u/2$ and $\mathfrak{n}(P_3, G) = \sum_u \binom{d_u}{2}$ inspires us to define

$$\begin{aligned} f(d_u) &= \left(\frac{1}{p} \right) \frac{p \binom{d_u}{2} + p \binom{d_u}{2}}{p \binom{d_u}{2} + \binom{d_u}{2}(2 - p)} + \left(\frac{1}{2p} \right) \frac{2q \binom{d_u}{2}}{p \binom{d_u}{2} + \binom{d_u}{2}(2 - p)} \\ &= \left(\frac{1}{2p} \right) \frac{2pd_u}{p(d_u - 1) + (2 - p)} + \left(\frac{1}{2p} \right) \frac{2q}{p(d_u - 1) + (2 - p)} \\ &= \frac{pd_u + q}{p(pd_u + 2q)}. \end{aligned}$$

With this choice of f , we will verify that the variance and covariance terms in (51) also yield the rate (17). Note that

$$\begin{aligned} \frac{q}{p} \sum_{u \neq v} d_{uv} \left[\frac{q}{pd_u + 2q} \right] \left[\frac{q}{pd_v + 2q} \right] &\leq \frac{q}{p} \sum_{u \neq v} \frac{d_{uv}}{pd_u + 2q} \leq \frac{dq}{p} \sum_u \frac{d_u}{pd_u + 2q} \\ &\leq \frac{Ndq}{p} \frac{e(G)}{pe(G) + qN} \leq \frac{Nd}{p^2} \wedge \frac{e(G)d}{p}, \end{aligned}$$

where the second last inequality follows from the concavity of $x \mapsto \frac{x}{px+2q}$ for $x \geq 0$. The variance term has the bound

$$\sum_{\{u,v\} \in E(G)} \left[\frac{(1 - pq(f(d_u) + f(d_v)))^2}{p^2} + pq(f^2(d_u) + f^2(d_v)) - 1 \right] \lesssim e(G) \left(\left(\frac{1}{p^2} \wedge d^2 \right) \vee \frac{1}{p} \right),$$

which follows from $\frac{(1-pq(f(d_u)+f(d_v)))^2}{p^2} \lesssim \frac{1}{p^2} \wedge (d_u^2 + d_v^2)$ and $pq(f^2(d_u) + f^2(d_v)) \lesssim \frac{1}{p}$. \blacksquare

Appendix J. Neighborhood sampling without colors

In this appendix we demonstrate the usefulness of the color information (namely, which vertices are sampled) in neighborhood sampling by showing that without observing the colors, the performance guarantees in Theorem 2 are no longer unattainable in certain regimes.

Theorem 28 *Let \mathcal{F} denote the collection of all forests. Consider the neighborhood sampling model without observing the colors $\{b_v : v \in V\}$. Then*

$$\inf_{\hat{e}} \sup_{\substack{G \in \mathcal{F}: \\ d(G) \leq d \\ e(G) \leq m}} \mathbb{E}_G |\hat{e} - e(G)|^2 \gtrsim mp(d \wedge m). \quad (53)$$

Proof

Let $M = m/k$, where $k = d \wedge m$ and set $\mathcal{F}_0 = \{G_{\underline{\theta}} : G_{\underline{\theta}} = S_{\theta_1} + \dots + S_{\theta_M}, \underline{\theta} = (\theta_1, \dots, \theta_M) \in [k]^M\}$. Note that for each $\underline{\theta} \in [k]^M$, $e(G_{\underline{\theta}}) = \|\underline{\theta}\|_1$. Thus, if $\underline{X} = (X_1, \dots, X_M)$, where $\{X_i\}$ are independent and $X_i \sim p\delta_{\theta_i} + q\text{Bin}(\theta_i, p)$ for $i \in [M]$, then

$$\inf_{\hat{e}} \sup_{\substack{G \in \mathcal{F}: \\ d(G) \leq d \\ e(G) \leq m}} \mathbb{E}_G |\hat{e} - e(G)|^2 \geq \inf_g \sup_{\underline{\theta} \in [d]^M} \mathbb{E}_{\underline{\theta}} \|\underline{\theta}\|_1 - g(\underline{X})^2.$$

By the minimax theorem,

$$\begin{aligned} \inf_g \sup_{\underline{\theta} \in [k]^M} \mathbb{E}_{\underline{\theta}} \|\underline{\theta}\|_1 - g(\underline{X})^2 &= \sup_{\underline{\theta} \in \pi} \inf_g \mathbb{E}_{\underline{\theta}} \|\underline{\theta}\|_1 - g(\underline{X})^2 = \sup_{\underline{\theta} \in \pi} \mathbb{E}_{\underline{X}} \mathbb{E}_{\underline{\theta}|\underline{X}} \|\underline{\theta}\|_1 - \mathbb{E}_{\underline{\theta}|\underline{X}} \|\underline{\theta}\|_1^2 \\ &\geq \sup_{\underline{\theta} \in \pi^{\otimes M}} \mathbb{E}_{\underline{X}} \mathbb{E}_{\underline{\theta}|\underline{X}} \|\underline{\theta}\|_1 - \mathbb{E}_{\underline{\theta}|\underline{X}} \|\underline{\theta}\|_1^2 = M \sup_{\theta \in \pi} \mathbb{E}_X \mathbb{E}_{\theta|X} |\theta - \mathbb{E}_{\theta|X} \theta|^2 \\ &= M \inf_g \sup_{\theta \in [d \wedge m]} \mathbb{E}_{\theta} |\theta - g(X)|^2 \asymp m \left(pk \vee \left(\frac{1}{p} \wedge k \right) \right) \\ &\gtrsim mp(d \wedge m), \end{aligned}$$

where $X \sim \delta_\theta + q\text{Bin}(\theta, p)$ and the second to last line follows from Lemma 30 below. \blacksquare

Remark 29 Note that when $p > (1/d)^{1/3}$ and $m \geq d$, the minimax lower bound (53) is strictly greater than the minimax risk in Theorem 14, thus confirming the intuition that the knowledge of which vertices are sampled provide useful information. On the other hand, the Horvitz-Thompson estimator (7) can be implemented without the color information and achieve the error bound $O(\frac{md}{p})$ in (9). Comparing with Theorem 1, we conclude that neighborhood sampling is at least as informative as subgraph sampling, even if the colors are not observed. This is intuitive because neighborhood sampling reveals more edges from the parent graph.

Lemma 30 Given $\theta \in [k]$, let X be distributed according to $p\delta_\theta + q\text{Bin}(\theta, p)$. Assume that $p \leq 1/2$. Then

$$\inf_g \sup_{\theta \in [k]} \mathbb{E}_\theta[|\theta - g(X)|^2] \asymp pk^2 \vee \left(\frac{k}{p} \wedge k^2\right). \quad (54)$$

Moreover, the minimax rate is achieved by the estimator $\hat{g}(X) = k \wedge \frac{X}{p}$.

Proof Denote the minimax risk by R . Let id denote the identity map. Given any estimator g , without loss of generality, we assume $g : \{0, \dots, k\} \rightarrow [0, k]$. Since $\mathbb{E}_\theta[(\theta - g(X))^2] = p(\theta - g(\theta))^2 + q\mathbb{E}_{X \sim \text{Bin}(\theta, p)}[(\theta - g(X))^2]$, we have

$$\sup_{\theta \in [k]} \mathbb{E}_\theta[|\theta - g(X)|^2] \geq p\|\text{id} - g\|_\infty^2. \quad (55)$$

Also, $(\theta - g(X))^2 \geq -(X - g(X))^2 + (\theta - X)^2/2$, and hence

$$\mathbb{E}_{X \sim \text{Bin}(\theta, p)}[(\theta - g(X))^2] \geq -\|\text{id} - g\|_\infty^2 + \frac{1}{2}(q^2\theta^2 + pq\theta).$$

Therefore

$$\sup_{\theta \in [k]} \mathbb{E}_\theta[|\theta - g(X)|^2] \geq -q\|\text{id} - g\|_\infty^2 + \frac{q}{2}(q^2k^2 + pqk). \quad (56)$$

Combining (55) and (56), we get

$$R \geq \frac{pq}{2}(q^2k^2 + pqk) \asymp pk^2. \quad (57)$$

Next by the minimax theorem,

$$R = \sup_\pi \inf_g \mathbb{E}_\pi[|\theta - g(X)|^2] = \sup_\pi \inf_g \left(\underbrace{p \mathbb{E}_{\theta \in \pi}[|\theta - g(\theta)|^2]}_{\in [0, k^2]} + q \mathbb{E}_{\theta \in \pi, X \sim \text{Bin}(\theta, p)}[|\theta - g(X)|^2] \right).$$

We also know that

$$\sup_\pi \inf_g \mathbb{E}_{\theta \in \pi, X \sim \text{Bin}(\theta, p)}[|\theta - g(X)|^2] = \inf_g \sup_{\theta \in [k]} \mathbb{E}_{\theta \in \pi, X \sim \text{Bin}(\theta, p)}[|\theta - g(X)|^2] \asymp \frac{k}{p} \wedge k^2.$$

Therefore we have

$$\frac{k}{p} \wedge k^2 \lesssim R \lesssim pk^2 + \frac{k}{p} \wedge k^2.$$

Combining with (57) yields the characterization (54). \blacksquare

Appendix K. Lower bounds for other motifs

Theorem 31 (Wedges) *For neighborhood sampling with sampling ratio p ,*

$$\inf_{\widehat{w}} \sup_{\substack{G: d(G) \leq d \\ w(G) \leq w}} \mathbb{E}_G |\widehat{w} - w(G)|^2 \asymp \frac{wd}{p^2} \wedge w^2.$$

Proof For the lower bound, consider two cases:

Case I: $p \leq 1/d$. Let $h = P_5$ and $h' = K_3 + K_2$. For each node in the original graph, we associate ℓ distinct isolated vertices and connect each pair of vertices by an edge if and only if they were connected in the original graph. Call these expanded graphs H and H' . Note that H and H' that have matching degree sequences $(2, 2, 2, 1, 1)$ and hence $\text{TV}(P_{\widehat{H}}, P_{\widehat{H}'}) = O(\ell^2 p^2)$. Furthermore, $s(P_3, H) \asymp s(P_3, H') \asymp |s(P_3, H) - s(P_3, H')| \asymp \ell^3$. If $\ell = c(d \wedge w^{1/3})$, then by Theorem 6 with $M = w/\ell^3$, $\inf_{\widehat{w}} \sup_{G \in \mathcal{G}(w, d)} \mathbb{E}_G |\widehat{w} - w(G)|^2 \gtrsim \frac{w\ell}{p^2} \wedge w^2 \asymp \frac{wd}{p^2} \wedge w^2$.

Case II: $p > 1/d$. We use Lemma 23 with $G = K_\ell$, $H_1 = K_{1/p} + K_{1/p}$, and $H_2 = K_{2/p}$. This gives us two graphs H and H' with $s(P_3, H) = |s(P_3, H) - s(P_3, H')| \asymp \ell/p^2$. By Theorem 6 with $M = w/(\ell/p^2)$, $\inf_{\widehat{w}} \sup_{G \in \mathcal{G}(w, d)} \mathbb{E}_G |\widehat{w} - w(G)|^2 \gtrsim \frac{w\ell}{p^2} \wedge w^2$. Let $\ell = cd$ if $\frac{d}{p^2} \leq w$ and $\ell = cp^2 w$ if $\frac{d}{p^2} > w$, for some small constant c . In either case, we find that $w(H) \leq w$, $w(H') \leq w$, and $\inf_{\widehat{w}} \sup_{G \in \mathcal{G}(w, d)} \mathbb{E}_G |\widehat{w} - w(G)|^2 \asymp \frac{wd}{p^2} \wedge w^2$. \blacksquare

Lower bound for motifs of size four It remains to show that holds the result in Theorem 12 that holds for K_4 , namely,

$$\inf_{\widehat{s}} \sup_{\substack{G: d(G) \leq d \\ s(h, G) \leq s}} \mathbb{E}_G |\widehat{s} - s(K_w, G)|^2 = \Theta \left(\frac{sd}{p^3} \wedge \frac{sd^2}{p^2} \wedge s^2 \right) \quad (58)$$

continues to hold for $h = \square, \triangle, \square$ and \circ . For the case of $p < 1/d$, the construction for K_4 in (29) works simultaneously for all motifs, because each motif is contained in one of H and H' and not the other. Next we consider the case of $p > 1/d$. The construction is ad hoc and similar to those in Theorem 1 and Theorem 3.

- For $h = \circ$, we use the clique construction: label the root as v_1 and the leaves as v_2, v_3, v_4 . Define the graph H as follows: Expand v_1 into a clique S_1 of size ℓ , and for $i = 2, 3, 4$, expand each v_i into a clique S_i of size $1/p$. Connect each pair of vertices $u_i \in S_i$ and $u_j \in S_j$ for $i \neq j$ if and only if v_i and v_j are connected in the motif h . This defines a graph H on $\ell + 3/p$ vertices. Repeat the same construction with

h replaced by $\triangleleft\circ$, where the degree-one vertex is v_1 . Note that if we remove the edges between the set of vertices $T \triangleq S_2 \cup S_3 \cup S_4$, for H and H' the resulting graph is isomorphic. Thus by Lemma Lemma 23, we have $\text{TV}(P_{\tilde{H}}, P_{\tilde{H}'}) \leq 1 - (1-p)^{3/p} \leq 0.9$ if $p \leq 1/2$. Furthermore, note that $s(\circ\circ, H') = 0$ and $s(\circ\circ, H) = \ell/p^3$. Finally, taking $\ell = c(d \wedge \frac{s}{\ell/p^3})$ for some small constant c and invoking Theorem 6, we obtain the desired lower bound $\frac{sd}{p^3} \wedge s^2$ in (58).

- For $h = \triangleleft\circ$, use the same construction as above with H and H' swapped.
- For $h = \triangleleft\circ$, we repeat the clique construction of H with v_1 being any of the degree-three vertices in h , and of H' with $h' = \triangleleft\circ$; in other words, we simply have $H' = K_{\ell+3/p}$.
- For $h = \square\circ$, we repeat the clique construction of H with v_1 being any vertex in h , and of H' with $h' = \circ\circ$, with v_1 being the degree-two vertices.