Journal of Machine Learning Research 18 (2017) 1-33 Submitted 10/16; Revised 5/17; Published 7/17

On the Propagation of Low-Rate Measurement Error to Subgraph Counts in Large Networks

Prakash Balachandran PRAKASHB@BU.EDU
Eric D. Kolaczyk KOLACZYK@BU.EDU
Department of Mathematics & Statistics
Boston University
Boston, MA 02215 USA

Weston D. Viles WESTON.D.VILES@DARTMOUTH.EDU
Department of Biomedical Data Science
Dartmouth College
Hanover, NH 03755 USA

Editor: Network Analysis Special Issue

Abstract

Our work in this paper is inspired by a statistical observation that is both elementary and broadly relevant to network analysis in practice—that the uncertainty in approximating some true graph $G = (V, E)$ by some estimated graph $\hat{G} = (V, \hat{E})$ manifests as errors in our knowledge of the presence/absence of edges between vertex pairs, which must necessarily propagate to any estimates of network summaries $\eta(G)$ we seek. Motivated by the common practice of using plug-in estimates $\eta(\hat{G})$ as proxies for $\eta(G)$, our focus is on the problem of characterizing the distribution of the discrepancy $D = \eta(\hat{G}) - \eta(G)$, in the case where $\eta(\cdot)$ is a subgraph count. Specifically, we study the fundamental case where the statistic of interest is $|E|$, the number of edges in G. Our primary contribution in this paper is to show that in the empirically relevant setting of large graphs with low-rate measurement errors, the distribution of $D_E = |\hat{E}| - |E|$ is well-characterized by a Skellam distribution, when the errors are independent or weakly dependent. Under an assumption of independent errors, we are able to further show conditions under which this characterization is strictly better than that of an appropriate normal distribution. These results derive from our formulation of a general result, quantifying the accuracy with which the difference of two sums of dependent Bernoulli random variables may be approximated by the difference of two independent Poisson random variables, i.e., by a Skellam distribution. This general result is developed through the use of Stein's method, and may be of some general interest. We finish with a discussion of possible extension of our work to subgraph counts $\eta(G)$ of higher order.

Keywords: Limit distribution, network analysis, Skellam distribution, Stein's method.

1. Introduction

The analysis of network data is widespread across the scientific disciplines. Technological and infrastructure, social, biological, and information networks are a few of the major network classes in which such analyses have been employed. However, despite the already substantial body of work in network analysis generally (e.g., see Jackson, 2010; Kolaczyk, 2009; Newman, 2010), with contributions from a variety of different fields of study, much work still remains to more fully develop the theory and methods of statistical analysis of network data, particularly for certain classes

of problems of a fairly fundamental nature. Here in this paper we pose and address a version of one such fundamental problem, that regarding the propagation of error through the process of network construction and summary.

In applied network analysis, a common *modus operandi* is to (i) gather basic measurements relevant to the interactions among elements in a system of interest, (ii) construct a graph-based representation of that system, with nodes serving as elements and links indicating interactions between pairs of elements, and (iii) summarize the structure of the resulting graph using a variety of numerical and visual tools. See Kolaczyk (2009, Chs 3 & 4) for background and several case studies illustrating this process. Key here is the point that the process of network analysis usually rests upon some collection of measurements of a more basic nature. For example, online social networks (e.g., Facebook) are based on the extraction and merging of lists of 'friends' from millions of individual accounts (Hansen et al., 2010, Ch 11.8). Similarly, biological networks (e.g., of gene regulatory relationships) are often based on notions of association (e.g., correlation, partial correlation, etc.) among experimental measurements of gene activity levels (Fogelberg and Palade, 2009). Finally, maps of the logical Internet traditionally have been synthesized from the results of surveys in which paths along which information flows are learned through a large set of packet probes (e.g., Cheswick et al., 2000).

Importantly, while it is widely recognized that there is measurement error associated with these and other common types of network constructions, most applied network analyses in practice effectively proceed as if there were in fact no error. There are at least two possible reasons for this current state of affairs: (1) there is comparatively little in the way of formal probabilistic analyses characterizing the propagation of such error and of statistical methods accounting for such propagation, and (2) in many settings (arguably due at least in part to (1)), much attention is given at the stages of measurement and network construction to trying to keep the rate of error 'low' in declaring the presence and absence of links between nodes.

Here we offer a formal treatment of the problem of propagation of error. We provide a general framework within which it is possible to characterize the manner in which errors made in assigning links between nodes accumulate in the reporting of certain functions of the network as a whole. Our treatment is probabilistic, wherein our goal is to understand the nature of the distribution induced on the graph functions by that of the errors in the graph construction. We anticipate that this contribution will provide part of a critical foundation upon which, in turn, it should be possible to develop a statistical treatment, e.g., the construction of appropriate confidence intervals for graph-based quantities of interest.

More formally, we consider a setting wherein an underlying (undirected) network-graph G possesses a network characteristic $\eta(G)$ of interest. While there are many types of functions $\eta(\cdot)$ used in practice to characterize networks (e.g., centralities, path-based metrics, output from methods of community detection, etc.) we restrict our attention here to the canonical problem of subgraph counting. That is, we are interested in the class of functions η of the form

$$\eta_H(G) = \frac{1}{|Iso(H)|} \sum_{H' \subseteq K_{n_v}, H' \cong H} 1\{H' \subseteq G\} \ , \tag{1}$$

where $n_v = |V(G)|$ is the number of vertices in G, K_{n_v} is the complete graph on n_v vertices, H is a graph of interest (i.e., copies of which we desire to count), and $H \subseteq G$ indicates that H is a subgraph of G (i.e., $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$). The value $|Iso(H)|$ is a normalization

factor for the number of isomorphisms of H. We will concentrate primarily on the fundamental case where $\eta(G) = |E|$, i.e., the number of edges in G.

If $\hat{G}$ is a network-graph resulting from an attempt to construct G from some collection of basic measurements, then the common practice of reporting the analogous characteristics of $\hat{G}$ is equivalent to the use of $\eta\left(\hat{G}\right)$, i.e., effectively a plug in estimator. Let the discrepancy between these two quantities be defined as $D = \eta\left(\hat{G}\right) - \eta\left(G\right)$, which in the case of counting edges reduces to $D_E = |\hat{E}| - |E|$. Our goal is to make precise probabilistic statements about the behavior of D under certain conditions.

Importantly, in the case where η is defined as a subgraph count, as in (1), D may be expressed as the difference of (i) the number of times the subgraph H arises somewhere in $\hat{G}$ but does *not* in fact exist in the same manner in G, and (ii) vice versa. Hence, D may be understood in this context to be the difference in total number of Type I and Type II errors, respectively. Intuitively, in the cases where a sufficiently low rate of error occurs on a large graph G, each of these two sums should have a Poisson-like behavior. This observation suggests that the propagation of low-rate measurement error to subgraph counts should behave, under appropriate conditions, like the difference of two Poisson random variables, i.e., a so-called Skellam distribution (Skellam, 1946).

Our contribution in this paper is to provide an initial set of results on the accuracy with which the Skellam distribution may be used in approximating the distribution of D, under the setting where the graph G is large and the rate of error is low. We consider the cases of both sparse and dense networks. The primary technical device used here is Stein's method (e.g, Barbour and Chen, 2005). Specifically, we present a Stein operator for the Skellam probability distribution and, in a manner consistent with the Stein methodology, we derive several bounds on the discrepancy between the distribution of the difference of two arbitrary sums of binary random variables to an appropriately parameterized Skellam distribution. The latter in turn is then used to establish in particular the rate of weak convergence of D_E to an appropriate Skellam random variable, under either independent or weakly dependent measurement errors.

As remarked above, there appears to be little in the way of a formal and general treatment of the error propagation problem we consider here. However, there are, of course, several areas in which the probabilistic or statistical treatment of uncertainty enters prominently in network analysis. The closest area to what we present here is the extensive literature on distributions of subgraph counts in random graphs. See Janson et al. (2011), for example, for comprehensive coverage. Importantly, there the graph G is assumed to emerge according to a (classical) random graph and uncertainty typically is large enough that normal limit theorems are the norm (although Poisson limit theorems also have been established). In contrast, in our setting we assume that G is a fixed, true underlying graph, and then study the implications of observing a 'noisy' version $\hat{G}$ of that graph, under various assumptions on the nature of the noise, which involves two specific types of error (i.e., Type I and II errors), the contributions of which are informed in part by the topology of G itself. An area related to this work in random graphs is the work in statistical models for networks, such as exponential random graph models (ERGMs). See Lusher et al. (2012) for a recent treatment. Here, while these models are inherently statistical in nature, the randomness due to generation of the graph G and due to observation of G—resulting in what we call $\hat{G}$—usually are combined into a single realization from the underlying distribution. And while subgraph counts do play a key role in traditional ERGMs, they typically enter as covariates in these (auto)regressive models. In a somewhat different direction, uncertainty in network construction due to sampling has also been

studied in some depth. See, for example, Kolaczyk (2009, Ch 5) or Ahmed et al. (2014) for surveys of this area. However, in this setting, the uncertainty arises only from sampling—the subset of vertices and edges obtained through sampling are typically assumed to be observed without error. Finally, we note that there just recently have started to emerge in the statistics literature formal treatments of the same type of graph observation error model that we propose here. There the emphasis is on producing estimators of network model parameters and/or classifiers (e.g., Priebe et al., 2015), for example, rather than on the type of basic network summary statistics that are the focus here.

The organization of this paper is as follows. In Section 2 we provide necessary background. In Section 3 we then provide a general set of results useful for our general problem. Specifically, we establish a bound for the Kolmogorov-Smirnov distance between the distribution of the difference of two arbitrary sums of binary random variables from a certain Skellam. This work is based on the application of Stein's method to the Skellam distribution, a first of its kind to the best of our knowledge, and the results therefore are of some independent interest as well. In Section 4 we then illustrate the way in which these general results may be used to understand the propagation of error in networks for counting edges. In doing so, several other general results are established. Some implications of these results on the problem of counting subgraphs of higher order are noted in Section 5, along with other related discussion. Proofs of our key results may be found in the appendices.

2. Background

In this section, we provide essential notation and background for our forthcoming general results.

2.1 Notation and Assumptions

By $G = (V, E)$ we will mean an undirected graph, with vertex set V of cardinality $|V|$ and edge set E of cardinality $|E|$. Vertex pairs in the complement of E, i.e., in E^c, will be referred to as non-edges. Much of the results that follow will be stated as a function of the number of vertices which, for notational convenience, we denote $n_v = |V|$. Let $\mu = \mu(G) = 2|E|/n_v$ correspond to the average degree of a vertex in G. We assume the vertex set V is known but that the edge set E is unknown. However, we assume there is information by which to construct an approximation to E or, more formally, by which to infer E, as a set $\hat{E}$, yielding an inferred graph $\hat{G} = (V, \hat{E})$.

While there are many ways in practice by which the set $\hat{E}$ is obtained, one principled way of viewing the task is as one of performing $\binom{n_v}{2}$ hypothesis tests, using the data underlying the graph construction process as input, one for each of the vertex pairs $\{i, j\}$ in the graph G. In some contexts, $\hat{G}$ is literally obtained through hypothesis testing procedures; for instance, in constructing some gene regulatory networks from microarray expression data. See Kolaczyk (2009, Ch 7), for example. Formally, in such cases we can think of $\hat{G}$ as resulting from a collection of testing problems

$$H_0 : \{i, j\} \notin E \text{ versus } H_1 : \{i, j\} \in E \ ,$$

for $\{i, j\} \in V^{(2)}$, where

$$V^{(2)} = \{\{i, j\} : i, j \in V; i < j\} \ .$$

These tests amount to a collection of $\binom{n_v}{2}$ binary random variables $\left\{Y_{ij} : \{i,j\} \in V^{(2)}\right\}$, where

$$Y_{ij} = \begin{cases} 1 & \text{if } H_0 \text{ is rejected} \\ 0 & \text{if } H_0 \text{ is not rejected.} \end{cases}$$

Note that the random variables Y_{ij} need not be independent and, in fact, in many contexts will most likely be dependent. Gene regulatory networks inferred by correlating gene expression values at each vertex i with that of all other vertices $j \in V \setminus \{i\}$ and maps of the logical Internet obtained through merging paths learned by sending traffic probes between many sources and destinations are just two examples where dependency can be expected.

Whether obtained informally or formally, we can define the inferred edge set $\hat{E}$ as

$$\hat{E} = \left\{\{i,j\} \in V^{(2)} : Y_{ij} = 1\right\} \ .$$

It is useful to think of the collection of random variables $\left\{Y_{ij} : \{i,j\} \in V^{(2)}\right\}$ as being associated with two types of errors. That is, we express the marginal distributions of the Y_{ij} in the form

$$Y_{ij} \sim \begin{cases} \text{Bernoulli}\left(\alpha_{ij}\right), \text{if } \{i,j\} \in E^c, \\ \text{Bernoulli}\left(1-\beta_{ij}\right), \text{if } \{i,j\} \in E, \end{cases} \tag{2}$$

where $E^c = V^{(2)} \setminus E$. Again pursuing the example of network construction based on hypothesis testing, α_{ij} can be interpreted as the probability of Type-I error for the test on vertex pair $\{i,j\} \in E^c$, while β_{ij} is interpreted as the probability of Type-II error for the test on vertex pair $\{i,j\} \in E$. For example, α_{ij} might be the probability of incorrectly inferring a regulatory relationship between two genes i and j, based on thresholding the empirical correlation between their expression values, as quantified through deep sequencing measurements over a series of experiments. Similarly, β_{ij} might be the probability of missing a logical link in the Internet topology between two routers i and j, due to artifacts in measurement technologies (e.g., traceroute).

Our interest in this paper is in characterizing the manner in which the uncertainty in the Y_{ij} propagates to subgraph counts on $\hat{G}$. More specifically, we seek to characterize the distribution of the difference

$$D = \frac{1}{|Iso(H)|} \sum_{H' \subseteq K_{n_v}, H' \cong H} \left[1\{H' \subseteq \hat{G}\} - 1\{H' \subseteq G\}\right] \ , \tag{3}$$

for a given choice of subgraph H. Naturally, this distribution will depend in no small part on context. Here we focus on a general formulation of the problem in which we make the following three assumptions.

(A1) *Large Graphs:* $n_v \to \infty$.

(A2) *Edge Unbiasedness:* $\sum_{\{i,j\} \in E^c} \alpha_{ij} = \sum_{\{i,j\} \in E} \beta_{ij} \ (\equiv \lambda)$.

(A3) *Low Error Rate:* $\lambda = \Theta\left(\mu\right)$.

Assumption (A1) reflects both the fact that the study of large graphs is a hallmark of modern applied work in complex networks and, accordingly, our desire to make statements that are asymptotic in n_v.

Our use of assumption (A2) reflects the understanding that a 'good' approximation $\hat{G}$ to the graph G should at the very least have roughly the right number of edges. The difference of the two sums defined in (A2) is in fact the expectation of the statistic D in (3) for the case where the subgraph being counted is just a single edge, i.e., it is the expected discrepancy between the number of observed edges $|\hat{E}|$ and the actual number of edges $|E|$. So (A2) states that this particular choice of D has expectation zero. Alternatively, (A2) may be interpreted as saying that the total numbers of Type I and Type II errors should be equal to a common value λ.

Finally, we capture the notion of a 'low' rate of error in $\hat{G}$ through assumption (A3). Specifically, we assume that the number of Type I or Type II errors in edge status that we expect throughout the network is roughly on par with the average number of edges incident to any given vertex in the network. This condition can be re-expressed in a useful manner with respect to n_v if, as is common in the literature, we distinguish between sparse and dense graphs. By the term *sparse* we will mean a graph for which $|E| = \Theta(n_v \log n_v)$, and by *dense*, $|E| = \Theta\left(n_v^2\right)$. Hence, assumption (A3) reduces to $\lambda = \Theta(\log n_v)$ in the case of sparse graphs, and to $\lambda = \Theta(n_v)$, in the case of dense graphs.

In addition, for convenience, we add to the core assumptions (A1)-(A3) a fourth assumption, upon which we will call periodically throughout the paper when desiring to simplify some of our expressions.

(A4) *Homogeneity:* $\alpha_{ij} \equiv \alpha$ and $\beta_{ij} \equiv \beta$, for $\alpha, \beta \in (0, 1)$.

In other words, we assume that the probability of making a Type I or II error (as the case may be) does not depend upon the specific (non)edge in question.

Lastly, for completeness, we recall the definition of the Skellam distribution. A random variable W defined on the integers is said to have a Skellam distribution, with parameters $\lambda_1, \lambda_2 > 0$, i.e., $W \sim \text{Skellam}(\lambda_1, \lambda_2)$, if

$$\mathbb{P}(W = k) = e^{-(\lambda_1+\lambda_2)} \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k\left(2\sqrt{\lambda_1\lambda_2}\right) \text{ for } k \in \mathbb{Z}, \tag{4}$$

where $I_k\left(2\sqrt{\lambda_1\lambda_2}\right)$ is the modified Bessel function of the first kind with index k and argument $2\sqrt{\lambda_1\lambda_2}$. The Skellam distribution may be constructed by defining W through the difference of two independent Poisson random variables, with means λ_1 and λ_2, respectively. The mean and variance of this distribution are given by $\mathbb{E}[W] = \lambda_1 - \lambda_2$ and $\text{Var}(W) = \lambda_1 + \lambda_2$. The distribution of W is in general asymmetric, with symmetry holding if and only if $\lambda_1 = \lambda_2$.

The main results we provide in this paper are in the form of bounds on the extent to which the distribution of random variables like the discrepancy D in (3) may be well-approximated by an appropriate Skellam distribution. For this purpose, we adopt the Kolmogorov-Smirnov distance to quantify the distance between distributions of two random variables, say, X_1 and X_2, i.e.,

$$ds_{KS}(X_1, X_2) \equiv \sup_{x \in \mathbb{R}} |\mathbb{P}(X_1 \leq x) - \mathbb{P}(X_2 \leq x)|.$$

2.2 Counting Edges

Generic subgraph counts, and the corresponding noise in obtaining them, can be quite varied in real applications. Accordingly, most of our specific results pertain to the fundamental case of counting

edges. That is, where the choice of subgraph H is simply a single edge, and therefore the function $\eta(G)$ in (1) is just the total number of edges in G, i.e., $\eta(G) = |E|$. We will consider two scenarios for this case, wherein the random variables Y_{ij} are independent or weakly dependent.

In the case where the edge noise is independent, the discrepancy

$$\begin{aligned} D_E &= |\hat{E}| - |E| \\ &= \sum_{\{i,j\} \in E^c} Y_{ij} - \sum_{\{i,j\} \in E} (1 - Y_{ij}) \end{aligned}$$

has expectation $\mathbb{E}[D_E] = \alpha|E^c| - \beta|E| = \lambda - \lambda = 0$ and variance $\sigma^2 = \alpha(1-\alpha)|E^c| + \beta(1-\beta)|E|$, and its behavior can be established using existing methods from the literature (i.e., essentially, Chen-Stein methods). However, we include it as an important base case, comparing results obtainable by our methods to those obtainable by more traditional techniques, in Section 4.1.

Alternatively, suppose that the variables Y_{ij} are *dependent*. The random variable D_E again has expectation zero, although its variance—and hence its asymptotic behavior—will differ from the independent case, in a manner dictated by the nature of the underlying dependency in the noise. It often can be expected in practice that the error associated with construction of the empirical graph $\hat{G}$ will involve dependency across (non)edges. For example, relations in gene regulatory networks are often declared based on sufficiently strong levels of association between gene-specific measurements (e.g., measures of gene expression). The comparison of the measurements for each gene with those of all of the others necessarily induces potential dependencies among the random variables Y_{ij}. However, a precise characterization of such dependency is typically problem-specific and, more often than not, nontrivial in nature. In Section 4.2 we will assume general dependency conditions in the spirit of traditional monotone coupling arguments, which will allow for further analysis and interpretation.

3. General Results on Approximation by Skellam

Recall the general form of our statistic of interest D in (3), as the difference of two sums of binary random variables. Under appropriate conditions it seems reasonable to expect that the distribution of D be well-approximated by a Skellam distribution. And for the simplest case, in which we are counting edges under independent noise, it is possible to show that this is in fact the case, through manipulation of existing results for approximating sums by Poisson distributions. Without independence, however, it is necessary to approach the problem directly, by explicitly using the Skellam distribution. In this section, we therefore provide the results of such an approach. This is a completely general treatment—devoid of the motivating context of counting subgraphs—and therefore also likely of some independent interest. In Section 4 we return to the problem of counting subgraphs under low-rate error and illustrate the use of the results presented here in this section through application to the case of counting edges.

Our approach in this section is through Stein's method. This choice is reminiscent, naturally, of the Chen-Stein treatment for Poisson approximations. However, the task is technically more involved at several points, as it requires handling a Stein function that is defined through a second-order difference, rather than the first-order difference encountered in the Poisson problem. Moreover, the kernel of the Skellam distribution includes a modified Bessel function of the first kind, which emerges in ways necessitating a somewhat delicate treatment.

3.1 A Stein Bound for the Skellam Distribution

Let U be a random variable defined as

$$U = \sum_{k=1}^{n} L_k - \sum_{k=1}^{m} M_k \ , \tag{5}$$

where $\{\{L_k\}_{k=1}^n, \{M_k\}_{k=1}^m\}$ is a collection of two sets of indicator random variables with $\mathbb{E}[L_k] = p_k$ for $k = 1, \ldots, n$ and $\mathbb{E}[M_k] = q_k$ for $k = 1, \ldots, m$. In the case of our subgraph counting problem, $U = D$, where D is defined in (3), although for the remainder of this current section U is defined generally.

Recall the definition of a Skellam random variable W in (4). We desire a bound on

$$d_{KS}(U, W) := \sup_{x \in \mathbb{R}} |\mathbb{P}(U \leq x) - \mathbb{P}(W \leq x)| \ , \tag{6}$$

quantifying how close the distribution of U is to that of W. In pursuing the standard paradigm for Stein's method, we first determine an operator $\mathcal{A}[f(k)]$ such that

$$\mathbb{E}\mathcal{A}[f(W)] = 0 \quad \text{if and only if} \quad W \sim \text{Skellam}(\lambda_1, \lambda_2)$$

for any bounded function $f : \mathbb{Z} \mapsto \mathbb{R}$. This operator need not be unique, but the theory only requires one. This is accomplished through the following result, the proof of which uses several properties of the modified Bessel function of the first kind, as detailed in the appendix, in Section A.1.

Theorem 1 *A Stein operator $\mathcal{A}$ for the Skellam (λ_1, λ_2) distribution is*

$$\mathcal{A}[f(k)] = \lambda_1 f(k+1) - k f(k) - \lambda_2 f(k-1) .$$

With this operator in hand, and again following the usual paradigm under Stein's method, we set

$$\mathcal{A}[f(k)] = g(k)$$

for a class of test functions $g(k)$, and allow that to implicitly define the function f. The choice of the test functions g is guided by the choice of the metric used to measure the distance between U and W. Since the metric we choose to measure the distance between U and W is given by $d_{KS}(U, W)$ in (6), we choose the test function $g := g_x$ given by

$$g_x(k) = 1\{k \leq x\} - \mathbb{P}(W \leq x) \tag{7}$$

for any $x \in \mathbb{R}$.

At this point it is common to exhibit a solution f defined by our choice of g. Instead, we forestall that step until later in this section, choosing rather to state a general result that will allow us to more quickly gain insight into the nature of the bounds we are able to obtain. Our result employs a minor variant of the notion of coupling that is common to the literature on Chen-Stein approximations.

Theorem 2 *Let U be as in (5) and let $\mathcal{L}(U)$ denote the law of U. Let*

$$\mathcal{L}\left(U_k^{(L)} + 1\right) = \mathcal{L}(U | L_k = 1) \quad \text{for} \quad k = 1, \ldots, n$$

and

$$\mathcal{L}\left(U_k^{(M)} - 1\right) = \mathcal{L}\left(U | M_k = 1\right), \quad \text{for} \quad k = 1 \ldots, m$$

be a collection of random variables all defined on a common probability space. If $\lambda_1 = \sum_{k=1}^{n} p_k$ *and* $\lambda_2 = \sum_{k=1}^{m} q_k$*, and* $W \sim Skellam(\lambda_1, \lambda_2)$*, then*

$$d_{KS}(U, W) \leq ||\Delta f|| \left\{ \sum_{k=1}^{n} p_k \mathbb{E}\left| U - U_k^{(L)} \right| + \sum_{k=1}^{m} q_k \mathbb{E}\left| U - U_k^{(M)} \right| \right\} \quad , \tag{8}$$

where

$$||\Delta f|| = \sup_{x \in \mathbb{R}} \sup_{j \in \mathbb{Z}} \left| f_x\left(j+1\right) - f_x\left(j\right) \right|$$

and f_x *is a solution to* $\mathcal{A}\left[f_x(k)\right] = g_x(k)$ *for* $k \in \mathbb{Z}$*.*

The proof of this result relies on elementary considerations of the equation $\mathcal{A}[f_x(k)] = g_x(k)$ and may be found in the appendix in Section A.2. The extent to which it allows one to obtain error estimates of practical interest in a particular setting will depend on the extent to which both the main expression within brackets in (8) and the preceding constant $||\Delta f||$ can be further controlled. While control of the former is problem dependent, control of the latter is not, and may be dealt with separately, as we do next. Afterwards, in Section 4, we illustrate the control of the bracketed expression in (8), in the context of the problem of counting edges introduced in Section 2.2.

3.2 Controlling the Constant Term

Controlling $||\Delta f||$ in (8) first requires understanding the solution $f_x(k)$. We provide a family of closed-form expressions for this solution in the following.

Theorem 3 *Let* g_k *be defined as in equation (7). If* f_x *is a bounded solution to the difference equation*

$$\lambda_1 f_x\left(k+1\right) - k f_x\left(k\right) - \lambda_2 f_x\left(k-1\right) = g_x\left(k\right)$$

for $k \in \mathbb{Z}$*, then* f_x *is given by*

$$f_x\left(m\right) = \begin{cases} (-1)^m \left(\sqrt{\frac{\lambda_2}{\lambda_1}}\right)^m I_m\left[(-1)^c \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^c \frac{1}{I_c} f\left(c\right) \right. \\ \left. + \frac{e^{\lambda_1+\lambda_2}}{\sqrt{\lambda_1\lambda_2}} \sum_{n=c}^{m-1} \frac{(-1)^{n+1}}{I_n I_{n+1}} \mathbb{P}\left(W \leq \min\left\{n, x\right\}\right) \mathbb{P}\left(W > \max\left\{n, x\right\}\right) \right] & \text{if } m > c \\ (-1)^m \left(\sqrt{\frac{\lambda_2}{\lambda_1}}\right)^m I_m\left[(-1)^c \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^c \frac{1}{I_c} f\left(c\right) \right. \\ \left. - \frac{e^{\lambda_1+\lambda_2}}{\sqrt{\lambda_1\lambda_2}} \sum_{n=m}^{c-1} \frac{(-1)^{n+1}}{I_n I_{n+1}} \mathbb{P}\left(W \leq \min\left\{n, x\right\}\right) \mathbb{P}\left(W > \max\left\{n, x\right\}\right) \right] & \text{if } m < c. \end{cases}$$

for any initial condition $(c, f_x\left(c\right))$ *with* $c \in \mathbb{Z}$ *and* $f_x\left(c\right) \in \mathbb{R}$*. Here* I_m *is used as short-hand for* $I_m(2\sqrt{\lambda_1\lambda_2})$*.*

The proof of this theorem is similar to that of solving a second order linear differential equation. An integrating factor is found, integration is performed with a boundary condition at $-\infty$, and then a second integration is performed with the initial condition. Details are provided in the appendix in Section A.3.

Leveraging our insight into f_x to control $||\Delta f||$ means producing a bound on the absolute differences $|\Delta f_x(j)| = |f_x(j+1) - f_x(j)|$ independent of $x \in \mathbb{R}$ and $j \in \mathbb{Z}$. Consider first the special case where λ_1 and λ_2 are equal, for which we are able to offer the following result.

Theorem 4 *Suppose that $\lambda_1 = \lambda_2 \equiv \lambda$, so that $\mathbb{E}[U] = 0$ and W is a Skellam(λ, λ) distribution, symmetric about zero. Assume $\lambda \geq 1$. Then*

$$||\Delta f|| \leq \frac{160}{2\lambda}.$$

The proof of this theorem is highly technical in nature, and relies on a concentration inequality for the Skellam(λ, λ) distribution (Balachandran et al., 2013b) with several other technical arguments. A sketch of the proof may be found in the appendix in Section A.4, while a detailed presentation is available in the online appendix.

Note that the bound in Theorem 4 is essentially the analogue of the classical result for Poisson approximation, in which, for sufficiently large λ, the term $1/\lambda$ is the standard factor. In both cases, therefore, the corresponding term $||\Delta f||$ is bounded by the inverse of the expected total number of counts, where here that is $\mathbb{E}[T_1 + T_2] = 2\lambda$.

The above result is of immediate relevance to the problem of counting edges, given assumption (A2), whether under the assumption that the edge noise is independent or dependent. We will make use of this result in the next section. For applications involving higher-order subgraphs, we can expect to need an extension of Theorem 4 to the general case of $\lambda_1 \neq \lambda_2$. For arbitrary $\lambda_1, \lambda_2 > 0$, we are unable to produce a satisfactory bound. However, supported by preliminary numerical work, we have the following conjecture.

Conjecture 5 *In general, for λ_1 and λ_2 sufficiently close and large,*

$$||\Delta f|| \leq \frac{C}{\lambda_1 + \lambda_2},$$

for some constant $C > 0$ independent of λ_1, λ_2.

4. Application of General Results to Counting Edges

We now illustrate the use of our general results for the problem of characterizing the propagation of low-rate measurement error to subgraph counts in large graphs, for the specific case of counting edges.

4.1 Edge Counts Under Independent Edge Noise

Recall the problem wherein the function of interest (1) counts the number of edges in G, i.e., $\eta(G) = |E|$, and the variables Y_{ij} in (2) are independent. In light of Theorems 2 and 4, we have the following result characterizing the behavior of the discrepancy D in (3), which here is simply $D_E = |\hat{E}| - |E|$.

Theorem 6 *Under assumptions (A1)-(A4) and independence among errors in declaring (non)edge status (i.e., among the Y_{ij}),*

$$d_{KS}\left(D_E \,,\, \textit{Skellam}(\lambda, \lambda) \right) \leq O\left(n_v^{-1}\right) \quad , \tag{9}$$

for both sparse and dense graphs G.

Proof of this result may be found in the appendix, in Section A.5. The theorem establishes a rate at which—in large networks, whether sparse or dense, with independent and homogeneous low-rate errors—the distribution of the discrepancy D_E tends to that of an appropriate Skellam distribution, i.e., symmetric and centered on zero, with variance 2λ. The same rate may be established using more standard arguments from Chen-Stein theory, the proof of which we also include in the appendix, for completeness. These latter arguments, of course, only hold in the case of independence assumed here, and do not extend generally to the case of dependence in the edge noise.

To put the rate established in the above theorem in better context, it is interesting to compare to the case where a normal distribution is used instead to approximate that of the discrepancy D_E. The following theorem, proof of which also may be found in Section A.5, provides both upper and lower bounds.

Theorem 7 *Let $\sigma^2 = Var(D_E)$. Under the same conditions as Theorem 6 , in the case of sparse graphs*

$$\Omega\left(\log^{-1} n_v\right) \leq d_{KS}\left(\, D_E/\sigma \,,\, N(0,1)\,\right)) \leq O\left(\log^{-1/2} n_v\right) \quad , \tag{10}$$

while in the case of dense graphs,

$$\Omega\left(n_v^{-1}\right) \leq d_{KS}\left(\, D_E/\sigma \,,\, N(0,1)\,\right)) \leq O\left(n_v^{-1/2}\right) \quad , \tag{11}$$

where $N(0,1)$ refers to a standard normal random variable.

These two theorems together indicate that in this context a Skellam approximation is clearly superior to a normal for sparse graphs, and they suggest that it can be better as well for dense graphs. These statements are supported by the results of a small simulation study, shown in Figure 1. There we compared the two approximations as n_v ranges from 100 to 1000 to $10,000$, for error rates λ defined to be constant, logarithmic, square root, or linear functions of n_v. For the sparse and dense cases, we let $|E|$ equal $n_v \log n_v$ and $n_v(n_v-1)/4$, respectively. Looking at the sparse case, for when $\lambda = \log n_v$, the Skellam approximation clearly dominates the normal. However, interestingly, this dominance continues even when the error rate is set equal to $n_v^{1/2}$. Only once the error rate is n_v do we see the normal approximation begin to overtake the Skellam approximation. Note that by this stage, $\beta = O(1)$, and so essentially there is no 'signal' standing out from the 'noise'. Similarly, looking at the dense case, we see that the Skellam approximation is better than the normal approximation at all error rates, including, in particular, when the error rate equals n_v, the case addressed by the above two theorems.

In summary, in the independent case, the Skellam distribution dominates the normal as an approximation when there can be expected to be a clear graph 'signal' standing out against the 'noise' induced by underlying low-rate measurement errors.

4.2 Edge Counts Under Dependent Edge Noise

Again, as just above, consider the context wherein counting edges is of interest, so that $\eta(G) = |E|$ and our goal is to characterize the accuracy with which $D_E = |\hat{E}| - |E|$ is approximated by a Skellam(λ, λ) random variable. Now, however, we assume that the error associated with construction of the empirical graph $\hat{G}$ will involve dependency across (non)edges. That is, the random variables Y_{ij} are now dependent. A precise characterization of such dependency is typically problem-specific and, more often than not, nontrivial in nature. Here, for the purposes of illustration, we instead provide certain results of a general nature, working from the bound (8) of Theorem 2.

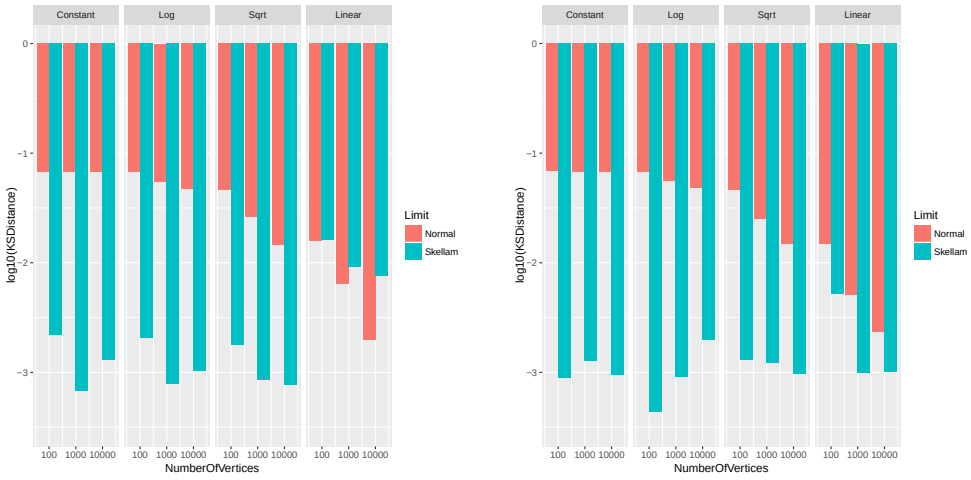

Figure 1: (Log)Kolmogorov-Smirnov distance between Skellam and standard normal approximations to the distribution of discrepancy D_E in edge counts under independent errors. *Left:* Sparse case. *Right:* Dense case.

Of the two terms in (8), the first term $||\Delta f||$ is again known to behave as $O(\lambda^{-1})$, by Theorem 4. On the other hand, control of the second term, in brackets, requires some care. For example, naive inter-change of absolute values and summations with expectation yields that

$$p_k \mathbb{E}\left|U - U_k^{(L)}\right| \leq p_k^2 + \sum_{k \neq j} p_k p_j + E[L_k L_j] + \sum_{\ell=1}^{m} p_k q_\ell + E[L_k M_\ell] \ ,$$

and similarly for $q_k \mathbb{E}\left|U - U_k^{(M)}\right|$. Unfortunately, it is straightforward to show that for the dependent error version of the problem considered in the previous subsection (i.e., involving independent and homogeneous low-rate errors on large-spare networks) the bound we obtain for $d_{KS}\left(D_E, \text{Skellam}(\lambda, \lambda)\right)$ will be no better than $O(\lambda)$—regardless of the nature of the dependency among the Y_{ij}.

One possible approach to a more subtle handling of these terms is motivated by considerations of hypothesis testing. Suppose that the L_k correspond to indicators of Type I error for n tests under their corresponding null hypotheses, and the M_k, to indicators of Type II error for m tests under their corresponding alternative hypotheses. Furthermore, suppose that the corresponding test statistics are all defined on the same scale and compared to the same threshold. Moreover, for simplicity, we assume these statistics all have non-negative values and that their distribution under the null sits to the left of that under the alternative, so that more extreme positive values tend to support the alternative. In this setting, if we know, for example, that $L_1 = 1$, we know that at least one rejection of a null hypothesis has occurred, indicating that the threshold sits to the left of the right-most extreme of the empirical null distribution. Accordingly, we are inclined to expect that there may be other such rejections of the null, i.e., other Type I errors. At the same time, we would expect fewer Type II errors, i.e., fewer M that equal 1. Conversely, if we see a Type II error, say $M_1 = 1$, it can be argued that we would be inclined to expect more Type II errors and, at

the same time, fewer Type I errors. These arguments are heuristic, of course, although intuitively appealing. From a practical perspective, their validity would need to be examined within specific applied contexts. However, that work lies beyond the scope of the present manuscript.

Together these high-level arguments suggest that a reasonable generic model for these errors is one in which there is positive correlation *within* the L's and M's, respectively, but negative correlation *between*. The conditions of the following theorem capture this notion, which in turn allow us to produce a sensible bound, improving on that of Theorem 8.

Theorem 8 *Let $\tilde{L}_j^{L_k}$ and $\tilde{M}_\ell^{L_k}$ be random variables distributed as L_j and M_ℓ respectively, conditional on $L_k = 1$. Similarly, let $\tilde{L}_j^{M_k}$ and $\tilde{M}_\ell^{M_k}$ be distributed as L_j and M_ℓ, conditional on $M_k = 1$. Suppose that*

i. $\tilde{L}_j^{L_k} \geq L_j$ and $\tilde{M}_\ell^{L_k} \leq M_\ell$, for $j \neq k$ and $\ell = 1, \ldots, m$, while

ii. $\tilde{L}_j^{M_k} \leq L_j$ and $\tilde{M}_\ell^{M_k} \geq M_\ell$, $j = 1 \ldots, n$ and $\ell \neq k$.

Then

$$d_{KS}(U, W) \leq ||\Delta f|| \left\{ \mathit{Var}(U) - (\lambda_1 + \lambda_2) \right\} \quad , \tag{12}$$

where $W \sim$ Skellam(λ_1, λ_2), with λ_1, λ_2 defined as in Theorem 2.

The proof of this theorem is given in the appendix, in Section A.6. We note that for a collection of binary random variables to satisfy conditions (i) and (ii) in the above theorem, it is sufficient, for example, to generate a vector of positively associated random variables $(L_1, \ldots, L_n, 1 - M_1, \ldots, 1 - M_m)$. The L's and the M's will then be positively associated within, but negatively associated between, which in turn implies the conditions (i.e., analogous to positive and negative relatedness, respectively). See Barbour and Chen (2005, p. 78), for example, for a brief summary of these latter notions and their relationships.

With this theorem, the following then holds for large networks with dependent and homogeneous errors, when the dependency is of the nature just defined.

Corollary 9 *Suppose that the collections of edge indicator random variables $\{Y_{ij}\}_{\{i,j\} \in E^c}$ and $\{Y_{ij}\}_{\{i,j\} \in E}$ satisfy conditions (i) and (ii) of Theorem 8 , playing the roles of the $L's$ and $M's$, respectively. Then under assumptions (A1)-(A2),*

$$d_{KS}\left(D_E \, , \, \mathit{Skellam}(\lambda, \lambda) \right) = O\left(\frac{\mathit{Var}(D_E) - 2\lambda}{2\lambda} \right) \quad .$$

This result can be compared to that of Theorem 6 , where the edge noise was independent and the error in approximating by a Skellam decayed like n_v^{-1}. By way of contrast, Corollary 9 tells us that in order to achieve a decay in approximation error like $O(f(n_v))$, we must have $\text{Var}(D_E) = 2\lambda\left(1 + O(f(n_v))\right)$.

More generally, the quality of the approximation of D_E by a Skellam will be influenced by the nature of the dependency in the errors, as the latter manifests itself through the overall variance $\text{Var}(D_E)$. The nature of that dependency is highly problem-specific, and a detailed investigation of possible cases is well beyond the scope of the present paper, although we provide some initial insight later as it arises in the context of counting higher-order subgraphs, through Theorem 10 in Section 5.

However, nontrivial insight also can be gained into the influence of the level of dependency on accuracy through numerical work under the following simple model.

For a vector of binary random variables $(L_1, \ldots, L_n, 1 - M_1, \ldots, 1 - M_m)$, let $S = D + m$, where $D = \sum_{i=1}^{n} L_i - \sum_{i=1}^{m} M_i$. We equip S with a distribution of the form

$$\mathbb{P}(S = k) \propto \binom{n+m}{k}^{\nu-1} \mathbb{P}(U + V = k) \quad , \tag{13}$$

for ν a real number, where U and V are binomial, with parameters (n, p) and (m, q), respectively. This distribution is a rescaling of that of the sum of two independent binomials, in a spirit analogous to the Conway-Maxwell binomial (COMB) distribution introduced recently by Kadane (2014). The COMB distribution is a simple extension of the binomial distribution that introduces dependency among the corresponding Bernoulli random variables. Our proposed distribution for S in (13) involves two binomial random variables, for which the corresponding Bernoulli random variables are dependent both within and between the two. Accordingly, we call this a COMB2 distribution.

Now impose assumption (A2) on this model. Since the assumption implies that $\mathbb{E}[D] = 0$, it follows that necessarily we must have $\mathbb{E}[S] = m$. Furthermore, the limiting Skellam distribution in Corollary 9 will be symmetric under this assumption. Symmetry can be imposed here on the distribution of S, and hence D, by taking $n = m$ and $q = 1 - p$. Therefore, we let $|E| = \binom{n_v}{2}/2 = n_v(n_v - 1)/4$ and $\alpha = 1 - \beta$. Note that this choice of $|E|$ means that our numerical work pertains to the case of dense graphs. (We are unable to exhibit a sparse variant of the COMB2 model with the necessary characteristics above.)

Note that when $\nu = 1$, the binary random variables $(L_1, \ldots, L_n, 1 - M_1, \ldots, 1 - M_m)$ are independent. On the other hand, proceeding along lines of reasoning similar to those in Kadane (2014), it can be argued that the COMB2 distribution, with the parameter constraints just described, renders the $(L_1, \ldots, L_n, 1 - M_1, \ldots, 1 - M_m)$ positively associated when $\nu < 1$, with the mass being transferred increasingly to the endpoints of the support of the distribution of D as $\nu \to -\infty$. As a result, per the discussion immediately following Theorem 8, the particular COMB2 distribution we have defined can be used to simulate network edge data in a way that satisfies the conditions of Corollary 9.

In Figure 2 are shown the results of numerical work calculating the Kolmogorov-Smirnov distance between the Skellam and standard normal approximations to the distribution of the discrepancy D_E in edge counts under the COMB2 distribution, for $\nu = 0, 0.5$, and 1.0. The noise levels used here are the same as used earlier in producing Figure 1. We see that the accuracy of the Skellam distribution decays slightly with increasing dependency in the errors, and with increasing noise levels.

5. Discussion

The propagation of uncertainty in network analysis is a topic that currently lags the field in development. Despite almost 15 years of work in the modern 'network science' era, on a vast array of topics, from researchers in many different disciplines, there remains a sizable gap in our understanding of how 'low level' errors (i.e., at the level of declaration of edge / non-edge status between vertex pairs) propagate to 'high level' summaries (e.g., subgraph counts, centralities, etc.). As a result, in most practical work, network summary statistics are cited without any indication of likely error.

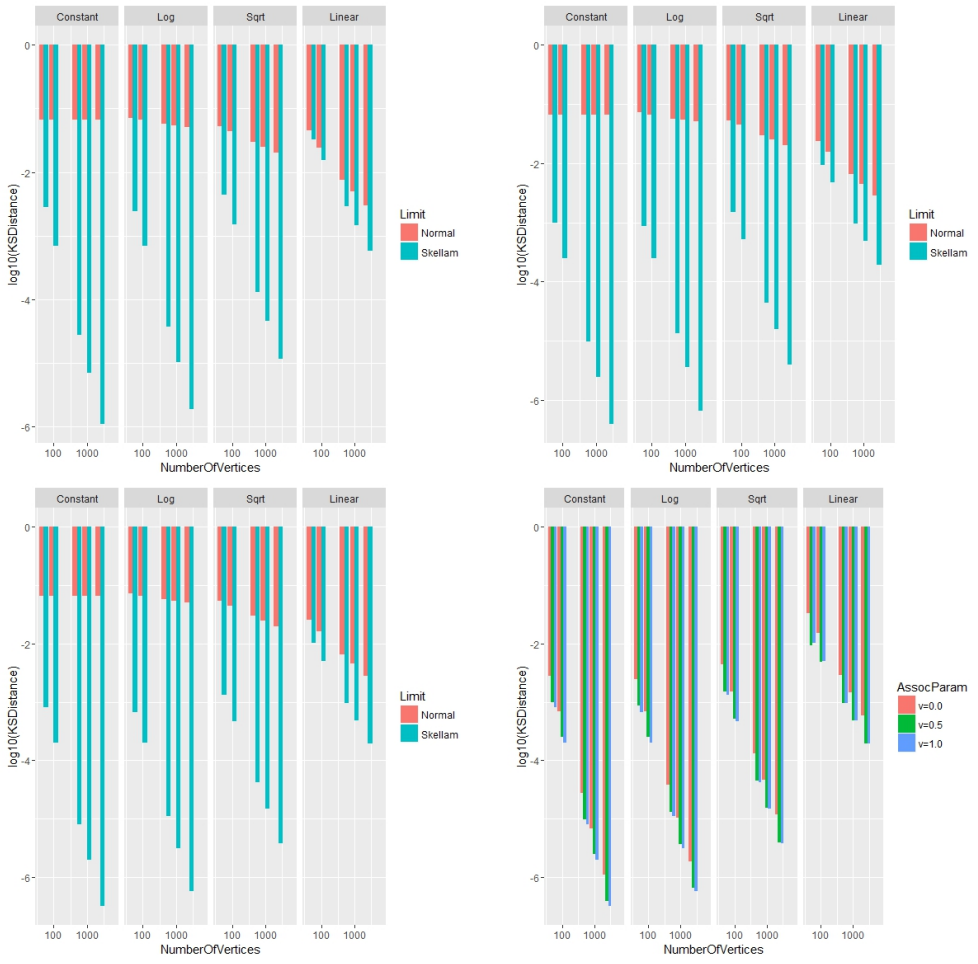

Figure 2: (Log)Kolmogorov-Smirnov distances to the distribution of the discrepancy D_E in edge counts under the COMB2 distribution, with $\nu = 0$ (top left), $\nu = 0.5$ (top right), and $\nu = 1.0$ (bottom left), for the Skellam and standard normal approximations. Also shown is a comparison of Skellam approximations as a function of ν (bottom right).

Our contributions in this paper are aimed at helping to begin laying a foundation for work in this area, with a focus on establishing an initial understanding of the distributional behavior of certain simple network summary statistics. Our choice to work with subgraph counts is both natural and motivated by convenience, whereas our emphasis on the specific case of large networks with low-rate measurement error is intended to capture a sizable fraction of what arguably is encountered in practice.

Ultimately, however, this initial work leaves us still well short of natural practical endpoints in applied statistics and machine learning for network science—for example, even to be able to do hypothesis testing or construct confidence intervals for graph-based parameters of interest (i.e., for $\eta(G)$). To close this gap is well beyond the scope of this paper. Indeed, we anticipate that it

will require the contributions of many papers in this emerging topic area, from researchers across a combination of probability, statistics and machine learning, and domain areas. Accordingly, we devote the remainder of this paper to a brief discussion of what we view as key challenges in terms of applied probability, statistical and machine learning methodology, and practical applications.

5.1 Challenges in Applied Probability

From the perspective of applied probability—including random graph theory—our problem formulation appears to be new. Our formulation is reminiscent of the type of 'signal plus noise' model commonly used in nonparametric function estimation and digital signal processing. Notably, in our formulation the true underlying graph G is fixed. This necessitates a different treatment than, say, traditional analysis of subgraph counts in classical Erdős-Rényi random graphs (e.g., Janson et al., 2011, Chs 3 & 6). In the special case where an Erdős-Rényi model is assumed, as well as assuming independence among the measurement errors, and the analysis is done without conditioning on G, then the problem could be viewed as involving a classical random graph wherein the (marginal) probability of an edge is a function not only of an initial edge frequency but also the Type I and Type II error rates. In general, however, either when G is fixed, as assumed in this paper, or from some other class of random graph models (e.g., various models with heterogeneous degree distributions), or when the measurement errors are dependent, the problem is more involved. By conditioning on G, our formulation allows us to focus our analysis firstly on a high-level notion of Type I and II errors among (non)edges, and then secondly on the manner in which the structure of the underlying graph G may interact with those errors.

We view our work in part as laying a key initial piece of the foundation on an important new class of problems in applied probability. However, we have provided a detailed analysis only for the most fundamental of subgraph count statistics, i.e., the number of edges in a network. Our initial work on extension to counts for subgraphs of higher order suggests that the problem becomes increasingly nontrivial. Specifically, the interaction of noise level, graph topology, and choice of subgraph would appear to need to be studied with care.

The following general theorem should be useful in exploring further in this direction.

Theorem 10 *Let H be a given subgraph of interest. Re-express the difference D in subgraph counts defined in equation (3) as*

$$D_H = \sum_{H' \notin \mathcal{C}_H} L_{H'} - \sum_{H' \in \mathcal{C}_H} M_{H'} \ ,$$

for $\mathcal{C}_H = \{H' \subseteq K_{n_v} \, : \, H' \cong H, H' \subseteq G\}$, where $L_{H'}$ and $M_{H'}$ are indicator variables of Type I and Type II error, respectively, for a subgraph H'. Under the assumption of independent edge noise,

$$d_{KS}(D_H, W) \leq ||\Delta f|| \left\{ \mathit{Var}(D_H) - (\lambda_1 + \lambda_2) \right\} \ , \tag{14}$$

where $W \sim$ Skellam(λ_1, λ_2), with $\lambda_1 = \sum_{H' \notin \mathcal{C}_H} p_{H'}$ and $\lambda_2 = \sum_{H' \in \mathcal{C}_H} q_{H'}$, for $p_{H'} = \mathbb{E}\left[L_{H'}\right]$ and $q_{H'} = \mathbb{E}\left[M_{H'}\right]$.

This result follows directly from application of Theorem 8 and the comment immediately following that theorem. In particular, each of the indicator random variables $L_{H'}$ and $1 - M_{H'}$ may be expressed as a product of $n_{v(H)}$ choose two binary random variables, where $n_{v(H)}$ is the order of the subgraph H. Since these products are non-decreasing functions of their arguments, and their

arguments are independent, it follows that the collection of random variables defined by the union of the $L_{H'}$ and the $1 - M_{H'}$ are positively associated (e.g., Esary et al., 1967).

Application of this theorem to specific choices of subgraphs H requires calculation or bounding of the two key elements within brackets in (14), i.e., $\text{Var}(D_H)$ and $\lambda_1 + \lambda_2$, as well as control of the constant $||\Delta f||$. For the case of independent edge noise (which, nevertheless, yields *dependent* indicator variables $L_{H'}$ and $M_{H'}$), there is some hope that the former two quantities may be bounded through straightforward calculations, at least for low-order subgraphs. As for the latter quantity, under Conjecture 5 this term is controlled by a term of order $(\lambda_1 + \lambda_2)^{-1}$, but this conjecture, while supported by numerical work, remains to be proven.

By way of illustration, consider the problem of counting two-stars on the one-dimensional toroidal lattice of degree $2r$, $r \geq 1$. That is, the graph G is composed of a ring, with each vertex connected to its r nearest neighbors in each direction. The subgraph H is simply a chain of length two, i.e., a subgraph on three vertices $\{i, j, k\}$ with edges $\{i, j\}$ and $\{i, k\}$. In this case, the total number of two-stars is $\eta(G) = n_v\binom{2r}{2}$. Under homogeneous Type I and Type II error rates α and β, the parameters λ_1 and λ_2 in Theorem 10 take the values

$$\lambda_1 = 2rn_v\left(n_v - 2r - 1\right)\alpha\left(1 - \beta\right) + n_v\binom{n_v - 2r - 1}{2}\alpha^2$$

and

$$\lambda_2 = n_v\binom{2r}{2}\left[1 - (1 - \beta)^2\right] \quad .$$

Here the two terms defining λ_1 derive from counting the number of vertex triples in G that possess only one edge or no edges among them, respectively. Under assumptions (A1) - (A4), and supposing that G is sparse, λ_1 and λ_2, are both $O(\log^2 n_v)$, so that neither of the sums in D_H dominates the other. However, λ_1 is in general not equal to λ_2, and so Conjecture 5 is pertinent here. At the same time, calculation of the variance in (14) is decidedly more involved, as it requires evaluation of the covariances $\text{Cov}\left(L_{H'}, L_{H''}\right)$, $\text{Cov}\left(M_{H'}, M_{H''}\right)$, and $\text{Cov}\left(L_{H'}, M_{H''}\right)$ over all pairs of vertex-triples, corresponding to putative subgraphs H' and H''. Under independent edge noise, the calculations simplify to some extent, in that these covariances are non-zero only in the case where the relevant pair of vertex triples is defined in terms of a total of only either three or four unique vertices. (When five or six unique vertices define the two triples, the indicators L and/or M are functions of two independent sets of edge indicators Y_{ij}.) Nevertheless, there still remain several subcases to enumerate for each of the three forms of covariance above in order to complete the calculation, rendering the exercise arguably feasible but tedious. Ultimately, however, since the denominator in (14) is $O\left(\log^2 n_v\right)$ under Conjecture 5, we see that under this particular noise level, with this particular graph topology, the convergence to Skellam is unlikely to be fast. Numerical experiments confirm this expectation.

This small illustration is sufficient to illustrate the point we assert above, i.e., that the interaction of noise level, graph topology, and choice of subgraph are likely to be nontrivial in the class of problems introduced here in this paper, and will likely require substantial and careful study in applied probability. Of course, looking even more broadly, by restricting our attention to subgraph counts we have only considered a portion of the full universe of network summary statistics. Clearly there are other classes of summaries of interest to be explored, and which can be expected to potentially display different behavior. Summaries relying on shortest path behavior (e.g., the diameter of a graph, between centrality of nodes, etc.) are a natural example.

5.2 Challenges in Statistics and Machine Learning Methodology

From the perspective of statistics and machine learning, the work in this paper is largely a necessary means to an end. In the context of network-based propagation of error, a central goal of interest in these fields would be, for example, to develop methods of hypothesis testing or confidence interval construction for graph-based parameters of interest (i.e., for $\eta(G)$) based on an observed graph $\hat{G}$. Alternatively, one might wish to use such network summaries (or perhaps even the networks themselves) as input into a regression or classifier. Here, again, much work remains to be done.

For parameter estimation and testing in the spirit of classical statistics (e.g., Statistics 101), it is typical to first understand the probability distribution underlying a quantity of interest. That is the goal of this manuscript and, for an important subclass of problems in this space, our results offer some initial insight. But knowledge of this distribution does not necessarily yield a tool for statistical inference. For example, to apply the Skellam results here in practice would require knowledge of the Type I and II error rates. If that knowledge is lacking (as generally is the case), then they must be estimated. How best to perform such estimation will presumably depend on context. To the best of our knowledge, this question has not been explored broadly in the literature. See Section 5.3 for additional comments in this direction.

In our work here, we have purposely taken an agnostic stance on the nature of G. However, it can be expected that additional leverage on the propagation of error problem—and, in particular, related problems of statistical inference—can be gained by incorporating additional model structure into the problem. For example, within the statistics and machine learning literatures, arguably the canonical methodology for producing an estimate $\hat{G}$ of a graph G from data is the neighborhood selection approach of Meinshausen and Bühlmann (2006) and the many derivatives thereof. In this approach, we assume we have available n independent and identically distributed observations of a n_v-dimensional random vector $\mathbf{X}$ that follows a multivariate normal distribution, with mean 0 and covariance matrix Σ. The graph G in this setting is taken to be the so-called concentration graph, defined so as to have an edge between a given vertex pair i and j if and only if the i, j-th entry of the concentration matrix Σ^{-1} is non-zero. Since it is typical to have $n \ll n_v$ in practice, methods of inference for this problem generally make use of complexity-penalized optimization. The original proposal in Meinshausen and Bühlmann (2006) makes use of ℓ_1-penalized regression. In what appears to be a first, the combination of Gaussian graphical models and hypothesis testing for network structure has been explored by Neykov et al. (2016) in recent work. There they characterize the complexity of what they call combinatorial inference problems through the use of minimax lower bounds. Algorithms are provided for tests that achieve those lower bounds.

Examples of other related work in this space include methodology for predicting network topology or attributes with models that explicitly include a component for network noise (e.g., Jiang et al., 2011; Jiang and Kolaczyk, 2012), the 'denoising' of noisy networks (e.g., Chatterjee, 2015), and the adaptation of methods for vertex classification using networks observed with errors Priebe et al. (2015). Also related is recent work by Cai et al. (2017) on the detection of ring lattice structure in the presence of random connections (i.e., small worlds). While there are undoubtedly other examples in this space that we have not cited here, it is nevertheless the case that we are currently still only at the start of methodological development of this type.

Finally, we note that network summary statistics have also long been used in practice as input to regressions and classifiers. A canonical example in the field of marketing is Hill et al. (2006), where the focus is on the identification of early adopters through the use of consumer networks,

with neighborhood summaries (e.g., proportion of neighbors adopting a product) being used as predictors. Similar use has been made in the social sciences. For example, Hahm et al. (2012) use vertex centrality measures from a friendship network as input in an effort to predict adolescents at risk of binge drinking. In these and other cases, it is certainly conceivable that the network summary measures are observed with error. In that case, incorporation of ideas like errors-in-variables regression might be appropriate (Fuller, 2009). Characterization of distributional properties for network summary measures, as in this paper, can be an important step towards understanding an appropriate error model for the predictors.

5.3 Challenges for Practical Applications

When all is said and done, probabilistic and statistical / machine learning development in the area of network-based propagation of error will only have significant practical impact when paired with an appropriate degree of context-specific understanding of the nature and magnitude of errors that arise in network construction. Certainly in most applied settings it is widely recognized by practitioners that there is measurement error associated with common types of network constructions. And in many settings the general issue has received nontrivial attention, such as, for example, in the context of protein-protein interaction networks (e.g., Hart et al., 2006) or social networks (e.g, Almquist, 2012).

In order to interface in a productive manner with theoretical and methodological development, such as offered in this paper, it will be important to be able to quantify the noise levels in applications. The difficulty of this task, of course, can vary greatly by context, as well as can the magnitude of the noise. In the case of protein-protein interaction networks, for example, it is understood that Type I and Type II error rates associated with traditional affinity binding experiments can be alarmingly high. Hart et al. (2006) summarized the fairly substantial literature on quantifying these rates, finding that values of $\alpha \approx 0.35$ and $\beta \approx 0.40$ were not atypical. In this case, such rates derive largely from experimental evidence. Alternatively, Balachandran et al. (2013a) employ empirical null principles (Efron, 2012) for correlation-based inference of a small gene regulatory network in yeast to estimate error rates $\alpha \approx 0.004$ and $\beta \approx 0.009$. But these later should be taken with a grain of salt. Cosgrove et al. (2010) show that such nominal estimates of error rates can be decidedly off from actual error rates when the underlying model assumptions are violated in ways that can often be expected in practice (e.g., due to experimental batch effects).

Finally, it is important to note that in some areas of network science the notion of 'ground truth'—in the sense of what constitutes a true (non)edge—can be a slippery concept. For example, there continues to be much discussion and debate in computational neuroscience as to what functional (or, alternatively, effective) connectivity represents, as inferred from data (e.g., from neuroimaging measurements), in contrast to anatomical (or, alternatively, physical or structural) connectivity. See, for instance, the review by Friston (2011). Such concerns further complicate the task of defining and quantifying appropriately the notions of Types I and II errors used in this paper.

Acknowledgments

We thank Kostas Spiliopoulos and Tasso Kaper for helpful discussions. This work was supported in part by AFOSR award 12RSL042 and NSF grant CNS-0905565.

Appendix A.

A.1 Proof of Theorem 1

Proof We begin with the operator,

$$\mathcal{A}\left[f\left(k\right)\right] = \lambda_1 f\left(k+1\right) - kf\left(k\right) - \lambda_2 f\left(k-1\right)$$

with the intent of showing that the random variable $W \sim \text{Skellam}\left(\lambda_1, \lambda_2\right)$ if and only if $\mathbb{E}\mathcal{A}\left[f\left(W\right)\right] = 0$ for any bounded function $f : \mathbb{Z} \mapsto \mathbb{R}$.

We begin with the necessity direction and the computation of

$$\begin{aligned}
\mathbb{E}\mathcal{A}\left[f\left(W\right)\right] &= \mathbb{E}\left[\lambda_1 f\left(W+1\right) - Wf\left(W\right) - \lambda_2 f\left(W-1\right)\right] \\
&\propto \sum_{k=-\infty}^{\infty}\left[\lambda_1 f\left(k+1\right) - kf\left(k\right) - \lambda_2 f\left(k-1\right)\right]\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k
\end{aligned}$$

where $\propto$ is to be read as "proportional to," and as shorthand, we write I_k for $I_k(2\sqrt{\lambda_1\lambda_2})$. By standard properties of I_k (e.g., Abramowitz and Stegun, 1972) we have that

$$I_{k-1} - I_{k+1} = \frac{k}{\sqrt{\lambda_1\lambda_2}} I_k$$

or, in other words,

$$\sqrt{\lambda_1\lambda_2}\frac{I_{k-1}}{I_k} - \sqrt{\lambda_1\lambda_2}\frac{I_{k+1}}{I_k} = k. \tag{15}$$

This means that

$$\begin{aligned}
\mathbb{E}\mathcal{A}\left[f\left(W\right)1\left\{W \le n\right\}\right] &\propto \sum_{k=-\infty}^{n}\left[\sqrt{\frac{\lambda_1}{\lambda_2}} f\left(k+1\right) - \frac{I_{k-1}}{I_k} f\left(k\right) + \right. \\
&\qquad \left. \frac{I_{k+1}}{I_k} f\left(k\right) - \sqrt{\frac{\lambda_2}{\lambda_1}} f\left(k-1\right)\right]\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k \\
&= \sum_{k=-\infty}^{n}\left[\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{k+1} I_k f\left(k+1\right) - \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_{k-1} f\left(k\right)\right] \\
&\quad + \sum_{k=-\infty}^{n}\left[\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_{k+1} f\left(k\right) - \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{k-1} I_k f\left(k-1\right)\right] \\
&= \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{n+1} I_n f\left(n+1\right) + \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^n I_{n+1} f\left(n\right).
\end{aligned}$$

Now, since f is bounded,

$$\lim_{n\to\infty}\left\{\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{n+1} I_n f\left(n+1\right) + \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^n I_{n+1} f\left(n\right)\right\} = 0$$

so that by monotone convergence,

$$\begin{aligned}\mathbb{E}\left[\mathcal{A}\left[f(W)\right]\right] &= \lim_{n\to\infty} \mathbb{E}\mathcal{A}\left[f\left(W\right)1\left\{W\leq n\right\}\right] \\ &= \lim_{n\to\infty}\left\{\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{n+1} I_n f\left(n+1\right) + \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{n} I_{n+1} f\left(n\right)\right\} \\ &= 0\end{aligned}$$

which proves the claim.

To prove sufficiency, we begin with $\mathbb{E}\mathcal{A}\left[f\left(W\right)\right]=0$ and suppose that $f_k\left(j\right)=1\left\{j=k\right\}$ for some $j\in\mathbb{Z}$ in which case

$$\lambda_1 p\left(k-1\right) - kp\left(k\right) - \lambda_2 p\left(k+1\right) = 0$$

where $p\left(k\right)=\mathbb{P}\left(W=k\right)$. An *ansatz* of

$$S\left(k\right)=\left(\sqrt{\tfrac{\lambda_1}{\lambda_2}}\right)^k I_k\left(2\sqrt{\lambda_1\lambda_2}\right) \quad \text{and} \quad T\left(k\right)=\left(\sqrt{\tfrac{\lambda_1}{\lambda_2}}\right)^k K_k\left(2\sqrt{\lambda_1\lambda_2}\right)$$

shows that S and T form two linearly independent solutions to this second order linear difference equation, where $I_k(x)$ and $K_k(x)$ are the modified Bessel functions of the first and second kinds. Thus, we know that the general solution is given by,

$$p\left(k\right)=C_1 S\left(k\right)+C_2 T\left(k\right)$$

for some constants $C_1, C_2\in\mathbb{R}$.

Now, to determine the constants C_1 and C_2 we appeal to the fact that $\sum_{k=-\infty}^{\infty} p\left(k\right)=1$. Since $I_k, K_k > 0$ for all $k\in\mathbb{Z}$ and $\sum_{k=-\infty}^{\infty} K_k=\infty$ it must be that $C_2=0$. Now, consider the generating function

$$e^{\frac{z}{2}(t+1/t)} = \sum_{k=-\infty}^{\infty} t^k I_k\left(z\right)$$

which means that

$$\begin{aligned} C_1 &= \frac{1}{\sum_{k=-\infty}^{\infty}\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k\left(2\sqrt{\lambda_1\lambda_2}\right)} \\ &= \frac{1}{e^{\sqrt{\lambda_1\lambda_2}\left(\sqrt{\frac{\lambda_1}{\lambda_2}}+\sqrt{\frac{\lambda_2}{\lambda_1}}\right)}} \\ &= e^{-(\lambda_1+\lambda_2)}\end{aligned}$$

so that

$$p\left(k\right)=e^{-(\lambda_1+\lambda_2)}\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k\left(2\sqrt{\lambda_1\lambda_2}\right)$$

so that $W\sim\text{Skellam}\left(\lambda_1,\lambda_2\right)$. ■

A.2 Proof of Theorem 2

Proof Given that f_x is a solution to $\mathcal{A}[f_x(k)] = g_x(k)$, we have

$$\lambda_1 f_x(k+1) - k f_x(k) - \lambda_2 f_x(k-1) = 1(k \leq x) - \mathbb{P}[W \leq x].$$

Substituting $k = U$ and taking expected values, we obtain,

$$|\mathbb{P}[U \leq x] - \mathbb{P}[W \leq x]| = |\mathbb{E}[\lambda_1 f_x(U+1) - U f_x(U) - \lambda_2 f_x(U-1)]| \quad . \tag{16}$$

Next, recall from (5) that $U = \sum_{k=1}^{n} L_k - \sum_{k=1}^{m} M_k$. Since $\lambda_1 = \sum_{k=1}^{n} p_k$ and $\lambda_2 = \sum_{k=1}^{m} q_k$, we have after conditioning on L_k and M_k,

$$\begin{aligned}
&|\mathbb{E}[\lambda_1 f_x(U+1) - U f_x(U) - \lambda_2 f_x(U-1)]| \\
&= \left| \sum_{k=1}^{n} \mathbb{E}[p_k f_x(U+1) - L_k f_x(U)] + \sum_{k=1}^{m} \mathbb{E}[M_k f_x(U) - q_k f_x(U-1)] \right| \\
&= \left| \sum_{k=1}^{n} p_k \left(\mathbb{E}[f_x(U+1)] - \mathbb{E}[f_x(U)|L_k = 1] \right) + \sum_{k=1}^{m} q_k \left(\mathbb{E}[f_x(U)|M_k = 1] - \mathbb{E}[f_x(U-1)] \right) \right| \\
&= \left| \sum_{k=1}^{n} p_k \left(\mathbb{E}\left[f_x(U+1) - f_x\left(U_k^{(L)} + 1 \right) \right] \right) + \sum_{k=1}^{m} q_k \left(\mathbb{E}\left[f_x\left(U_k^{(M)} - 1 \right) - f_x(U-1) \right] \right) \right| \\
&\leq \sum_{k=1}^{n} p_k ||\Delta f|| \mathbb{E}\left| U - U_k^{(L)} \right| + \sum_{k=1}^{m} q_k ||\Delta f|| \mathbb{E}\left| U - U_k^{(M)} \right| \\
&= ||\Delta f|| \left[\sum_{k=1}^{n} p_k \mathbb{E}\left| U - U_k^{(L)} \right| + \sum_{k=1}^{m} q_k \mathbb{E}\left| U - U_k^{(M)} \right| \right].
\end{aligned}$$

Combining this with (16) yields the result. ■

A.3 Proof of Theorem 3

Proof First, consider the solution to

$$\lambda_1 f(k+1) - k f(k) - \lambda_2 f(k-1) = g(k), \tag{17}$$

for some bounded function $g : \mathbb{Z} \mapsto \mathbb{R}$, with the boundary condition

$$\lim_{k \to -\infty} \left(\sqrt{\frac{\lambda_1}{\lambda_2}} \right)^k I_k f(k) = 0. \tag{18}$$

We use (15) to substitute for k in (17). Then, multiplying both sides of (17) by $\left(\sqrt{\lambda_1/\lambda_2}\right)^k I_k$, we obtain,

$$\begin{aligned}
&\lambda_1 \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k f(k+1) - \lambda_1 \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{k-1} I_{k-1} f(k) \\
&+ \lambda_2 \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{k+1} I_{k+1} f(k) - \lambda_2 \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k f(k-1) \\
&= \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k g(k)
\end{aligned}$$

which is the same as,

$$\begin{aligned}
&\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{k+1} I_k f(k+1) - \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_{k-1} f(k) \\
&+ \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_{k+1} f(k) - \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{k-1} I_k f(k-1) \\
&= \frac{1}{\sqrt{\lambda_1 \lambda_2}} \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k g(k).
\end{aligned}$$

Notice that we have grouped terms together so that summing over k yields a telescoping sum. So, summing over $k \in \{-\infty, \ldots, n\}$ and using the boundary condition (18),

$$\left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^{n+1} I_n f(n+1) + \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^n I_{n+1} f(n) = \frac{1}{\sqrt{\lambda_1 \lambda_2}} \sum_{k=-\infty}^{n} \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k g(k).$$

Now, multiplying both sides by $(-1)^{n+1}/(I_n I_{n+1})$ and summing over $n \in \{c, c+1, \ldots, m\}$ for $m > c$ and over $n \in \{m, m+1, \ldots, c-1\}$, for some initial condition $c \in \mathbb{Z}$ and $f(c) \in \mathbb{R}$, we obtain

$$f(m) = \begin{cases}
(-1)^m \left(\sqrt{\frac{\lambda_2}{\lambda_1}}\right)^m I_m \left[(-1)^c \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^c \frac{1}{I_c} f(c) \right. \\
\left. + \frac{1}{\sqrt{\lambda_1\lambda_2}} \sum_{n=c}^{m-1} \frac{(-1)^{n+1}}{I_n I_{n+1}} \sum_{k=-\infty}^{n} \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k g(k)\right] & \text{if } m > c \\
(-1)^m \left(\sqrt{\frac{\lambda_2}{\lambda_1}}\right)^m I_m \left[(-1)^c \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^c \frac{1}{I_c} f(c) \right. \\
\left. - \frac{1}{\sqrt{\lambda_1\lambda_2}} \sum_{n=m}^{c-1} \frac{(-1)^{n+1}}{I_n I_{n+1}} \sum_{k=-\infty}^{n} \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k g(k)\right] & \text{if } m < c.
\end{cases}$$

Note that if

$$g(k) = g_x(k) = 1\{k \le x\} - \mathbb{P}(W \le x)$$

then

$$\sum_{k=-\infty}^{n} \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k g(k) = \begin{cases} e^{\lambda_1+\lambda_2} \mathbb{P}(W \le n) \mathbb{P}(W > x) & \text{if } n \le x \\ e^{\lambda_1+\lambda_2} \mathbb{P}(W \le x) \mathbb{P}(W > n) & \text{if } n \ge x \end{cases}$$

since, for example if $n \leq x$

$$\begin{aligned}
\sum_{k=-\infty}^{n} \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^k I_k g(k) &= e^{\lambda_1+\lambda_2} \sum_{k=-\infty}^{n} \mathbb{P}(W=k)\, g(k) \\
&= e^{\lambda_1+\lambda_2} \sum_{k=-\infty}^{n} \mathbb{P}(W=k) \left[1\{k \leq x\} - \mathbb{P}(W \leq x)\right] \\
&= e^{\lambda_1+\lambda_2} \left[\mathbb{P}(W \leq \min\{n,x\}) - \mathbb{P}(W \leq x)\, \mathbb{P}(W \leq n)\right] \\
&= e^{\lambda_1+\lambda_2} \mathbb{P}(W \leq n) \left[1 - \mathbb{P}(W \leq x)\right] \\
&= e^{\lambda_1+\lambda_2} \mathbb{P}(W \leq n)\, \mathbb{P}(W > x).
\end{aligned}$$

The case that $n \geq x$ is similar. This means that

$$f_x(m) = \begin{cases} (-1)^m \left(\sqrt{\frac{\lambda_2}{\lambda_1}}\right)^m I_m \left[(-1)^c \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^c \frac{1}{I_c} f(c) \right. \\ \left. + \frac{e^{\lambda_1+\lambda_2}}{\sqrt{\lambda_1\lambda_2}} \sum_{n=c}^{m-1} \frac{(-1)^{n+1}}{I_n I_{n+1}} \mathbb{P}(W \leq \min\{n,x\})\, \mathbb{P}(W > \max\{n,x\}) \right] & \text{if } m > c \\ (-1)^m \left(\sqrt{\frac{\lambda_2}{\lambda_1}}\right)^m I_m \left[(-1)^c \left(\sqrt{\frac{\lambda_1}{\lambda_2}}\right)^c \frac{1}{I_c} f(c) \right. \\ \left. - \frac{e^{\lambda_1+\lambda_2}}{\sqrt{\lambda_1\lambda_2}} \sum_{n=m}^{c-1} \frac{(-1)^{n+1}}{I_n I_{n+1}} \mathbb{P}(W \leq \min\{n,x\})\, \mathbb{P}(W > \max\{n,x\}) \right] & \text{if } m < c. \end{cases}$$

■

A.4 Proof of Theorem 4

Our proof of Theorem 4 is highly involved, from an analysis perspective, but the overall program can be stated in a relatively succinct manner. Accordingly, we sketch here the overall program behind our proof and refer the interested reader to the online appendix for a detailed account.

Proof Recall that we are trying to obtain a bound on $|\Delta f_x(j)| = |f_x(j+1) - f_x(j)|$ independent of $x \in \mathbb{R}$ and $j \in \mathbb{Z}$. From Theorem 3, we have the solution to the Stein equation, however to use it to bound $|\Delta f_x(j)|$, we need to simplify it further. For ease of notation, we simply refer to f instead of f_x and g instead of g_x.

First, note that we have the freedom to choose the initial condition $(c, f(c))$. Making the choice that $c = \lambda_2 - \lambda_1$, and hence that $c = 0$ under the assumption that $\lambda_1 = \lambda_2$, we are able to simplify our expression for f in Theorem 3 to read, in the case that $m > 0$, as

$$\begin{aligned}
f(m+1) = (-1)^{m+1} I_{m+1} \Bigg[& \frac{1}{I_0} f(0) \\
& - \frac{e^{2\lambda}}{\lambda} \frac{1}{I_0 I_1} \mathbb{P}(W \leq \min\{0,x\})\, \mathbb{P}(W > \max\{0,x\}) \\
& - \frac{e^{2\lambda}}{\lambda} \sum_{n=0}^{m-1} \frac{(-1)^{n+1}}{I_{n+1} I_{n+2}} \mathbb{P}(W \leq \min\{n+1,x\})\, \mathbb{P}(W > \max\{n+1,x\}) \Bigg]
\end{aligned}$$

and, in the case that if $m < 0$, as

$$f(m-1) = (-1)^{m-1} I_{m-1} \Bigg[\frac{1}{I_0} f(0) - \frac{e^{2\lambda}}{\lambda} \frac{1}{I_0 I_1} \mathbb{P}(W \leq \min\{-1, x\}) \mathbb{P}(W > \max\{-1, x\}) + \frac{e^{2\lambda}}{\lambda} \sum_{n=m}^{-1} \frac{(-1)^{n+1}}{I_{n-1} I_n} \mathbb{P}(W \leq \min\{n-1, x\}) \mathbb{P}(W > \max\{n-1, x\}) \Bigg].$$

Finally, for the case $m = 0$, we have

$$\begin{aligned} f(0) &= \frac{e^{2\lambda}}{2\lambda} \frac{1}{I_0 + I_1} [\mathbb{P}(W \leq \min\{0, x\}) \mathbb{P}(W > \max\{0, x\}) \\ &\quad + \mathbb{P}(W \leq \min\{-1, x\}) \mathbb{P}(W > \max\{-1, x\})]. \end{aligned}$$

Next, through manipulation of the arguments in the sums defining the above expressions for f, exploiting properties of the modified Bessel functions I_k, and applying the triangle inequality, we are able to produce bounds on the differences $|f(m+1) - f(m)|$ of the form

$$\begin{aligned} |f(m+1) - f(m)| &\leq \frac{\mathbb{P}(W \leq x)}{\lambda} \Bigg\{ \left| \sum_{n=1,3,\ldots}^{m-1} \frac{I_{m+1}}{I_{n+2}} - \frac{I_m}{I_{n-1}} \right| \\ &\quad + \sum_{n=1,3,\ldots}^{m-1} H(n) \left| \frac{I_{m+1} I_n}{I_{n+1} I_{n+2}} - \frac{I_m}{I_{n+1}} - \frac{I_{m+1}}{I_{n+1}} + \frac{I_m}{I_{n-1}} \right| \Bigg\} \\ &\quad + \Bigg| I_m \frac{1}{I_0} f(0) + I_{m+1} \Bigg\{ \frac{1}{I_0} f(0) \\ &\quad - \frac{e^{2\lambda}}{\lambda} \frac{1}{I_0 I_1} \mathbb{P}(W \leq \min\{0, x\}) \mathbb{P}(W > \max\{0, x\}) \Bigg\} \Bigg| \quad , \end{aligned} \tag{19}$$

if m is even, and

$$\begin{aligned} |f(m+1) - f(m)| &\leq \frac{\mathbb{P}(W \leq x)}{\lambda} \Bigg\{ \left| \sum_{n=1,3,\ldots}^{m-1} \frac{I_{m+1}}{I_{n+2}} - \frac{I_m}{I_{n-1}} \right| \\ &\quad + \sum_{n=1,3,\ldots}^{m-1} H(n) \left| \frac{I_{m+1} I_n}{I_{n+1} I_{n+2}} - \frac{I_m}{I_{n+1}} - \frac{I_{m+1}}{I_{n+1}} + \frac{I_m}{I_{n-1}} \right| \\ &\quad + |H(m+1) - H(m)| \} \\ &\quad + \Bigg| I_m \frac{1}{I_0} f(0) + I_{m+1} \Bigg\{ \frac{1}{I_0} f(0) \\ &\quad - \frac{e^{2\lambda}}{\lambda} \frac{1}{I_0 I_1} \mathbb{P}(W \leq \min\{0, x\}) \mathbb{P}(W > \max\{0, x\}) \Bigg\} \Bigg| \quad , \end{aligned} \tag{20}$$

if m is odd. Here $H(n) = \mathbb{P}(W > n)/\mathbb{P}(W = n)$ is the inverse of the hazard function of the Skellam distribution (and is not to be confused with our use of H in the main body of the paper as a subgraph of the graph G).

Note that (19) is defined by three key terms, while (20) has the same three, augmented by the addition of a fourth, i.e., $|H(m+1) - H(m)|$. Through a series of arguments (the result for each of which is presented as a separate proposition in the online appendix), we are able to control each of these terms as follows. First, we show that

$$\sup_{m \in \mathbb{N}^+} \left| \sum_{n=1,3,\ldots}^{m-1} \frac{I_{m+1}}{I_{n+2}} - \frac{I_m}{I_{n-1}} \right| \leq 5 \ .$$

Next we show that

$$\sum_{n=1,3,\ldots}^{m-1} H(n) \left| \frac{I_{m+1} I_n}{I_{n+1} I_{n+2}} - \frac{I_m}{I_{n+1}} - \frac{I_{m+1}}{I_{n+1}} + \frac{I_m}{I_{n-1}} \right| \leq 73 \ ,$$

for $\lambda \geq 1$. And furthermore, we show that

$$\begin{aligned} & \left| I_m \frac{1}{I_0} f(0) + I_{m+1} \left\{ \frac{1}{I_0} f(0) - \frac{e^{2\lambda}}{\lambda} \frac{1}{I_0 I_1} \mathbb{P}(W \leq \min\{0, x\}) \, \mathbb{P}(W > \max\{0, x\}) \right\} \right| \\ \leq & \ \frac{\mathbb{P}(W \leq x)}{\lambda} \ . \end{aligned}$$

Finally, it is clear that

$$\begin{aligned} & H(m) - H(m+1) \\ = & \ \frac{\mathbb{P}(W > m)}{\mathbb{P}(W = m)} - \frac{\mathbb{P}(W > m+1)}{\mathbb{P}(W = m+1)} \\ = & \ \frac{1}{\mathbb{P}(W = m)} \left[\mathbb{P}(W > m) - \mathbb{P}(W > m+1) \right] \\ = & \ \frac{\mathbb{P}(W = m+1)}{\mathbb{P}(W = m)} \\ \leq & \ 1 \end{aligned}$$

and, at the same time $H(m) - H(m+1) \geq 0$ so, we have that we may bound the magnitude of this final term by 1.

As a result of all of the above, we may conclude that

$$|f(m+1) - f(m)| \leq \frac{80}{\lambda}$$

for $m > 0$. Or, equivalently, we may express the right-hand side above as $160/2\lambda$.

The argument for the case of $m < 0$ involves similar reasoning, as described in the online appendix. ■

A.5 Proof of Theorems 6 and 7

A.5.1 PROOF OF THEOREM 6

Proof The terms $\mathbb{E}\left|U - U_k^{(L)}\right|$ and $\mathbb{E}\left|U - U_k^{(M)}\right|$ in (8) measure the dependence of U on the events $L_k = 1$ and $M_k = 1$, respectively. In the context of the empirical graph $\hat{G}$, the random variables L are equal to Y_{ij}, for $\{i,j\} \in E^c$, while the random variables M are equal to Y_{ij}, for $\{i,j\} \in E$. With the Y_{ij} assumed independent, $U_k^{(L)}$ and $U_k^{(M)}$ are independent of their respective events, and so we obtain

$$d_{KS}(U,W) \leq ||\Delta f|| \left[\sum_{k=1}^{n} p_k^2 + \sum_{k=1}^{m} q_k^2\right] \quad . \tag{21}$$

Accordingly, and drawing on definitions and the result of Theorem 4,

$$\begin{aligned} d_{KS}\left(D_E, \text{Skellam}(\lambda,\lambda)\right) \leq \frac{1}{\lambda}\left[\sum_{(i,j)\in E^c} \alpha^2 + \sum_{(i,j)\in E} \beta^2\right] &= \frac{|E^c|\alpha^2 + |E|\beta^2}{|E^c|\alpha} \\ &= \frac{|E^c|\alpha^2 + |E|\left(\frac{|E^c|}{|E|}\right)^2\alpha^2}{|E^c|\alpha} \\ &= \alpha + \frac{|E^c|}{|E|}\alpha \\ &= \alpha + \frac{\binom{n_v}{2} - |E|}{|E|}\alpha \\ &= \frac{\binom{n_v}{2}}{|E|}\alpha \quad . \end{aligned}$$

Noting that $\alpha = \lambda/|E^c|$, and recalling that $|E^c| = \Theta\left(n_v^2\right)$ under both sparse and dense graphs G, the last quantity above is seen to behave like $\lambda/|E|$ which, under assumption (A3) and our definition of sparse and dense in Section 2 , reduces to $O\left(n_v^{-1}\right)$. So the bound in (9) is established.

Note that the right-hand side of (21) is analogous to the classical form of the bound for individual sums of independent indicator random variables (e.g., Barbour and Chen, 2005). As remarked in the main text, for this particular case of independent Y_{ij}, those more classical techniques could also be used to produce the result of Theorem 6. Specifically, Let $T_1, T_2, \tilde{T}_1$, and $\tilde{T}_2$ be independent random variables supported on the integers. Denote by $d_{TV}(X_1, X_2)$ the total-variation distance between two random variables X_1 and X_2. Then

$$\begin{aligned} d_{KS}\left(T_1 - T_2, \tilde{T}_1 - \tilde{T}_2\right) &\leq d_{TV}\left(T_1 - T_2, \tilde{T}_1 - \tilde{T}_2\right) \\ &\leq d_{TV}\left((T_1, T_2), (\tilde{T}_1, \tilde{T}_2)\right) \\ &\leq d_{TV}\left(T_1, \tilde{T}_1\right) + d_{TV}\left(T_2, \tilde{T}_2\right) \quad , \end{aligned}$$

where the first inequality exploits the fact that total-variation distance provides an upper bound on Kolmogorov-Smirnov distance, and the second and third inequalities follow from Lemmas 3.6.3

and 3.6.2 of Durrett (2010), respectively. Now define

$$T_1 = \sum_{\{i,j\}\in E^c} Y_{ij} \quad \text{and} \quad T_2 = \sum_{\{i,j\}\in E} (1 - Y_{ij}) \quad ,$$

and let $\tilde{T}_1$ and $\tilde{T}_2$ be independent Poisson random variables with common mean λ. Setting $\lambda = |E^c|\alpha = |E|\beta$, and applying to each of $d_{TV}\left(T_1, \tilde{T}_1\right)$ and $d_{TV}\left(T_2, \tilde{T}_2\right)$ the standard Stein bounds for Poisson approximation to sums of independent indicators (e.g., Barbour and Chen, 2005, Eqn 2.6), we again obtain that

$$d_{KS}\left(D_E, \text{Skellam}(\lambda,\lambda)\right) \leq \frac{1}{\lambda}\left[\sum_{(i,j)\in E^c} \alpha^2 + \sum_{(i,j)\in E} \beta^2\right] = \frac{|E^c|\alpha^2 + |E|\beta^2}{|E^c|\alpha} \quad ,$$

and the rest follows. ■

A.5.2 PROOF OF THEOREM 7

To establish the bounds in (10) and (11), we use the following result from Stein's method for the normal distribution (e.g., Barbour and Chen, 2005).

Theorem 11 *Let $\xi_1, \ldots, \xi_n$ be independent random variables which have zero means and finite variances $\mathbb{E}\left[\xi_i^2\right] = \sigma_i^2$, $1 \leq i \leq n$, and satisfy $\sum_{i=1}^n \sigma_i^2 = 1$. If $F_n(x)$ is the cumulative distribution function of $\sum_{i=1}^n \xi_i$, then, for every $\epsilon > 0$,*

$$\frac{1 - e^{-\frac{\epsilon^2}{4}}}{40}\sum_{i=1}^n \mathbb{E}\left[\xi_i^2 I_{\{|\xi_i|>\epsilon\}}\right] - \sum_{i=1}^n \sigma_i^4 \leq \sup_{x\in\mathbb{R}} |F_n(x) - \Phi(x)| \leq 7\sum_{i=1}^n \mathbb{E}\left[|\xi_i|^3\right].$$

Proof of Theorem 7 We apply Theorem 11, with $\xi_i = X_i/\sigma$ where X_i is a term in one of the sums of D_E, to establish each of our upper and lower bounds in turn.

Upper Bounds in (10) and (11): First, note that since

$$\sum_{i=1}^n \mathbb{E}\left[|\xi_i|^3\right] = \frac{\sum_{i=1}^n \mathbb{E}\left[|X_i|^3\right]}{\sigma^3} \quad ,$$

and

$$\mathbb{E}\left[|X_i|^3\right] = \alpha(1-\alpha)\left[(1-\alpha)^2 + \alpha^2\right] \quad \text{or} \quad \beta(1-\beta)\left[(1-\beta)^2 + \beta^2\right] \quad ,$$

with n understood to be either $|E^c|$ or $|E|$, it follows that

$$\begin{aligned}
\sum_{i=1}^n \mathbb{E}\left[|\xi_i|^3\right] &= \frac{\alpha(1-\alpha)\left[(1-\alpha)^2 + \alpha^2\right]|E^c| + \beta(1-\beta)\left[(1-\beta)^2 + \beta^2\right]|E|}{(\alpha(1-\alpha)|E^c| + \beta(1-\beta)|E|)^{\frac{3}{2}}} \\
&\leq \max\{(1-\alpha)^2 + \alpha^2, (1-\beta)^2 + \beta^2\}\frac{\alpha(1-\alpha)|E^c| + \beta(1-\beta)|E|}{(\alpha(1-\alpha)|E^c| + \beta(1-\beta)|E|)^{\frac{3}{2}}} \\
&= \frac{\max\{(1-\alpha)^2 + \alpha^2, (1-\beta)^2 + \beta^2\}}{(\alpha(1-\alpha)|E^c| + \beta(1-\beta)|E|)^{\frac{1}{2}}} \\
&= \frac{\max\{(1-\alpha)^2 + \alpha^2, (1-\beta)^2 + \beta^2\}}{\sqrt{2-(\alpha+\beta)}} \cdot \frac{1}{\sqrt{\alpha|E^c|}} \quad ,
\end{aligned}$$

where in the last equality we have used the fact that $\beta = (|E^c|/|E|)\,\alpha$ follows from (A2). Finally, note that

$$(1-\alpha)^2 + \alpha^2 = 1 - 2\alpha + 2\alpha^2 = 1 - 2\alpha(1-\alpha) \leq 1$$

and the same holds for $(1-\beta)^2 + \beta^2$, since $0 \leq \alpha, \beta \leq 1$, so that

$$\sum_{i=1}^{n} \mathbb{E}\left[|\xi_i|^3\right] \leq \frac{1}{\sqrt{2-(\alpha+\beta)}} \cdot \frac{1}{\sqrt{\alpha|E^c|}} \ .$$

This immediately implies, after another application of $\beta = (|E^c|/|E|)\alpha$,

$$\sup_{x\in\mathbb{R}} |F_n(x) - \Phi(x)| \leq \frac{7}{\sqrt{2-(\alpha+\frac{|E^c|}{|E|}\alpha)}} \cdot \frac{1}{\sqrt{\alpha|E^c|}} \ .$$

Using $\alpha = \lambda/|E^c|$, and invoking the assumption of low-rate measurement error in (A3) and the definitions of sparse and dense graphs in Section 2, the upper bounds in (10) and (11) follow.

Lower bound in (10) and (11): First, note that since $\xi_i = X_i/\sigma$, $\sigma_i^2 = \alpha(1-\alpha)/\sigma^2$ or $\sigma_i^2 = \beta(1-\beta)/\sigma^2$. Thus,

$$\begin{aligned}
\sum_{i=1}^{n} \sigma_i^4 &= \frac{(\alpha(1-\alpha))^2|E^c| + (\beta(1-\beta))^2|E|}{(\alpha(1-\alpha)|E^c| + \beta(1-\beta)|E|)^2} \\
&= \frac{(\alpha(1-\alpha))^2|E^c| + \left(\frac{|E^c|}{|E|}\alpha(1-\beta)\right)^2|E|}{(2-(\alpha+\beta))^2(\alpha|E^c|)^2} \\
&= \frac{1}{|E^c|} \cdot \frac{1}{(2-(\alpha+\beta))^2} \cdot \left[(1-\alpha)^2 + \frac{|E^c|}{|E|}(1-\beta)^2\right] \\
&= \frac{1}{(2-(\alpha+\beta))^2}\left[\frac{(1-\alpha)^2}{|E^c|} + \frac{(1-\beta)^2}{|E|}\right]
\end{aligned} \ ,$$

where in the second equality, we have used $\beta = (|E^c|/|E|)\alpha$.

Next, choose $\epsilon = 1/(2\sigma)$. Note that this is the midpoint of the intervals

$$\left(\frac{\alpha}{\sigma}, \frac{1-\alpha}{\sigma}\right), \quad \text{and} \quad \left(\frac{\beta}{\sigma}, \frac{1-\beta}{\sigma}\right)$$

if $\alpha, \beta < 1/2$ and of the intervals

$$\left(\frac{1-\alpha}{\sigma}, \frac{\alpha}{\sigma}\right), \quad \text{and} \quad \left(\frac{1-\beta}{\sigma}, \frac{\beta}{\sigma}\right) \ .$$

if $\alpha, \beta \geq 1/2$. In either case, these are the endpoints of the interval formed by the values of $|\xi_i| = |X_i|/\sigma$.

Due to the symmetry in these intervals about $\frac{1}{2}$ we may, without loss of generality, assume $\alpha, \beta < 1/2$. In doing so, and using $\beta = (|E^c|/|E|)\alpha$,

$$\begin{aligned}
&\frac{1-e^{-\frac{\epsilon^2}{4}}}{40}\sum_{i=1}^{n}\mathbb{E}\left[\xi_i^2 I_{\{|\xi_i|>\epsilon\}}\right] \\
&= \frac{1-e^{-\frac{\epsilon^2}{4}}}{40}\cdot\frac{(1-\alpha)^2\alpha|E^c| + (1-\beta)^2\beta|E|}{\alpha(1-\alpha)|E^c| + \beta(1-\beta)|E|} \\
&= \frac{1-e^{-\frac{\epsilon^2}{4}}}{40}\cdot\frac{(1-\alpha)^2 + (1-\beta)^2}{2-(\alpha+\beta)} \\
&= e^{-\frac{\epsilon^2}{4}}\frac{e^{\frac{\epsilon^2}{4}}-1}{40}\cdot\frac{(1-\alpha)^2 + (1-\beta)^2}{2-(\alpha+\beta)} \\
&\geq e^{-\frac{\epsilon^2}{4}}\frac{\epsilon^2}{160}\cdot\frac{(1-\alpha)^2 + (1-\beta)^2}{2-(\alpha+\beta)} \\
&= e^{-\frac{\epsilon^2}{4}}\frac{1}{640}\cdot\frac{1}{\alpha(1-\alpha)|E^c| + \beta(1-\beta)|E|}\cdot\frac{(1-\alpha)^2 + (1-\beta)^2}{2-(\alpha+\beta)} \\
&= e^{-\frac{\epsilon^2}{4}}\frac{1}{640}\cdot\frac{1}{\alpha|E^c|}\cdot\frac{(1-\alpha)^2 + (1-\beta)^2}{(2-(\alpha+\beta))^2} \quad .
\end{aligned}$$

Combining the two sets of expressions above, the lower bound becomes

$$\begin{aligned}
&\frac{1-e^{-\frac{\epsilon^2}{4}}}{40}\sum_{i=1}^{n}\mathbb{E}\left[\xi_i^2 I_{\{|\xi_i|>\epsilon\}}\right] - \sum_{i=1}^{n}\sigma_i^4 \\
\geq\ & e^{-\frac{\epsilon^2}{4}}\frac{1}{640}\cdot\frac{1}{\alpha|E^c|}\cdot\frac{(1-\alpha)^2 + (1-\beta)^2}{(2-(\alpha+\beta))^2} - \frac{1}{(2-(\alpha+\beta))^2}\left[\frac{(1-\alpha)^2}{|E^c|} + \frac{(1-\beta)^2}{|E|}\right] \\
=\ & \frac{1}{(2-(\alpha+\beta))^2}\left[(1-\alpha)^2\left(\frac{\exp\left(-\frac{1}{16}\cdot\frac{1}{\alpha|E^c|}\cdot\frac{1}{2-(\alpha+\beta)}\right)}{640}\cdot\frac{1}{\alpha|E^c|} - \frac{1}{|E^c|}\right)\right. \\
&\left. +(1-\beta)^2\left(\frac{\exp\left(-\frac{1}{16}\cdot\frac{1}{\alpha|E^c|}\cdot\frac{1}{2-(\alpha+\beta)}\right)}{640}\cdot\frac{1}{\alpha|E^c|} - \frac{1}{|E|}\right)\right] \quad . \qquad (22)
\end{aligned}$$

But for sufficiently large n_v, the exponential term in (22) behaves like $\exp[-1/(16\lambda)] \approx 1 - (1/16\lambda)$. Substituting accordingly and simplifying to ignore the various terms tending to a constant in large n_v, the expression in (22) can be seen to behave asymptotically like

$$\frac{1}{4}\left[\left(\frac{1}{640\lambda} - \frac{1}{|E^c|}\right) + \left(\frac{1}{640\lambda} - \frac{1}{|E|}\right)\right] \quad .$$

Again, by the assumption of low-rate measurement error in (A3) and the definitions of sparse and dense graphs given in Section 2, appropriate substitution of the values for λ, $|E|$, and $|E^c|$ yield the lower bounds in (10) and (11). This completes the proof of Theorem 7. ■

A.6 Proof of Theorem 8

Proof The proof follows by rewriting each of the two sums bracketed in (8), and then aggregating terms. Under condition (i) of the theorem,

$$\begin{aligned}\left|U - U_k^{(L)}\right| &= \left|L_k + \sum_{j\neq k} L_j - \sum_{\ell} M_\ell - \left(\sum_{j\neq k} \tilde{L}_j^{L_k} - \sum_{\ell} \tilde{M}_\ell^{L_k}\right)\right| \\ &= \sum_{j\neq k} \tilde{L}_j^{L_k} - \sum_{\ell} \tilde{M}_\ell^{L_k} - \left(\sum_{j\neq k} L_j - \sum_{\ell} M_\ell\right) - L_k \\ &= U_k^{(L)} - U \ .\end{aligned}$$

Similarly, under condition (ii) of the theorem, $\left|U - U_k^{(M)}\right| = U - U_k^{(M)}$.

In the absence of having to deal directly with the absolute values, we find that

$$\sum_{k=1}^{n} p_k E\left|U - U_k^{(L)}\right| = \sum_{k=1}^{n}\sum_{k\neq j} E\left[L_k L_j\right] - \sum_{k,\ell} E\left[L_k M_\ell\right] - E\left[U\right]\lambda_1$$

and

$$\sum_{k=1}^{m} q_k E\left|U - U_k^{(M)}\right| = \sum_{k=1}^{m}\sum_{\ell\neq k} E\left[M_k M_\ell\right] - \sum_{k,\ell} E\left[L_k M_\ell\right] + E\left[U\right]\lambda_2 \ .$$

As a result, the bracketed term in (8) takes the form

$$\begin{aligned}&\sum_{k=1}^{n} p_k E\left|U - U_k^{(L)}\right| + \sum_{k=1}^{m} q_k E\left|U - U_k^{(M)}\right| \\ =\ &\sum_{k=1}^{n}\sum_{k\neq j} E\left[L_k L_j\right] + \sum_{k=1}^{m}\sum_{\ell\neq k} E\left[M_k M_\ell\right] - 2\sum_{k,\ell} E\left[L_k M_\ell\right] - \left(E\left[U\right]\right)^2 \\ =\ &E\left[U^2\right] - \left(E\left[U\right]\right)^2 - \sum_{k=1}^{n} E\left[L_k^2\right] - \sum_{\ell=1}^{m} E\left[M_\ell^2\right] \\ =\ &\text{Var}(U) - \sum_{k=1}^{n} E\left[L_k\right] - \sum_{\ell=1}^{m} E\left[M_\ell\right] \\ =\ &\text{Var}(U) - (\lambda_1 + \lambda_2) \\ =\ &\text{Var}(U) - \text{Var}(W) \ .\end{aligned}$$

■

References

Milton Abramowitz and Irene A Stegun. *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*. Number 55. Courier Dover Publications, 1972.

Nesreen K Ahmed, Jennifer Neville, and Ramana Kompella. Network sampling: from static to streaming graphs. *ACM Transactions on Knowledge Discovery from Data*, 8:1–56, 2014.

Zack W. Almquist. Random errors in egocentric networks. *Social Networks*, 34(4):493–505, 2012.

Prakash Balachandran, Edoardo Airoldi, and Eric Kolaczyk. Inference of network summary statistics through network denoising. *arXiv preprint arXiv:1310.0423*, 2013a.

Prakash Balachandran, Weston D. Viles, and Eric D Kolaczyk. Exponential-type inequalities involving ratios of the modified Bessel function of the first kind and their applications. arxiv:1311.1450, 2013b.

Andrew D Barbour and Louis Hsiao Yun Chen. *An Introduction to Stein's Method*, volume 4. World Scientific, 2005.

T. Tony Cai, Tengyuan Liang, and Alexander Rakhlin. On detection and structural reconstruction of small-world random networks. *IEEE Transactions on Network Science and Engineering*, 2017.

Sourav Chatterjee. Matrix estimation by universal singular value thresholding. *Annals of Statistics*, 43(1):177–214, 2015.

Bill Cheswick, Hal Burch, and Steve Branigan. Mapping and visualizing the internet. In *USENIX Annual Technical Conference, General Track*, pages 1–12. Citeseer, 2000.

Elissa J. Cosgrove, Timothy S. Gardner, and Eric D. Kolaczyk. On the choice and number of microarrays for transcriptional regulatory network inference. *BMC bioinformatics*, 11(1):454, 2010.

Rick Durrett. *Probability: Theory and Examples*. Cambridge University Press, 2010.

Bradley Efron. *Large-scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*, volume 1. Cambridge University Press, 2012.

James D Esary, Frank Proschan, David W Walkup, et al. Association of random variables, with applications. *The Annals of Mathematical Statistics*, 38(5):1466–1474, 1967.

Christopher Fogelberg and Vasile Palade. Machine learning and genetic regulatory networks: a review and a roadmap. In *Foundations of Computational, Intelligence Volume 1*, pages 3–34. Springer, 2009.

Karl J Friston. Functional and effective connectivity: a review. *Brain Connectivity*, 1(1):13–36, 2011.

Wayne A Fuller. *Measurement Error Models*, volume 305. John Wiley & Sons, 2009.

Hyeouk Chris Hahm, Eric Kolaczyk, Jisun Jang, Theadora Swenson, and Asma Moiz Bhindarwala. Binge drinking trajectories from adolescence to young adulthood: the effects of peer social network. *Substance use & misuse*, 47(6):745–756, 2012.

Derek Hansen, Ben Shneiderman, and Marc A Smith. *Analyzing Social Media Networks with NodeXL: Insights from a Connected World*. Morgan Kaufmann, 2010.

G. Traver Hart, Arun K. Ramani, and Edward M. Marcotte. How complete are current yeast and human protein-interaction networks? *Genome Biology*, 7(11):1, 2006.

Shawndra Hill, Foster Provost, and Chris Volinsky. Network-based marketing: Identifying likely adopters via consumer networks. *Statistical Science*, pages 256–276, 2006.

Matthew O Jackson. *Social and Economic Networks*. Princeton University Press, 2010.

Svante Janson, Tomasz Luczak, and Andrzej Rucinski. *Random Graphs*, volume 45. John Wiley & Sons, 2011.

Xiaoyu Jiang and Eric D. Kolaczyk. A latent eigenprobit model with link uncertainty for prediction of protein–protein interactions. *Statistics in Biosciences*, 4(1):84–104, 2012.

Xiaoyu Jiang, David Gold, and Eric D. Kolaczyk. Network-based auto-probit modeling for protein function prediction. *Biometrics*, 67(3):958–966, 2011.

Joseph B Kadane. Sums of possibly associated Bernoulli variables: The Conway-Maxwell-binomial distribution. *arXiv preprint arXiv:1404.1856*, 2014.

Eric D Kolaczyk. *Statistical Analysis of Network Data*. Springer, 2009.

Dean Lusher, Johan Koskinen, and Garry Robins. *Exponential Random Graph Models for Social Networks: Theory, Methods, and Applications*. Cambridge University Press, 2012.

Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, pages 1436–1462, 2006.

Mark Newman. *Networks: an Introduction*. Oxford University Press, 2010.

Matey Neykov, Junwei Lu, and Han Liu. Combinatorial inference for graphical models. *arxiv preprint arXiv:1608.03045*, 2016.

Carey E. Priebe, Daniel L. Sussman, Minh Tang, and Joshua T. Vogelstein. Statistical inference on errorfully observed graphs. *Journal of Computational and Graphical Statistics*, 24(4):930–953, 2015.

John G Skellam. The frequency distribution of the difference between two Poisson variates belonging to different populations. *Journal of the Royal Statistical Society. Series A (General)*, 109(Pt 3):296, 1946.