

ECONOMETRICS OF SAMPLED NETWORKS

ARUN G. CHANDRASEKHAR[†]

RANDALL LEWIS[‡]

ABSTRACT. A growing literature studies social networks and their implications for economic outcomes. This paper highlights, examines, and addresses econometric problems that arise when a researcher studies these network effects using sampled network data. In applied work, researchers generally construct networks from data collected from a partial sample of nodes. For example, in a village, a researcher asks a random subset of households to report their social contacts. Treating this sampled network as the true network of interest, the researcher constructs statistics to describe the network or specific nodes and employs these statistics in regression or GMM analysis. This paper shows that even if nodes are selected randomly, partial sampling leads to non-classical measurement error and therefore bias in estimates of the regression coefficients or GMM parameters. The paper presents the first in-depth look at the impact of missing network data on the estimation of economic parameters. We provide analytical and numerical examples to illustrate the severity of the biases in common applications. We then develop two new strategies to correct such biases: a set of analytical corrections for commonly used network statistics and a two-step estimation procedure using graphical reconstruction. Graphical reconstruction uses the available (partial) network data to predict what the full network would have been and uses these predictions to mitigate the biases. We derive asymptotic theory that allows for each network in the data set to be generated by a different network formation model. Our analysis of the sampling problem as well as the proposed solutions are applied to rich network data of Banerjee, Chandrasekhar, Duflo, and Jackson (2011) from 43 villages in Karnataka, India.

JEL CLASSIFICATION CODES: C21, C49, C52, C83

KEYWORDS: Social networks, Sampling of networks, Measurement error

1. INTRODUCTION

A growing literature examines social networks and their implications for economic outcomes (see e.g., Jackson, 2008b, 2009a,b for an extensive survey of the literature). A network represents a set of connections (edges) among a collection of agents (nodes). For example, in a village network, nodes may represent households and edges may represent risk-sharing ties between households. Applied researchers typically construct a network from data that has been collected from a partial sample

Date: November 10, 2011 – Preliminary and Incomplete.

We thank Abhijit Banerjee, Victor Chernozhukov, Esther Duflo, Matthew O. Jackson, Anna Mikusheva, Whitney Newey, and Rob Townsend for encouragement and guidance. We are also grateful to Daron Acemoglu, Isaiah Andrews, Yann Bramouille, Emily Breza, Denis Chetverikov, Albert Fang, Paul Goldsmith-Pinkham, Ben Golub, Jerry Hausman, Nathan Hendren, Guido Imbens, Anton Kolotilin, Horacio Larreguy, Ben Olken, Adam Sacarny, Paul Schrimpf, and Jan Zilinsky. Essential feedback was provided by participants at the Calvó-Armengol Prize Workshop 2010, NEUDC 2010, and MIT workshops. Chandrasekhar is grateful for support from the NSF GRFP.

[†]Massachusetts Institute of Technology. Email: arunc@mit.edu.

[‡]Yahoo! Research. Email: rlewis@yahoo-inc.com.

of nodes rather than from all nodes in the network. Henceforth, such a network will be called the “sampled network”. This sampled network is naively treated as the true network of interest. The researcher uses a collection of sampled networks to estimate how network structure affects economic outcomes. This paper highlights, examines, and addresses econometric problems that arise when a researcher studies these network effects using sampled network data.

Concrete examples of network-based regressions in applied work include [Kremer and Miguel \(2007\)](#), who study the diffusion of deworming pill take-up, and [Hochberg et al. \(2007\)](#), who regress fund performance on measures of network importance of venture capital firms.¹ The applied work typically has low sampling rates (the share of nodes sampled), with a median of 25%, and 2/3 of the papers having a sampling rate below 51% (see Figure 3). Despite the prevalence of partial sampling, its implications for the estimation of economic parameters are rarely considered. One exception is [Conley and Udry \(2010\)](#) who study the diffusion of information among pineapple farmers in Ghana. Aware of the sampling problem, they conduct robustness exercises.

Our goal is to analyze the effect of using sampled network data on the estimation of parameters in network models of economic behavior. Henceforth, we call these the “economic parameters” without meaning to suggest that network formation is not economic.² In general, we are interested in parameters in a generalized method of moments (GMM) model, motivated by theory, describing the behavior of nodes in a network. The biases in estimates of economic parameters have not yet been systematically dealt with. While GMM is a general framework, two common classes of models allow us to explicitly characterize biases and are easier to work with due to their linearity: regressions of economic outcomes on network characteristics and regressions of a node’s outcomes on its network neighbors’ outcomes. After characterizing the biases, we propose two new strategies to correct such biases: a set of analytical corrections for commonly used network statistics and a two-step estimation procedure using graphical reconstruction that can be applied more broadly.

We focus on a running example throughout the paper: the diffusion of microfinance in 43 villages in rural Karnataka, India ([Banerjee et al., 2011](#)). A microfinance institution (MFI) based in Bengaluru expanded into these villages. Upon entering a village, the MFI informed certain households about its intentions and asked them to encourage others to join. The researcher wants to estimate how networks affect the diffusion of microfinance participation through these villages.

The present paper makes two core contributions. *Our first contribution is to highlight and analyze the biases in estimates of economic parameters when using sampled network data.* We develop analytical examples for commonly used network statistics, motivated by a number of applied questions concerning diffusion of information, social collateral, and risk-sharing. Next, we derive the corresponding biases that emerge when each of these statistics is used in regression. We show that the standard argument for attenuation due to classical measurement error does not apply;

¹There are numerous other examples. [Kinnan and Townsend \(2011\)](#) study whether whether households that are socially closer to credit sources smooth consumption better. [Leider et al. \(2009\)](#) and [Goeree et al. \(2010\)](#) study the effect of social proximity between pairs on the offers made in dictator games. [Alatas et al. \(2011\)](#) and examines whether networks with better diffusion properties actually induced greater information spreading. [De Giorgi et al. \(2010\)](#) study how network neighbors’ major choices affect a student’s own major choice.

²Parameters which describe the process by which networks are formed certainly are economic. We reserve “economic parameters” in our environment for parameters that describe a process that occurs on fixed networks.

coefficients may expand, attenuate, or switch signs depending on the network statistic of interest. In addition, we consider a model in which a node’s outcome depends on its peers’ outcomes and a node’s peer group is defined by the set of its social connections (Bramouille et al., 2009; De Giorgi et al., 2010). We show that the instrumentation technique used in the literature to overcome the reflection problem (Manski, 1993) in such models is invalid since the measurement error in the instrument will be correlated with the measurement error in the endogenous variable. Similarly, we consider GMM estimation of the Jackson and Rogers (2007b) model of diffusion and show that sampling the network induces expansion bias in the diffusion parameter. We supplement our analysis with numerical evidence for a wide array of examples to illustrate how sensitive econometric estimation is to the sampling of a network. In our numerical experiments, we estimate many models across a number of network statistics. At a sampling rate of $1/3$ we find that the estimates of the economic parameters have a mean absolute bias of 90% with a maximum of 260% for network-level regressions and a mean absolute bias of 63% with a maximum of 91% for node-level regressions.

Our second contribution is to develop two strategies to alleviate the biases: analytical corrections that apply to commonly used network statistics and two-step estimation using graphical reconstruction, which uses the observed part of the data to probabilistically reconstruct the missing part and then estimate the economic parameter accordingly.

First, by explicitly characterizing the biases, we derive simple bias corrections when the problem is tractable. We discuss several corrections and explore their reliability in addressing the biases. While computationally simple and easy to implement, these methods are typically limited to network-level regressions and are dependent on the particular network statistic of interest. Thus we develop a second, more general method that works well in practice – estimation by graphical reconstruction – to consistently estimate the economic parameter. This technique does not limit the researcher to network-level regressions nor to specific and tractable network statistics.

Consider the case where a researcher wants to perform a network-level regression of the rate of microfinance participation in a village on the average path length of the network.³ Without the entire network, the researcher falsely codes some existing links between individuals as if they do not exist. Graphical reconstruction builds on the simple idea that replacing every regressor for each village with a conditional expectation of the regressor delivers a consistent estimate of the regression coefficient. In our example, instead of using the mismeasured average path length of each network, the researcher ought to use the conditional expectation of the average path length, given the observed data. This requires integrating over all the missing data, given the observed information and sampling scheme, as opposed to treating missing links as if they did not exist. Furthermore, because different village networks likely formed in heterogenous manners, researchers ought to respect this heterogeneity in their analysis. By treating every network as an independent, but not identically distributed, random variable, we estimate the conditional expectation of the average path length in every network and consistently estimate the regression coefficient.

In practice, the researcher will have to estimate the distribution of missing links. We propose a two-step procedure. In the first stage, the researcher fits a potentially different model of network

³The average path length is the mean of all shortest paths between all pairs of nodes in a network.

formation to each network in the sample by making use of the observed data. Having done so, the researcher uses the network formation models to take draws of networks from their respective distributions, conditional on the observed information. Using these draws, the researcher estimates the conditional expectation of the regressor or moment in a GMM setting. In the second stage these conditional expectations are used in the usual way to estimate the economic parameter of interest. [Conley and Udry \(2010\)](#) perform a robustness exercise where they estimate missing neighborhood data in their regression model, which is an instance of graphical reconstruction.⁴

This two-step procedure is useful for several reasons. First, it allows the researcher to capture realistic heterogeneity by estimating a different model of network formation for every network. Second, our theoretical frame is general and we establish sufficient conditions ensuring that a desired class of network formation models can be used in graphical reconstruction. To build intuition, we make the analogy with panel data. Every network (individual in a panel) is independent, but the edges within a network (outcomes for an individual across time periods time periods) exhibit dependence. Under regularity conditions, a large network, similar to a large time series, may contain enough information such that the researcher can use the observed data to accurately estimate the distribution which generates the network formation process. The technical challenge that we overcome is to control an incidental parameter problem, where a parameter for every network must be estimated.⁵ Third, in our numerical experiments, it performs well. Even at a sampling rate of $1/3$, the median bias is 5.7% for network-level regressions and 1.4% for node-level regressions. The median reduction in bias is 62%. Each of the 96 estimated parameters exhibits reduction in bias when the reconstruction estimator is applied. Fourth, in addition to regression of economic outcomes on network statistics, the methodology can be applied to GMM models and even indexed GMM models where we may have a family of moment functions indexed by some parameter which presents technical challenges. Covering these cases is essential to network analysis as natural models, such as stopping time models for diffusion, may carry an index.

Given that estimation by graphical reconstruction requires integrating over missing links, this procedure demands more data than the analytical corrections. In addition to having a collection of sampled networks and outcome variables, we assume that the researcher has a covariates for each node (or pair of nodes) that will be predictive in the network formation models. Examples of such covariates include GPS coordinates, ethnicity, and caste, which are often readily available in development applications and are obtained during the listing process in each enumeration area. For a non-development example, consider school networks, where it is straightforward to obtain school rosters and demographic data for the entire collection of students.

To demonstrate another practical application of our results, we describe how researchers can employ our framework to make better decisions in collecting sampled network data, given their budget constraints. We provide an algorithm to assess the trade-off between the number of networks in a sample and the sampling rate a researcher uses. This exercise is similar in spirit to power

⁴The present paper develops a general theoretical framework, along with asymptotic analysis, that nests this strategy. We believe that estimates from graphical reconstruction ought to be used not only for robustness checks but also as estimates in their own right that exhibit substantially less bias.

⁵This is similar to non-linear fixed effects in panel settings ([Hahn and Kuersteiner, 2004](#); [Hahn and Newey, 2004](#)).

calculations frequently used in applied field work. First, the researcher obtains 100%-sampled network data for a small number of randomly chosen villages, using a pilot budget. Second, the researcher performs a numerical experiment by simulating outcome data from a specification that the researcher anticipates studying. In our microfinance example, the researcher simulates outcome data as a function of the path length from the initially informed households by assuming a regression coefficient and an R-squared. Third, the researcher draws, with replacement, a set of villages sampled at a given rate such that her budget is exhausted. By applying graphical reconstruction, the researcher can assess the mean-squared error minimizing choice of sampling rate.

We then apply our analysis of the sampling problem as well as the proposed solutions to sampled network data, collected in part by the authors, from 43 villages in Karnataka, India. [Banerjee et al. \(2011\)](#) study the diffusion of microfinance and, inspired by this analysis, we study natural specifications motivated by diffusion theory. We examine parameter estimates using the raw sampled data and compare them to those obtained by applying graphical reconstruction or analytical corrections. We find that applying our methods at times greatly changes parameter estimates and economic inferences. For instance, the impact of the network importance of initially informed households on the microfinance take-up rate in the village is under-estimated by 33% using the raw sampled network data when compared to using graphical reconstruction. In addition, a regression of a node's take-up decision on its neighbors' decisions shows that endogenous network effects may be severely under-estimated (with a 60% bias relative to the corrected estimate) or even switch signs (with a 166% bias relative to the corrected estimate). Moreover, regression coefficients in several specifications are not significantly different from zero at conventional levels when using the raw sampled data but are significantly different when applying the reconstruction estimator.

Related literatures across a number of fields including economics, epidemiology, statistics, sociology, and computer science have noted problems due to partial network data. The classical literature begins with [Granovetter \(1973\)](#), [Frank \(1980, 1981\)](#), and [Snijders \(1992\)](#) who identify how average degree and clustering are affected by several modes of random sampling. [Rothenberg \(1995\)](#) provides an excellent overview of the literature. More recently, the literature has focused on two classes of numerical experiments, typically with a single network.⁶ The first class documents biases that emerge when estimating parameters in a network formation model with partial data (e.g., in economics, [Santos and Barrett, 2008](#)). Second, the literature numerically describes behavior of certain network statistics under sampling (e.g., in epidemiology, [Ghani et al., 1998](#) and sociology, [Kossinets, 2006](#)). [Handcock and Gile \(2010\)](#) offer the straightforward solution to the first problem: by augmenting the likelihood to include the sampling scheme one can, in expectation, recover the correct network formation parameter.⁷ Finally, [Golub and Jackson \(2010b\)](#) explain how selection bias affects the estimation of a diffusion process and solve the puzzle raised by [Liben-Nowell and Kleinberg \(2008\)](#) as to why chain letter data exhibited very long average paths.

⁶[Santos and Barrett \(2008\)](#) also provide an extensive discussion of survey methodology and [Thompson \(2006\)](#) discusses sampling methodology and inferences on the degree distribution and network size.

⁷In simulations [Huisman \(2009\)](#) shows that for small amounts of missing data, an ad hoc method of imputing edges may reduce the bias in the network characteristics, though this is insufficient for larger amounts of missing data.

Our work builds on the above literature, with several key differences. First, the literature typically has not focused on nor developed a methodology to consistently estimate parameters from models of behavior on networks with sampled data. The substantive distinction here must be stressed. We are not interested in recovering the structural properties of the unobserved part of the network *per se*; instead, our goal is to understand the biases in estimation of these economic parameters and develop a method to recover them. Second, while augmented likelihood techniques for missing data are well-known in econometrics and statistics, we note that a collection of networks provides the researcher with a unique opportunity to set up the reconstruction problem in a manner which respects the substantial heterogeneity across networks. That is, a number of technical assumptions needed to control incidental parameter problems (e.g., nonlinear panel with fixed effects, [Hahn and Kuersteiner, 2004](#); [Hahn and Newey, 2004](#)) become very palatable in the network context, given that each network carries within it tremendous amounts of information. Consequently, graphical reconstruction focuses on conditional expectations of network regressors or moments to consistently estimate economic parameters when graphs are drawn from heterogeneous network formation models. This environment generates distinct technical challenges.

Finally, we note that partially observed networks constitute a special case of a broader class of network measurement problems.⁸ Our framework for analyzing the nature of biases in economic parameter estimates may provide intuition in these environments as well. If the researcher mis-specifies the network, biases will emerge. For example, if one is interested in economic behavior of individuals in a social network, but one defines the network based on social media data, biases will be induced as individuals may meaningfully interact with a subset of their social media neighbors. Another example is survey fatigue. Imagine that an individual forgets to name each of her connections independently with probability $1 - p$. Studying the OR network (where we denote two nodes as connected if either names the other) yields a network with only $1 - (1 - p)^2$ of the average number of friendship ties, while studying the AND network (where we denote two nodes as connected if both name each other) produces a network with only p^2 of the average number of friendship ties. A straightforward argument extends the results in section 3, with minor modification, to characterize the biases present in this example. Another common environment that induces bias is survey top-coding, where the survey limits the number of edges an individual can name.

The rest of the paper is organized as follows. Section 2 establishes the framework. The main results are in sections 3 and 4. Section 3 provides analytical examples of bias along with corrections. Section 4 discusses graphical reconstruction estimation. Section 5 contains numerical experiments which supplement sections 3 and 4. Section 6 applies the results to a study of the diffusion of microfinance. In section 7 we offer an algorithm for researchers to trade off the sampling rate against the number of networks. Section 8 concludes. All proofs are in the appendices.

2. FRAMEWORK

In this section we establish the framework. Section 2.1 introduces the notation, section 2.2 presents the econometric environment, and section 2.3 previews the asymptotic frame.

⁸Observed network data may miss information for a variety of reasons (see [Kossinets, 2006](#) for a discussion).




FIGURE 1. (a) $G = (V, E)$. (b) the sampled nodes, S , in red. (c) the induced subgraph, $G^{l^S} = (S, E^{l^S})$. (d) G , highlighting the sampled nodes and the edges that are induced if each sampled node reports all of its links from the census, E^S . (e) the star subgraph, G^S .

2.1. Notation and setup. A network or a *graph* is a pair $G = (V, E)$ consisting of a set V of *nodes* and a set E of *edges*, with $n := |V|$. Nodes i and j are either connected or unconnected (the graph is unweighted) and if i is connected to j , then j is connected to i (the graph is undirected). Most of what follows in this paper is applicable to directed and weighted graphs, though following the bulk of the applied research we restrict our attention to the undirected, unweighted case. A graph with n nodes is a member of the set of all undirected, unweighted graphs, denoted by \mathcal{G}_n .

A graph is represented by its *adjacency matrix*, $A := A(G)$. It is a matrix of 0s and 1s that depicts whether two nodes are connected, where $A_{ij} = \mathbf{1}\{ij \in E\}$ with the convention that $A_{ii} = 0$. We denote the *neighborhood* of i , the set of nodes it is connected to, by $N_i := \{j \in V : ij \in E\}$. Researchers are interested in economic models where an economic behavior or outcome is predicted by network statistics. We let $w(G)$ represent a d_w -dimensional vector of these network statistics. Since the data set may contain multiple networks, we use R to denote the number of graphs. The researcher is interested in *economic parameter* β_0 .

2.1.1. Sampling. Typically researchers obtain one of two types of sampled network data. First, the researcher may survey a set of m nodes and ask each node about the social connections with the other $m - 1$ nodes in that data set. We call this the *induced subgraph*, as it restricts the network among those who are sampled. Second, the researcher may have a list of the nodes in the network (e.g., a household census list in a village). From this list, a sample of m nodes may be surveyed. These nodes can name their social connections, not only to other $m - 1$ surveyed nodes, but connections to anyone from the list of n nodes. We call this the *star subgraph*.

Let S be the set of surveyed nodes, randomly chosen from V , with $m = |S|$. Let $m = \lfloor \psi n \rfloor$, where ψ is the *sampling rate*. The researcher obtains a subgraph of the graph in question. There are two potential resulting networks: the induced subgraph $G^{l^S} = (S, E^{l^S})$, which consists of the sampled nodes and the edges restricted to the set of surveyed nodes (E^{l^S}), and the star subgraph $G^S = (V, E^S)$, where E^S are edges such that at least one of the nodes is in S .

Figure 1 provides an illustration of the problem that this paper intends to address. Figure 1(a) displays G , the target network, Figure 1(c) shows the induced subgraph and Figure 1(e) depicts the star subgraph. We will also write $A = (A^{obs}, A^{mis})$ to denote the observed and missing part of the adjacency matrix, which are random variables under the sampling procedure. Although this framework idealizes the random sampling used in many applied contexts, our setting can easily be extended to other sampling methods such as independent edge sampling or snowball sampling.⁹

2.2. Econometric Models. The researcher intends to study economic behavior on R networks, $\{G_r : r = 1, \dots, R\}$. For simplicity, we assume every network has n nodes. An economic process has taken place on every network and can be described by an econometric model depending on an *economic parameter* β_0 . Returning to the microfinance example, information about microfinance has been introduced to certain households in every village and households decide to participate as the information propagates throughout the villages. Our goal is to estimate an economic parameter. We could easily do so if the networks were fully observed. The general framework for analyzing such models is to presume that a conditional moment restriction is satisfied,

$$(2.1) \quad \mathbb{E} [m(y, w(G); \beta_0) | G] = 0.$$

where $y \in \mathbb{R}^{d_y}$ is an outcome random variable, $m(\cdot, \cdot; \cdot)$ is a moment function, $w(\cdot)$ is function on \mathcal{G}_n , and $\beta \in \mathcal{B}$ is a parameter with true value β_0 .

Examples include discrete choice models, stopping time models (e.g., Iyer and Puri, forthcoming), quantile regression (e.g., Angelucci et al., 2010), and network-based matching models (e.g., Aral and Walker, Forthcoming; Banerjee et al., 2011). More generally, our results apply to indexed GMM models with parameter $\beta_0(u)$ where $u \in \mathcal{U}$ (e.g., time in a stopping time model or quantile in quantile regression). Partial sampling will generally generate biases as the moment will be a nonlinear function of the network statistic, so the estimated parameter will be inconsistent.

While GMM is a general framework, two common classes of econometric models with network data are easier to analyze due to their linearity. The first class consists of models wherein economic outcomes are regressed on network characteristics. The second class consists of models where a node's outcome depends on its network neighbors' outcomes.

Regression of Economic Outcomes on Network Characteristics. A researcher wants to study how network structure affects the economic outcome of interest, y , in regressions of the form¹⁰

$$(2.2) \quad y = \alpha + w(G) \beta_0 + \epsilon.$$

The researcher can estimate this regression at various observation levels. At the graph level, the data is $\{(y_r, w(G_r)) : r = 1, \dots, R\}$ where $w(G_r)$ is a d_w -vector of network statistics (e.g., average degree, clustering) and the regression contains R observations. In our example, the researcher may regress the microfinance take-up rate in a village on the average network importance of the

⁹Our graphical reconstruction solutions often apply to missing-at-random samples, where the probability of graph information being missing is independent of the missing data itself (Rubin, 1976).

¹⁰A vector of demographic covariates may be included, though we omit it for simplicity.

random set of households which were initially informed about microfinance. We anticipate that the centrality of these initial nodes correlates positively with take-up rates.

At the node level, the data is $\{(y_{ir}, w_{ir}(G_r)) : i = 1, \dots, n, r = 1, \dots, R\}$ where $w_{ir}(G_r)$ is a d_w -vector of statistics (e.g., degree of i , eigenvector centrality of i) and the regression has nR observations.¹¹ In our example, the researcher regresses a household's decision to join microfinance on its centrality. Theory suggests that central nodes will be more likely to learn new information. Similarly, one may estimate regressions at the edge level. Here $w_{ij}(G_r)$ is a d_w -vector of edge level statistics (e.g., social distance between the nodes) and the regression contains $\binom{n}{2} \cdot R$ observations.

Using sampled networks, the researcher runs regressions of the form

$$y = \alpha + w(\bar{G})\beta + u,$$

where \bar{G} is either G^{IS} or G^S , depending on the sampling scheme. In general, the measurement error will not be classical and may result in attenuation bias, expansion bias, or even sign switching. Sections 3.1 contain examples of common and economically meaningful network statistics where such biases exist and section 5 provides further numerical evidence on these biases.

Regression of Economic Outcomes on Network Neighbors' Outcomes. In a social equilibrium model, an economic outcome, y_i , depends on exogenous covariates of the individual, x_i , as well as the outcome of i 's peer group, $\{y_j : j \in N_i\}$. In our running example, y_i is the microfinance meeting attendance rate of a household and x_i represents whether the researcher has exogenously informed the household. Estimating such a model is difficult in the usual way (Manski, 1993), but with network data, assuming exogeneity of x_i as in the above examples, Bramoulle et al. (2009) and De Giorgi et al. (2010) show that the model may be identified as the peer groups for individuals are overlapping but not identical.

Formally, let $y = (y_1, \dots, y_n)'$ be the vector of outcome variables, $x = (x_1, \dots, x_n)'$ be the vector of exogenous covariates and $\iota = (1, \dots, 1)'$. A researcher is interested in estimating

$$(2.3) \quad y = \alpha_0 \iota + \rho_0 w(G)y + \gamma_0 x + \delta_0 w(G)x + \epsilon,$$

where $w(G)$ is a (possibly weighted) adjacency matrix that describes how much y_i is affected by others in the network. The economic parameter is $\beta_0 = (\rho_0, \gamma_0, \delta_0)$. Due to sampling, the researcher mistakenly estimates the model,

$$(2.4) \quad y = \alpha \iota + \rho w(\bar{G})y + \gamma x + \delta w(\bar{G})x + u,$$

where \bar{w} is defined analogously with \bar{G} either G^{IS} or G^S . The neighborhoods will be misspecified and the estimator exhibits bias. We discuss this model in Section 3.2.

2.3. Random Graphs and Asymptotic Framework.

Random Graphs. Until now we have described an economic process, such as diffusion, occurring on a given collection of networks. Consider the example of a regression of y on network covariate $w(G)$. With missing data the researcher does not observe the true network statistic. In section 3 we

¹¹With missing data, there will be $O(nR)$ observations. For instance with G^{IS} , one has $mR = \psi nR$ observations.

demonstrate the biases induced by using $w(\bar{G})$ where \bar{G} is the star or induced subgraph. Section 4 develops graphical reconstruction. We think of the network as the realization of a random network formation process. Consider a simple but commonly used model: the probability that individuals i and j are connected, conditional on covariate z_{ij} , is given by

$$P(A_{ij} = 1 | z_{ij}, \theta_0) = \Phi(z'_{ij}\theta_0),$$

where Φ is some link function. Thinking of the network as a random graph allows us to compute the conditional expectation of the regressor $w(G)$ given the observed portion of the network A^{obs} and the sampling scheme: $E[w(G)|A^{obs}; \theta_0]$. If we knew the distribution of G we could compute this expectation. By properties of conditional expectation using $E[w(G)|A^{obs}; \theta_0]$ as a regressor allows us to consistently estimate β_0 .

Formally, each network G_r is a random graph that is independently though not identically distributed over the space \mathcal{G}_{n_R} . We model the random networks as a triangular array of independent but not identically distributed random graphs, $G_{1,R}, \dots, G_{R,R}$. Each graph $G_{r,R}$ is a random draw from a distribution $P_{r,R}(G_r; \theta_{0r})$ over \mathcal{G}_{n_R} , where $\theta_{0r} \in \Theta_{r,R}$ is a parameter governing the distribution. In what follows, we omit the R subscript indexing the triangular array.

Asymptotic Frame. Graphical reconstruction requires estimating a conditional expectation for every network. Since the parameter θ_{0r} for each network is unknown we must be able to consistently estimate all of these together. Intuitively, we need conditions such that every network has enough information in it so that its parameter can be precisely estimated. This is similar to panel data with non-linear fixed effects, where both the number of individuals and the number of periods grow.

Formally, we will assume that $n_R \rightarrow \infty$ as $R \rightarrow \infty$. The rate requirements of n and R are discussed in Section 4. Moreover, $\Theta_{r,R}$ is typically finite dimensional, though we discuss an example where its dimension grows as $R \rightarrow \infty$. We assume conditions on n , R , and the random graph models such that every network asymptotically contains enough information to estimate θ_{0r} very well. In turn, we can estimate the conditional expectation very accurately and therefore recover the economic parameter β_0 .

Finally, we employ the following notation throughout the paper. $E[\cdot]$ denotes expectation, $\mathbb{E}_n[\cdot]$ the empirical expectation,¹² $\|\cdot\| = \|\cdot\|_2$ the ℓ^2 -norm, $\|\cdot\|_\infty$ the sup-norm, and $\ell^\infty(U)$ the space of bounded functions on \mathcal{U} . Also, $f_n \in \Theta(g_n)$ means $\exists k_1, k_2 > 0, n_0$ such that $\forall n > n_0$ $|g_n| \cdot k_1 \leq |f_n| \leq |g_n| \cdot k_2$, and the falling factorial is given by $(n)_j = n(n-1)\dots(n-j+1)$.

3. ANALYTICAL EXAMPLES OF BIAS

In this section we provide several analytical examples which demonstrate the biases due sampled network data. We provide three classes of examples: regression of economic outcomes on network statistics in section 3.1, regression of outcomes on network neighbors' outcomes in section 3.2, and a nonlinear GMM model of diffusion in section 3.3.

¹²For $a = (a_1, \dots, a_n)$, $\mathbb{E}_n[a_i] = \frac{1}{n} \sum_i a_i$. Similarly, $\mathbb{E}_R[a_r] = \frac{1}{R} \sum_r a_r$ and $\mathbb{E}_{n,R}[a_{i,r}] = \frac{1}{n} \sum_r \sum_i a_{i,r}$.

3.1. Regression of Economic Outcomes on Network Characteristics. To gain an intuition, we relate our problem to general measurement error. If the researcher is interested in a regression

$$y_r = w_r \beta_0 + \epsilon_r$$

but instead uses mismeasured regressors \bar{w}_r , the resulting estimator satisfies $\text{plim } \hat{\beta} = \beta_0 \frac{\text{cov}(\bar{w}, w)}{\text{var}(\bar{w})}$. Expansion, attenuation, and sign-switching bias are all possible. In our environment, $w_r = w(G_r)$, the relevant network statistic, but due to sampling the researcher uses $\bar{w}_r = w(\bar{G}_r)$, where \bar{G}_r is the star or induced subgraph. Consequently, we are primarily interested in the covariance of the network statistic with its true value, under the sampling scheme.

The covariance is typically not tractable to characterize. However, we partition our examples into two sets: those for which the mismeasurement has a scaling effect, in expectation and those that have a nonlinear effect. The scaling effect roughly means that $E[\bar{w}|w] = \pi w + o(1)$, where $\pi = \pi(\psi)$ is some known deterministic function. Thus,

$$\text{plim } \hat{\beta} = \beta_0 \cdot \underbrace{\pi^{-1}}_{\text{Scale}} \cdot \underbrace{\frac{\text{var}(w)}{\text{var}(w) + \text{var}(v)\pi^{-2}}}_{\text{Classical attenuation}}$$

where $v = \bar{w} - E[\bar{w}|w]$. There are two sources of biases: a scale effect which depends purely on $\pi(\psi)$ and a dispersion effect which generates attenuation. Average degree and graph clustering are commonly used network statistics that exhibit scale transformations. However, more general statistics such as path length and eigenvalues are not merely scaled in this manner. Through the remainder of the section we use $\hat{\beta}(\bar{G})$ to denote the estimator using the raw network data \bar{G} and $\tilde{\beta}(\bar{G})$ to denote the analytical correction, should one exist.

3.1.1. *Scaling transformations due to sampling.*

Average Degree. The degree of a node is its number of connections. Degree is a common measure of network importance. Kremer and Miguel (2007), Hochberg et al. (2007), Angelucci et al. (2009), Banerjee et al. (2011), Alatas et al. (2011), among others, use degree as a regressor. In addition to its ubiquitous use in applied work, average degree is a transparent, linear example which easily demonstrates the problem and has an intuitive (though not generalizable) solution.

We define $d(G_r) := \frac{1}{n} \sum_{i \in V} \sum_{j \in V} A_{ijr}$ as the average degree for graph r and use $d(\bar{G}_r)$ for $\bar{G} \in \{G^S, G^{|S}\}$ to denote the average degree computed with sampled data. As above, we are interested in the conditional expectation under a sampling scheme. One can show that

$$E[d(G_r^S)|G_r] = (1 - (1 - \psi)^2 + o(1)) d(G_r) \text{ and } E[d(G_r^{|S})|G_r] = (\psi + o(1)) d(G_r).$$

Intuitively, a node will miss links in proportion to the number of links it has. For instance, in the star subgraph an edge only appears in the data if at least one of the two nodes are sampled, yielding $1 - (1 - \psi)^2$. Notice that similar computations can be performed for other random sampling schemes the researcher may face. We define regressors to use in the analytical correction

$$\tilde{w}(G_r^S) := m^{-1} \sum_{i \in S} \sum_{j \in V} A_{ijr} \text{ and } \tilde{w}(G_r^{|S}) := \psi^{-1} d(G_r^{|S}).$$

The former involves simply constructing the average degree among the randomly sampled nodes, for which there is no missing data,¹³ and the latter simply accounts for the scaling effect. Notice that the former solution is highly non-generic; while it works for degree, it will work for nearly no other network statistic since other network characteristics will depend on other links as well. In addition, such a solution typically will not work for other sampling schemes which may not leave a portion of the neighborhoods intact. We define the limiting cross-network variance in the degree and the disturbance terms,

$$\sigma_{\langle d \rangle}^2 := \text{plim}_{R \rightarrow \infty} a_R^{-2} \mathbb{E}_R [d(G_r) - \mathbb{E}_R d(G_r)]^2 \text{ and } \sigma_{v_j}^2 := a_R^{-2} \text{plim}_{n \rightarrow \infty} \mathbb{E}_R v_{rj}^2 \text{ for } j \in \{S, |S\}$$

where $v_{rj} := d(\bar{G}_r) - \mathbb{E}[d(\bar{G}_r)|G_r]$ and (a_R) is some sequence of normalizing constants.

PROPOSITION 3.1. *Assume that $m = \lfloor \psi n \rfloor$ nodes are uniformly randomly sampled from each graph and the data $(y_{rR}, d(G_{rR}))_{r \leq R}$ is a triangular array satisfying the regularity conditions of Assumption A.2. Then,*

$$\begin{aligned} (1) \quad & \widehat{\beta}(G^{|S|}) \xrightarrow{P} \frac{\beta_0}{\psi} \cdot \frac{\sigma_{\langle d \rangle}^2}{\sigma_{\langle d \rangle}^2 + \psi^{-2} \sigma_{v|S}^2} \text{ and } \widehat{\beta}(G^S) \xrightarrow{P} \frac{\beta_0}{1 - (1 - \psi)^2} \cdot \frac{\sigma_{\langle d \rangle}^2}{\sigma_{\langle d \rangle}^2 + (\psi(2 - \psi))^{-2} \sigma_{vS}^2}. \\ (2) \quad & \widetilde{\beta}(G^S) \xrightarrow{P} \beta_0 \text{ and if } \sigma_{v|S}^2 \rightarrow 0, \widetilde{\beta}(G^{|S|}) \xrightarrow{P} \beta_0. \end{aligned} \quad ^{14}$$

In expectation, the average degree will be scaled down as a function of the sampling rate, since only a share of social connections are observed. Because the regressors are scaled down the coefficient expands, while dispersion around this expectation induces attenuation. The behavior of the analytical corrections is straightforward. The case with G^S is clear so we discuss $G^{|S|}$. The analytical correction directly removes the rescaling bias, leaving only the attenuation term involving $\sigma_{v|S}^2$.

Graph Clustering. The clustering coefficient of a node is the fraction of its neighbors that are themselves connected to each other. Bloch et al. (2008) study risk-sharing on social networks and find that networks which have theoretical properties that lend themselves to higher levels of risk-sharing tend to have higher clustering, which motivates clustering as a regressor with risk-sharing data.¹⁵ Moreover, under certain assumptions, Möbius and Szeidl (2006) and Karlan et al. (2009) use a model of trust and social collateral to microfound clustering as a measure of social capital. The clustering of the graph (denoted $c(G)$) is the ratio of the number of triangles (three nodes where every node is connected to every other node, denoted $\rho(G)$) to the number of connected triples (three nodes that have at least two social connections between them, denoted $\tau(G)$). That is, the *graph clustering* is $c(G) := \frac{\rho(G)}{\tau(G)}$, where

$$\rho(G) := 3 \sum_i \sum_{k>i} \sum_{j \neq i,k} A_{ij} A_{jk} A_{ki} \text{ and } \tau(G) := \sum_i \sum_{k>i} \sum_{j \neq i,k} A_{ij} A_{jk}.$$

The analysis of graph clustering is similar to degree. One can show that

$$\mathbb{E} [\rho(G_r^S) | G_r] = (3\psi^2(1-\psi) + \psi^3 + o(1)) \cdot \rho(G_r) \text{ and } \mathbb{E} [\tau(G_r^S) | G_r] = (\psi(1-\psi)^2 + 3\psi^2(1-\psi) + \psi^3 + o(1)) \cdot \tau(G_r).$$

¹³One could also simply account for the scaling effect.

¹⁴In Appendix D we demonstrate how this can be done without assuming $\sigma_{v|S}^2 \rightarrow 0$ by instead estimating $\widehat{\sigma}_{v|S}^2$.

¹⁵One may also be interested in the *support* of a graph defined in Jackson et al. (Forthcoming). The analysis of the bias in using support as a regressor will follow in a similar manner

This comes from the fact that to obtain a triangle one needs to sample at least two of the nodes, while to obtain a connected triple, one can additionally sample the middle node. Meanwhile, for the induced subgraph we have $E[\rho(G_r^{|S})|G_r] = (\psi^3 + o(1)) \cdot \rho(G_r)$ and $E[\tau(G_r^{|S})|G_r] = (\psi^3 + o(1)) \cdot \tau(G_r)$ since all three nodes must be sampled.

We assume a sequence of graphs which have non-zero clustering such that the sequence also has non-zero clustering in the limit. We define

$$v_{rj} := c(\bar{G}_r) - E[\rho(\bar{G}_r)|G_r]/E[\tau(\bar{G}_r)|G_r] \text{ for } j \text{ and } \bar{G} \text{ indexing the sampling scheme.}$$

In turn, $\sigma_{\langle c \rangle}^2$, $\sigma_{v|S}^2$, and $\sigma_{v|S}^2$ are defined analogously as in section 3.1.1. For adjusted least squares we use $\tilde{w}(G^S) = \left(\frac{\psi(3-2\psi)}{1+\psi(1-\psi)}\right)^{-1} c(G^S)$ and $\tilde{w}(G^{|S}) \equiv c(G^{|S})$. The analytical correction is identical to the sampled statistic in the case of the induced subgraph as the graph clustering of the induced subgraph consistently estimates the underlying graph clustering.

PROPOSITION 3.2. *Assume that $m = \lfloor \psi n \rfloor$ nodes are uniformly randomly sampled from each graph, the data $(y_{rR}, c(G_{rR}))_{r \leq R}$ is a triangular array from a deterministic sequence of graphs satisfying $\sup_{R \geq 1} \sup_{r \leq R} a_R \rho_r \vee a_R \tau_r \leq \bar{\alpha}$, $\inf_{R \geq 1} \inf_{r \leq R} a_R \rho_r \vee a_R \tau_r \geq \underline{\alpha}$, and the regularity conditions of Assumption A.2. Then*

$$(1) \quad \hat{\beta}(G^{|S}) \xrightarrow{P} \beta_0 \cdot \frac{\sigma_{\langle c \rangle}^2}{\sigma_{\langle c \rangle}^2 + \sigma_{v|S}^2} \text{ and } \hat{\beta}(G^S) \xrightarrow{P} \left(\frac{\psi(3-2\psi)}{1+\psi(1-\psi)}\right)^{-1} \beta_0 \cdot \frac{\sigma_{\langle c \rangle}^2}{\sigma_{\langle c \rangle}^2 + \left(\frac{\psi(3-2\psi)}{1+\psi(1-\psi)}\right)^{-2} \sigma_{v|S}^2},$$

$$(2) \quad \text{if } \sigma_{v|S}^2, \sigma_{v|S}^2 \rightarrow 0 \text{ then } \tilde{\beta}(G^{|S}) \xrightarrow{P} \beta_0 \text{ and } \tilde{\beta}(G^S) \xrightarrow{P} \beta_0.$$

As random sampling yields on average the same share of connected edges between each possible triangle, $c(G^{|S})$ consistently estimates $c(G)$. Dispersion about the mean yields the attenuation bias in regression. For G^S , there is an expansion bias owing to the fact that a triangle appears as an intransitive triad (where exactly two of the three nodes are connected) when an edge of the triangle is missing due to sampling. The usual attenuation term also applies. A perhaps surprising consequence is that using the unbiased estimator of clustering, $c(G^{|S})$, in a regression may perform worse than using $\tilde{w}(G^S)$.

3.1.2. Nonlinear transformations due to sampling.

Average Path Length and Graph Span. The path length between two nodes i and j is given by the minimum number of steps taken on the graph to get from i to j , denoted $\gamma(i, j) := \min_{l \in \mathbb{N} \cup \infty} [A^k]_{ij} > 0$. If there is no such finite path, we put $\gamma(i, j) = \infty$. The average path length of a graph is the mean taken over all finite paths,

$$\gamma(G) := \sum_{i, j: \gamma(i, j) < \infty} \gamma(i, j) / \left| \{(i, j) \in V^2 : \gamma(i, j) < \infty\} \right|.$$

Models of diffusion of information, flows of finance, risk-sharing, nepotism, and other phenomena, build on the principle that the farther apart agents are, the less is transmitted between them. For example, [Kinnan and Townsend \(2011\)](#) study how the network distance to a bank affects consumption smoothing. Other papers that use path length or average path length include [Golub and Jackson \(2010\)](#) who simulate diffusion processes; [Leider et al. \(2009\)](#) and [Goeree et al. \(2010\)](#)

who study dictator games between members of a school; [Alatas et al. \(2011\)](#) who look at the diffusion of information about poverty; and [Banerjee et al. \(2011\)](#) who study the diffusion of microfinance.

The average path length of a network is known to be a very difficult object to study analytically.¹⁶ Both the economics and statistical physics literatures study an object we term the *graph span*, which mimics average path length. [Jackson \(2008a\)](#) shows that for a general family of random graph models the ratio of the graph span to average path length asymptotically almost surely is one. The statistical physics literature uses such an approximation as well (e.g., [Newman et al., 2001](#); [Watts and Strogatz, 1998](#), [Watts and Strogatz, 1998](#)). These literatures motivate the study of the graph span as a regressor. Let $d_2(G) := \frac{1}{n} \sum_{i=1}^n \sum_{j>i}^n \sum_{k \neq i,j} A_{ij} A_{jk}$ be the average number of second neighbors.¹⁷ The *graph span* is

$$\ell(G) := \frac{\log n - \log d(G)}{\log d_2(G) - \log d(G)} + 1.$$

Larger networks have higher spans. Networks that are more expansive in the sense that the number of second neighbors far exceeds the number of neighbors have lower spans; it takes fewer steps to walk across the network. We first describe some useful properties of average degree and average number second neighbors for randomly sampled graphs and then study regression bias.

LEMMA 3.1. *Put $k(\psi) = \psi + \psi^2 - \psi^3$. For any sequence of random graphs $(G_n)_{n \in \mathbb{N}}$ satisfying, as $n \rightarrow \infty$, $d(G)/a_{1n} \xrightarrow{P} c_1$, $d_2(G)/a_{2n} \xrightarrow{P} c_2$, $a_{1n}, a_{2n} \in o(n)$, and $c_1, c_2 > 0$,*

- (1) $|d(G^{I^S}) - \psi d(G)| = o_P(1)$, $|d_2(G^{I^S}) - \psi^2 d_2(G)| = o_P(1)$,
- (2) $|d(G^S) - (1 - (1 - \psi)^2)d(G)| = o_P(1)$, $|d_2(G^S) - k(\psi)d_2(G)| = o_P(1)$.

This observation is general in the sense that it only requires that degree and the number of second neighbors to grow sufficiently slowly, which is reasonable for realistic applications. We now study the behavior of regressions with $\ell(G^{I^S})$ and $\ell(G^S)$. If a sufficiently high fraction of nodes are sampled, the estimator exhibits attenuation. Meanwhile, if a sufficiently low fraction of nodes are sampled, the estimator may exhibit sign-switching.

Put $\tilde{d}(\tilde{G})$ as the corrected estimate of degree from section 3.1.1, $\tilde{d}_2(G^S) = d_2(G^S)/k(\psi)$, and $\tilde{d}_2(G^{I^S}) = d_2(G^{I^S})/\psi^2$. Define

$$\tilde{\ell}(G^S) := \frac{\log n - \log \tilde{d}(G^S)}{\log \tilde{d}_2(G^S) - \log \tilde{d}(G^S)} + 1 \text{ and } \tilde{\ell}(G^{I^S}) := \frac{\log(\psi^{-1}m) - \log \tilde{d}(G^{I^S})}{\log \tilde{d}_2(G^{I^S}) - \log \tilde{d}(G^{I^S})} + 1,$$

which we use in the adjusted least squares estimators.

PROPOSITION 3.3. *Assume that $m = \lfloor \psi n \rfloor$ nodes are uniformly randomly sampled from each graph and the data sequence satisfies the regularity conditions of Assumption A.3. Then*

- (1) $\hat{\beta}$ is sign-consistent with attenuation if $\psi \in (c_1, 1)$ or $k(\psi)/(1 - (1 - \psi)^2) \in (c_1, 1)$

$$\text{plim } |\hat{\beta}(G^{I^S})| < |\beta_0| \text{ and } \text{plim } |\hat{\beta}(G^S)| < |\beta_0|,$$

¹⁶Bollobas (2001) approaches path length from an exact analytical perspective but only for a very specific random graph family. This approach is not suitable for gaining intuition for broader classes of graphs.

¹⁷Notice this defines second neighbor in the sense of taking a random node and then counting the number of neighbors of each of the node's neighbors. The definition is different from counting the number of distinct nodes at path length two from a given node, which would be $\frac{1}{n} \sum_i \sum_{k>i} \sum_{j \neq i,k} A_{ij} A_{jk} (1 - A_{ik})$.

- (2) $\hat{\beta}$ may be sign-inconsistent otherwise, and
 (3) analytical corrections are consistent, $\text{plim} \tilde{\beta}(G^{l^S}) = \beta_0$ and $\text{plim} \tilde{\beta}(G^S) = \beta_0$,

where $\xi_{rR} := d(G_{rR})/d_2(G_{rR})$, $0 < c_0 < c_1 < 1$, $c_0 := \inf_{R \geq 1} \inf_{r \leq R} \xi_{rR}$, and $c_1 := \sup_{R \geq 1} \sup_{r \leq R} \xi_{rR}$.

Sampling a network thins out the set of edges, resulting in a higher graph span. As the graph span approximates behavior of average path length, it captures the intuition that due to sampling, paths on graphs seem longer than they truly are. The expansion of the graph span has a slope effect on $\hat{\beta}$, and as $\log \psi < 0$ and $\log(k(\psi)/(1 - (1 - \psi)^2)) < 0$, the effect is either attenuation unless the sampling rate is too low, in which case sign-switching becomes a possibility. One must proceed with caution when discussing cases where the sampling probability is too low. In this case the network can shatter, yielding “islands” of disconnected sets of nodes which have short average path lengths within the set but have infinite distance across the sets.¹⁸ Since average path length is defined as a mean conditional on all finite paths, this is precisely where sign-switching may occur in practice. Alatas et al. (2011) contains an example where this happens in Indonesian networks.

Spectral Functions. Spectral functions are network statistics that relate to the set of eigenvalues of matrices which represent the graph, such as the adjacency matrix. They are useful in characterizing properties of the network. The distribution of eigenvalues has applications to models of information diffusion and risk-sharing as well. The number of k -length walks that cycle back to the original node correspond to k -th moment of the eigenvalue distribution, denoted $\mu^k(G)$,

$$\mu^k(G) = n^{-1} \sum_{i_1, \dots, i_k \in V^k} A_{i_1 i_2} \dots A_{i_k i_1} = n^{-1} \text{Tr}(A^k)$$

where $V^k = V \times \dots \times V$ (Barabasi and Albert, 1999). Given that the graph spectrum carries a great deal of information about the diffusive properties of a network, it is a useful regressor.¹⁹

There are several applications of spectral statistics in economic theory. For instance, the first eigenvalue of the adjacency matrix, $\lambda_{\max}(G)$, describes how well the graph diffuses information (e.g., Bollobás et al., 2010).²⁰ In models of social learning Golub and Jackson (2009, 2010a) show that the second eigenvalue of a weighted adjacency matrix is related to the time it takes to reach consensus; similar results are shown in DeMarzo et al. (2003). Ambrus et al. (2010) also characterize the risk-sharing capacity of a network as a function of the expansiveness of the network; it is well-known in network theory that this maps into the eigenvalues of a transformation of the adjacency matrix (Chung, 1997). It is difficult to precisely characterize the behavior of these spectral regressors, though we present bounds on their behavior under sampling.

PROPOSITION 3.4. *For an arbitrary graph G , we have*

- (1) $\mu^k(G^S) < \mu^k(G)$ and $E[\mu^k(G^{l^S})|G] = \sum_{j=2}^k \frac{(m-1)_j}{(n-1)_j} \eta_j < \mu^k(G)$, where η_j is the number of sets of j -distinct nodes that are counted.

¹⁸One can check that a graph H with $d_2(H)/d(H) < 1$ cannot be connected. The sign-switching case requires at least some $d_2/d < 1$ which we note the researcher can immediately detect.

¹⁹For a discussion of how spectral distributions correlate with graphical properties, see Barabasi and Albert (2002).

²⁰In a percolation process the threshold probability above which a giant component emerges is precisely $1/\lambda_{\max}$. For another intuition, if A is diagonalizable, then the dominant factor in $\|A^k\|$ is λ_{\max}^k .

$$(2) \lambda_{\max}(G^{|S}) \leq \lambda_{\max}(G^S) \leq \lambda_{\max}(G).$$

Since λ_{\max} can be thought of as measuring the number of walks through the graph (and with missing edges there are fewer walks) we expect expansion bias in $\hat{\beta}$ when using these regressors.²¹ This means that networks will appear to be more diffusive than they actually are.

3.2. Regression of Outcomes on Network Neighbors' Outcomes. We discuss the impact of sampled networks on regressions of nodes' outcomes on network neighbors' outcomes. The models we consider are developed in Bramouille et al. (2009) and De Giorgi et al. (2010) and naturally extend the models discussed in Manski (1993) to a network setting. Blume et al. (2011) contains an extensive review of the literature. The network allows for nodes to have overlapping but not identical peer groups.

The model is given by (2.3) and we are interested in $\beta_0 = (\rho_0, \gamma_0, \delta_0)$. There are two natural examples for how neighbors' outcomes ought to affect a node's outcome. First, every node's outcome may be affected by the average outcome of its neighbors.²² Second, every node's outcome may be affected by the total sum of its neighbor's outcomes.²³ The reduced form is

$$y = \alpha\iota/(1 - \rho_0) + \gamma_0x + (\gamma_0\rho_0 + \delta_0) \sum_{k=0}^{\infty} \rho_0^k w^{k+1}x + \sum_{k=0}^{\infty} \rho_0^k w^{k+1}\epsilon.$$

Since a node's neighborhood outcome, wy , is the endogenous regressor, the reduced form suggests that extended neighborhood effects, powers w^kx ($k \geq 2$), can be used as instruments for wy . We focus on the instrument $Z = [\iota, x, wx, w^2x]$.²⁴ Setting $X = [\iota, wy, x, wx]$, the estimator is $(X'P_ZX)^{-1}X'P_Zy$.

Identification comes from intransitive triads.²⁵ If i and j are connected and j and k are connected, but i and k are not connected, then k affects i only through j . As such, x_k is used as an instrument for y_j 's effect on y_i . We caution that this identification strategy convincingly works only when x is randomly assigned (e.g., Ngatia, 2011; Dupas 2010) as identification crucially depends on exogeneity of x .

We examine the estimation of (2.4) using $\bar{w} = w(G^S)$ or $\bar{w} = w(G^{|S})$ with instrument $Z_{\bar{G}} = [\iota, x, \bar{w}x, \bar{w}^2x]$. We show that the exclusion restriction is invalid when using sampled network data, even if the covariates are exogenous and the usual identification requirements are met if the full network data was available.²⁶

PROPOSITION 3.5. *Assume $\gamma_0\rho_0 + \delta_0 \neq 0$ and $w^2 \neq 0$, so 2SLS is valid for (2.3). Then 2SLS with*

²¹Whether there is expansion bias depends on how the eigenvalues shrink across the initial distribution. For instance, if the contraction is by translation, the regression slope would clearly not change. Numerical with simulated networks and empirical data provide evidence of expansion bias.

²²We can write the model as $y_i = \alpha + \beta\mathbb{E}_{N_i}[y_j] + \gamma x_i + \delta\mathbb{E}_{N_i}[x_j] + \epsilon_i$ as $\mathbb{E}_{N_i}[y_j] = \sum_{j \in N_i} y_j/d_i = \sum_j y_j A_{ij}/d_i$.

²³We discuss the first case, though clearly by mimicking the argument the results follow for the second.

²⁴Other estimation strategies are suggested in the literature, on the basis of efficiency (Bramouille et al., 2009; Lee et al., 2009). They require the validity of the instrument Z .

²⁵Bramouille et al. (2009) provide formal identification conditions.

²⁶De Giorgi et al. (2010) are aware that measurement error may cause problems in this model and conduct a numerical robustness exercise.

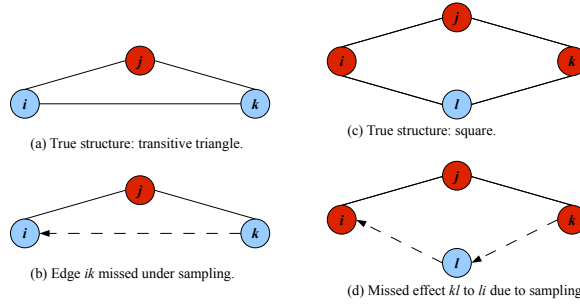


FIGURE 2. Red nodes are sampled. (a) and (c) show examples of true network structures. (b) shows the star subgraph from (a) when j is sampled. The effect of k on i is missed. (d) shows the induced subgraph from (c) when i , j , and k are sampled. The effect of k on i through l is missed.

- (1) $w(G^S)$ generically yields $E[Z_{G^S}u_{G^S}] \neq 0$,²⁷
- (2) $w(G^{lS})$ generically yields $E[Z_{G^{lS}}u_{G^{lS}}] \neq 0$,
- (3) an analytical correction, using $w(G^S)$ while restricting the second stage to $i \in S$, yields $E[Z_{G^S}u_{G^S}] = 0$.

Sampling induces an errors-in-variables problem, wherein the neighborhood effect is mismeasured since the neighborhoods themselves are misspecified. Though typically one uses instruments to address such a problem, here the instrument is correlated with the measurement error in the regressor, as the instrument involves powers of the mismeasured adjacency matrix. As such, the exclusion restriction is violated.

Figure 2 provides two examples where invalid instrumentation is generated. Figure 2(a,b) show that if j is sampled but i and k are not, the sampled network falsely suggests that k is a valid instrument for j 's effect on i . Similarly, figure 2(c,d) show a case with the induced subgraph, where k instrumenting for j 's effect on i will be invalid as the other channels through which k affects i are not accounted for due to sampling. In this case, the channel through l is omitted.

With G^S data, however, we propose a simple analytical correction. For $i \in S$, notice that $[\bar{w}x]_i = [wx]_i$ and $[\bar{w}y]_i = [wy]_i$. Consequently, there is no measurement error in the second stage for these observations. As only the first stage contains measurement error, uncorrelated with the second stage residual, such an exercise satisfies the exclusion restriction.

Finally, we note that the results presented in this section also have implications for the opposite case, wherein the researcher incorrectly adds edges to the network architecture and then estimates peer effects using modified network data. For instance, if the researcher assigns friendship links among students who share classes, when in fact the peer effect channel is through friendship links only, then for analogous reasons the estimation will be biased.

3.3. A Model of Diffusion. Having discussed several examples of network-based regressions, we now turn to a model of diffusion examined in Jackson and Rogers (2007b) which we discuss in

²⁷We say generically in the sense that given (G, x, β_0) , only a finite set of $\psi \in [0, 1]$ satisfy $E[Z_{G^S}u_{G^S}] = 0$.

the context of our microfinance example.²⁸ The researcher is interested in estimating this diffusion model which satisfies equilibrium moment equations. There are two states: whether or not a household endorses microfinance in a weekly village gathering. A non-endorsing household with d_i links may choose to endorse with probability $\nu_0 d_i \sigma_i$ where ν_0 is a transmission parameter and σ_i is the fraction of i 's neighbors that have decided to endorse. However, an endorsing household may naturally decide not to endorse, which can happen with probability δ_0 . Jackson and Yariv (2007) extend this model to a number of strategic environments.

The model is identified up to parameter $\beta_0 := \nu_0/\delta_0$, which is the transmission to recovery rate. Let $P(d)$ denote the degree distribution and $\rho(d)$ the share of nodes with degree d that endorse. Finally, $\bar{\rho}_* := \sum_d \rho(d)P(d)$ is the average endorsement rate in the network and the researcher observes $y := \bar{\rho} + \epsilon$, with ϵ an exogenous zero mean shock.

The second neighbors endorsement rate is given by $\sigma = (Ed)^{-1} \sum_d \rho(d)P(d) \cdot d$. Jackson and Rogers (2007b) use a mean-field approximation to derive a steady state equation,

$$\rho(d) = \frac{\beta_0 \sigma d}{1 + \beta_0 \sigma d}.$$

The equilibrium satisfies

$$(3.1) \quad \sigma(\beta_0) = (Ed)^{-1} \sum_d \frac{\beta_0 \sigma(\beta_0) d^2}{1 + \beta_0 \sigma(\beta_0) d} P(d).$$

By combining (3.1) with the definition of $\bar{\rho}$, we find that $h(G_r; \beta) := \sum_d \frac{\beta \sigma_r(\beta) d}{1 + \beta \sigma_r(\beta) d} P(d) = \bar{\rho}$. Therefore the researcher can use as moments

$$m(y_r, G_r; \beta) := y_r - h(G_r; \beta),$$

and estimate β_0 via nonlinear least squares. Jackson and Rogers (2007b) show that an equilibrium with non-zero endorsement rate exists only if $\beta > Ed/Ed^2$. The ratio of squared degree to degree, similar to what we have encountered when studying graph span, again becomes an important feature of the network. We put $\zeta := Ed^2/Ed$. Note that the typical summand of $h(G_r, \beta)$ is monotone and convex in d . Therefore, stochastic dominance relations among various distributions $P(d)$ will play a central role. We show that due to sampling of networks the researcher will overestimate the transmission parameter. An intuition for this is provided by the case of the star subgraph. This form of subsampling leads to a degree distribution that will be first order stochastically dominated by the true distribution. Therefore, the sampled network seems as if it has poorer diffusive properties; to generate the same average endorsement rate, the parameter governing the diffusion process must be higher. In addition, we show that the diffusion with the true parameter β_0 occurring on the sampled network may have no non-zero equilibria. When β_0 is close enough to the threshold $1/\zeta(G)$, the partially sampled network will make threshold ratio $1/\zeta(\bar{G})$ rise and therefore β_0 may appear to be less than $1/\zeta(\bar{G})$.

²⁸They study the SIS (susceptible, infected, susceptible) model of epidemiology, which they and others show have applications in a wide variety of economic contexts.

PROPOSITION 3.6. *Assume we have a triangular array (y_{Rr}, G_{Rr}) with degree distributions $P_{Rr}(d)$ and (i) (3.1) holds in expectation for each r , (ii) β_0 is such that there is a positive endorsement in every equilibrium, (iii) \mathcal{B} is a compact subset of $(0, \infty)$, (iv) (ϵ_r) are iid zero mean finite variance disturbances, and (v) $\limsup_{R \rightarrow \infty} \sup_{r \leq R} \sup_d |P_{Rr}(d) - P_{\infty r}(d)| = 0$.*

- (1) *The parameter estimates exhibit expansion bias: $\text{plim } \hat{\beta}(G^S) > \beta_0$ and $\text{plim } \hat{\beta}(G^{IS}) > \beta_0$.*
- (2) *For all r , β_0 is outside the range generating positive endorsement rate in the estimated equilibrium, with probability approaching one, under the following additional assumptions. Put $\delta_r := \beta_0 - 1/\zeta_r > 0$ and assume*
 - (a) *for star subgraphs, $\liminf_{R \rightarrow \infty} \zeta_r > 1 + \psi$ and $\limsup_{R \rightarrow \infty} \delta_r < (1 - \psi) \cdot \frac{1 - \zeta_r^{-1}(1 + \psi)}{\zeta_r + (1 - \psi^2)}$,*
 - (b) *or for induced subgraphs, $\limsup_{R \rightarrow \infty} \delta_r < (1 - \psi) \cdot \frac{1 + \zeta_r^{-1}}{\psi \zeta_r + (1 - \psi)}$.*

It is easy to see that for the star subgraph, an analytical solution to the bias is to use the degree distribution of the sampled nodes. However, this is a highly non-generic solution. The induced subgraph, for instance, does not allow this approach nor do other sampling schemes (e.g., randomly chosen edges, etc.) A natural question to ask is whether we may use the sampled degree distribution, such as $P^{IS}(d)$, to obtain $P(d)$. We note that this will not be straightforward to do, in general, because it generates an ill-posed inverse problem. The researcher is faced with an under-determined system; while we can describe how $P(d)$ maps into $P^{IS}(d)$ due to sampling, there appears to be no unique inverse. Graphical reconstruction, however, will provide a way to address the problem.

3.4. From Analytical Examples to Graphical Reconstruction. In this section we have analytically examined biases that emerge from sampled networks. We focused on three main network statistics: degree, graph span, and clustering. This choice was motivated by a number of applied questions concerning diffusion of information, network importance, risk-sharing, and social collateral. By analytically characterizing the biases, we were able to describe the mechanics of the non-classical measurement error and construct analytical corrections to eliminate the biases, under regularity conditions. The analytical study required us to focus on graph-level regressions; moreover, to be consistent, the analytical corrections focused on eliminating a slope effect of the biases, but needed to assume away or estimate a dispersion effect.

We also examined a model where a node's outcome depends on its neighbors' outcomes and demonstrated that a network-based instrumentation method violates the exclusion restriction when the network is sampled. With certain data structures, we provided a simple solution. Furthermore, we extended our analysis to a GMM model of diffusion and pointed out how the estimated parameters would exhibit expansion bias.

In general, our discussion has been on a case-by-case basis in this section. We have mostly focused on graph-level regressions and have been only able to examine very tractable network statistics. Numerous network statistics such as betweenness centrality, eigenvector centrality, and the aforementioned spectral statistics do not permit easy analytical examination nor corrections. The next section provides a more general method to estimate the economic parameter. Though the method is not limited to graph level regressions nor tractable network statistics, it comes at the cost of requiring more data and putting more structure on the problem by assuming models.

4. GRAPHICAL RECONSTRUCTION ESTIMATION

In this section, we discuss a two-step estimation procedure to consistently estimate economic parameters from linear regression and GMM models. In our asymptotic frame, both the size of each network and the number of networks grow. Every network is a draw from a distribution governed by its own parameter θ_{0r} . This will force us to control an incidental parameter problem. Clearly, we can nest the special case where every network is drawn from the same distribution, $\theta_{0r} = \theta_0$ for every r , and thereby assume away the incidental parameter problem.

We present an informal overview of our method in section 4.1. In section 4.2, we present the asymptotic distribution of $\hat{\beta}$ under high level assumptions on $\hat{\theta}_r$ and detail the estimation procedure in section 4.3. We discuss low level conditions for $\hat{\theta}_r$ by studying several key classes of network formation models, which also shed light on the limits of our approach, in section 4.4. Section 4.5 reflects on the interplay between network formation models and graphical reconstruction.

4.1. Informal Overview. In our overview we describe our procedure for regression,

$$y_{ir} = \alpha_0 + w_{ir}(G_r)\beta_0 + \epsilon_{ir}.$$

We assume that the researcher has the following data. First, she has outcome data for every node in every graph, $\{y_{ir} : i = 1, \dots, n, r = 1, \dots, R\}$, such as whether household i in village r participates in microfinance.²⁹ Second, she has a set of partially observed graphs, $\{G_r^S : r = 1, \dots, R\}$ or $\{G_r^{IS} : r = 1, \dots, R\}$. Third, she has variables which are predictive in a network formation model $\{z_r : r = 1, \dots, R\}$.³⁰ For instance, the researcher may have basic demographic characteristics such as religion, caste, household amenities, occupation or geographic location. This data structure is relatively innocuous and common in numerous applications. In development, when deciding how to draw a random sample to administer treatments, researchers usually conduct a listing in each enumeration area. This requires obtaining a census of the economic units, which can be done directly (e.g., Townsend, 2007; Suri, 2011; Banerjee et al., 2011) or indirectly by obtaining census information from the village representatives (e.g., Macours, 2003; Takasaki et al., 2000).³¹ It is well-known that obtaining GPS and basic demographic data during enumeration is cheap; the bulk cost of a network survey is the network module itself. For a different example, consider school networks where it is straightforward to obtain rosters and demographic data for all students. The full set of observed data is (y_r, A_r^{obs}, z_r) , consisting of y_r the vector of outcome data, A_r^{obs} the observed part of the graph, and z_r the vector of network formation covariates. The missing data for each network is A_r^{mis} and recall $G_r = (A_r^{obs}, A_r^{mis})$.

Every network is thought of as a realization of a random network formation process, drawn from a distribution which depends on z_r and parameter $\theta_{0r} \in \Theta_r$. To estimate β_0 we use an argument based on conditional expectations. If θ_{0r} were known for all r , we could estimate a conditional

²⁹In what follows it is not necessary for y_{ir} to be observed for every node, but it simplifies notation.

³⁰E.g., $z_r = \{z_{ir} : i = 1, \dots, n\}$ or $z_r = \{z_{ij,r} : i, j \in V\}$ where z_{ir} or $z_{ij,r}$ are covariates for nodes or pairs.

³¹Researchers can either collect simple covariate data from all nodes or from representatives who carry information.

expectation of $w_{ir}(G_r)$ given the observed data,

$$\mathcal{E}_{ir}(A_r^{obs}, z_r; \theta_{0r}) := \mathbb{E} \left[w_{ir}(G_r) | A_r^{obs}, z_r; \theta_{0r} \right].$$

By the properties of conditional expectation, using \mathcal{E}_{ir} in the regression instead of w_{ir} yields consistent estimation of β_0 . The least squares estimator is given by³²

$$\hat{\beta}_{ols} = \left(\sum_{r=1}^R \sum_{i=1}^n \mathcal{E}_{ir}(\hat{\theta}_r) \mathcal{E}_{ir}(\hat{\theta}_r)' \right)^{-1} \cdot \sum_{r=1}^R \sum_{i=1}^n \mathcal{E}_{ir}(\hat{\theta}_r) y_{ir}.$$

A similar but more involved result is true for GMM. Notice $\hat{\beta}_{ols}$ depends on $\hat{\theta}_r$ for all r .

To control the estimation of $\hat{\theta}_r$, we need to argue not only that it is consistent for θ_{0r} , but uniformly so. That is, $\sup_r \|\hat{\theta}_r - \theta_{0r}\| = O_P(a_R^{-1} \cdot R_n^{1/b})$, where a_R is the rate of convergence of $\hat{\theta}_r$ to θ_{0r} for every r , and $b > 1$ is the number of moments that the network formation model has. This imposes a rate requirement on the problem which says that the network-formation parameter needs to be estimated fast enough: $\sqrt{nR} \cdot a_R^{-1} \cdot R^{1/b} \rightarrow 0$.

The consistency of $\hat{\theta}_r$ follows from assumptions on the model of graph formation and the sampling procedure. With missing-at-random data, under assumptions on the graph model, a consistent estimator exists. Consider a model where an edge forms independently, conditional on covariates,

$$\mathbb{P}(A_{ijr} = 1 | z_r; \theta_{0r}) = \Lambda(f(z_{ir}, z_{jr})' \theta_{0r}),$$

where $\Lambda(\cdot)$ is some link function (e.g., logistic or normal), z_i is a vector of covariates for vertex i , and f is a vector-valued function. For instance, f may be the difference between characteristics of two nodes $f(z_i, z_j) = \|z_i - z_j\|$. If the sampling procedure is orthogonal to the network formation, a random subset of the $\binom{n}{2}$ pairs of nodes is observed. Therefore, $\hat{\theta}_r$ is consistent.

This model converges with $a_R = n$, since we have on the order of $n(n-1)/2$ observations. The requirement becomes $n^{-1/2} R^{1/2+1/b} \rightarrow 0$, so the number of networks must grow sufficiently slower than the number of nodes. In other models, the rate a_R may be different (e.g., $n/\log n$, n^τ for $\tau \in [1/2, 2)$, $\sqrt{n/\log n}$). If the rate is too slow, the requirement for node level regression may not be met, though usually the requirement for graph level regressions will be satisfied.

4.2. Formal Theory for $\hat{\beta}$. We begin by establishing that $\hat{\beta}$ is consistent and asymptotically normal. The main theorem is stated in section 4.2.2, under regularity conditions, including simple high level assumptions about the behavior of $\hat{\theta}_r$, which we will verify in section 4.4. In section 4.2.1 we discuss the regularity conditions in depth.

We have already introduced the regression environment. We consider the GMM environment of (2.1). Relative to regression, in GMM the value of y affects the conditional expectation of w . Observe (2.1) implies an unconditional moment restriction holds:

$$0 = \mathbb{E}m(X; \beta_0) = \sum_{G \in \mathcal{G}_n} \mathbb{E}[m(X; \beta_0) | G] \mathbb{P}_{\theta_0}(G)$$

³²For notational simplicity, assume the regressors are demeaned.

where $X = (y, w(G))$. Let x_r denote the triple of observed data, $x_r := (y_r, A_r^o, z_r)$. By iterated expectations, the conditional function

$$(4.1) \quad \mathcal{E}_{ir}(x_r; \beta_0, \theta_0) := \mathbb{E}[m(X_{ir}; \beta_0) | x_r; \beta_0, \theta_0]$$

satisfies $\mathbb{E}\mathcal{E}_{ir}(x_r; \beta_0, \theta_0) = 0$. Given an observed data series $\{(X_{ir}, z_{ir}) : i = 1, \dots, n, r = 1, \dots, R\}$ and an estimator $\widehat{\theta}_r$ of θ_{0r} , the estimator is

$$\widehat{\beta}_{\text{gmm}} := \underset{\beta \in \mathcal{B}}{\text{argmin}} \left(\mathbb{E}_{n,R} \mathcal{E}_{ir}(x_r; \beta, \widehat{\theta}_r) \right)' \widehat{W} \left(\mathbb{E}_{n,R} \mathcal{E}_{ir}(x_r; \beta, \widehat{\theta}_r) \right)$$

where \widehat{W} is a consistent estimator of W .³³

In order to compute the conditional moment in (4.1) we need to be able to integrate with respect to a conditional probability for every graph in our sample, $\mathbb{P}_{\beta_0, \theta_0}(A_r^{mis} | x_r)$. Computing the expectation requires a reweighting term,

$$\mathcal{E}_{ir}(x_r; \beta_0, \theta_{0r}) = \sum_{A_r^m} m(X_{ir}; \beta_0) \mathbb{P}_{\beta_0, \theta_{0r}}(A_r^{mis} | x_r),$$

with $\mathbb{P}_{\beta_0, \theta_{0r}}(A_r^{mis} | x_r) \propto f_{\beta_0}(y_r | G_r) \mathbb{P}_{\theta_{0r}}(A_r^{mis} | A_r^{obs}, z_r)$. To be able to utilize this approach, the researcher must make assumptions on the distribution of y given G .³⁴

4.2.1. Regularity Conditions. The main results are presented in section 4.2.2, to which the reader may skip ahead if desired. In this section we discuss the regularity conditions on which the results depend. Let $\mathbb{P}(G_r | z_r; \theta_r)$ be the distribution of the graph G_r given covariates z_r .

ASSUMPTION 1 (Random Graph Model and First Stage Estimation).

- (1) $\forall r$, Θ_r is a compact subset of \mathbb{R}^{d_θ} ; G_r is a \mathcal{G}_n -valued random graph with $\mathbb{P}(G_r | z_r; \theta_r) \in C^2(\Theta_r)$ at every $(G, z) \in \mathcal{G}_n \times \mathcal{Z}$; $\bar{H}_{r,R} := \sup_z \max_{G, \theta_r} \left| \frac{\partial}{\partial \theta} \mathbb{P}(G | z; \theta) \right|$, $\sup_R \sup_r \bar{H}_{R,r} < \infty$.
- (2) The first stage estimation satisfies for some sequence of normalizing constants (a_R) , $b > 1$, and $r \leq R$, $a_R \cdot (\widehat{\theta}_r - \theta_{0r}) = O_P(1)$ and $\sup_{r \leq R} \|\widehat{\theta}_r - \theta_{0r}\| = O_P(a_R^{-1} \cdot R^{1/b})$.
- (3) For node level analysis $a_R^{-1} \cdot \sqrt{nR^{1+2/b}} \rightarrow 0$ and for graph level analysis $a_R^{-1} \cdot \sqrt{R^{1+2/b}} \rightarrow 0$.
- (4) $\beta_0(u)$ is an interior point of \mathcal{B} , a compact subset of \mathbb{R}^{d_β} , for every $u \in \mathcal{U}$.

Condition 1 ensures that the random graph family is smooth enough in the parameter, so small deviations from the true parameter do not result in very different probability distributions. Condition 2 is a high-level condition on the first stage estimation which we will microfound in section 4.4. It guarantees that we can uniformly replace the estimated network formation parameter for every graph in the sequence with its true value. Condition 3 is a rate requirement which relates the rate of estimation of the network formation process to the rate of estimating the economic model of interest. Condition 4 is a standard interiority condition. Let h denote a random variable.

³³In the case of maximum likelihood where \mathcal{E} is the conditional score, $W = I$.

³⁴With an index $u \in \mathcal{U}$, $\widehat{\beta}(u) := \underset{\beta \in \mathcal{B}}{\text{argmin}} \left(\mathbb{E}_{n,R} \mathcal{E}_{ir}(x_r; \beta, \widehat{\theta}_r, u) \right)' \left(\mathbb{E}_{n,R} \mathcal{E}_{ir}(x_r; \beta, \widehat{\theta}_r, u) \right)$.

DEFINITION 4.1. A sequence of measurable (potentially matrix-valued) functions $\{\phi_{i,r}(h_{ir}; \alpha) : i = 1, \dots, n_R, r = 1, \dots, R\}$ satisfies an envelope condition over $\alpha \in \mathcal{A}$ if there exist measurable functions $L_{i,r}(h_{ir})$, with $\|\phi_{i,r}(h_{ir}; \alpha)\| \leq L_{i,r}(h_{ir})$ for every h_{ir} and α , and $\sqrt{n}\mathbb{E}_n L_{i,r,R}$ has uniformly integrable v th moment for $v \geq 2$.

DEFINITION 4.2. A sequence of measurable (potentially matrix-valued) functions $\{\phi_{i,r}(h_{ir}; \alpha) : i = 1, \dots, n_R, r = 1, \dots, R\}$ is Lipschitz continuous in $\alpha \in \mathcal{A}$ if there exist measurable functions $M_{i,r}(h_{ir})$ with $\|\phi_{i,r}(h_{ir}; \alpha) - \phi_{i,r}(h_{ir}; \bar{\alpha})\| \leq M_{i,r}(h_{ir}) \|\alpha - \bar{\alpha}\|$ for every h_{ir} and $\alpha, \bar{\alpha} \in \mathcal{A}$, and $\sqrt{n}\mathbb{E}_n M_{i,r,R}$ has uniformly integrable v th moment for $v \geq 2$.

In addition, we use $\mathcal{I}_{h|x}(\alpha) := \mathbb{E} \left[\frac{\partial}{\partial \alpha'} \log f(h|x; \alpha) \frac{\partial}{\partial \alpha} \log f(h|x; \alpha) | x; \alpha \right]$ to denote the conditional information matrix with random variable $h|x$, density or pmf $f(h|x)$, and parameter α .

Turning to the economic model, observe that the network statistic $w(G)$ may be growing or shrinking in n . For instance, the eigenvector centrality declines as it is a unit norm object. The degree of a node may be $\Theta(1)$, $\Theta(\log n)$, or $\Theta(n)$ depending on the graph family. In what follows, in regression we assume that the model is such that all regressors are rescaled at the appropriate rate: if they exhibit growth or shrinkage at b_n , we assume that the models are specified using $\tilde{w} := b_n^{-1}w$ as regressors. Similarly, in GMM we assume that the moments and network statistics, both of which may depend on R , are appropriately rescaled. For two reasons we present regularity conditions for least squares, GMM, and GMM with an index separately, though they essentially can be nested. First, least squares does not require assuming the joint distribution of y and G . Second, GMM conditions are more transparent than the more general case where parameters carry an index. After presenting the assumptions we discuss what they mean for networks.

ASSUMPTION 2 (Linear Regression).

- (1) $\mathbb{E}[\epsilon|w] = 0$, $\mathbb{E}[\epsilon\epsilon'|w] = \Omega$, *p.d.* with $\sup_R \lambda_{\max}(\Omega) < \infty$
- (2) $\mathbb{E}[\|w_{ir}\|^k | x_r; \theta_r]$ and $\|\mathcal{I}_{w_{ir}|x_r}(\theta_r)\|$ for $k = 1, 2$ satisfy the envelope condition with $L_{i,r}(x_r)$.
- (3) $\sup_{R \geq 1} \sup_{r \leq R} \text{var}(\sqrt{n}\mathbb{E}_n w_{ir}(G_r)) < C_1 < \infty$ and $\inf_{R \geq 1} \inf_{r \leq R} \text{var}(\sqrt{n}\mathbb{E}_n \mathcal{E}_{ir}(x_r; \theta_{0r})) > C_0 > 0$, uniformly over the array.

Define $g_R(\beta) := \mathbb{E}_{n,R} \text{Em}(y_{ir}, w_{ir}(G_r); \beta)$ and $f(m|x; \beta, \theta)$ be known up to parameters.

ASSUMPTION 3 (GMM).

- (1) $\widehat{W} = W + o_P(1)$, W is *p.s.d.* and the model satisfies $\lim_{R \rightarrow \infty} W g_R(\beta) = 0$ only if $\beta = \beta_0$.
- (2) The limits $\lim_{R \rightarrow \infty} \mathbb{E}_{n,R} [\mathbb{E} \mathcal{E}_{ir}(x_r; \theta_r, \beta)]$ and $\lim_{R \rightarrow \infty} \mathbb{E}_{n,R} \left[\mathbb{E} \frac{\partial}{\partial \beta'} \mathcal{E}_{ir}(x_r; \theta_r, \beta) \right]$ exist uniformly over $\mathcal{B} \times \prod_{r \in \mathbb{N}} \Theta_r$.
- (3) $\|\mathcal{I}_{m_{ir}|x_r}(\beta, \theta_r)\|$, $\mathbb{E} \left[\left\| \frac{\partial}{\partial \beta'} m(X_{ir}; \beta) \right\| | x_r; \theta_r, \beta' \right]$, and $\mathbb{E} \left[\|m(X_{ir}; \beta)\|^k | x_r; \theta_r, \beta' \right]$ for $k = 1, 2$ satisfy the envelope condition with envelope $L_{i,r}(x_r)$.
- (4) $m(X; \beta)$ is continuously differentiable on the interior of \mathcal{B} for every $X \in \mathcal{X}$ and both $m(X; \beta)$ and $\frac{\partial}{\partial \beta'} m(X; \beta)$ satisfy the Lipschitz condition with constant $M_{i,r}(X_{ir})$, where $\mathbb{E}[M_{i,r}(y_{ir}, w_{ir}) | x_r] \leq L_{i,r}(x_r)$.
- (5) $\sup_{R \geq 1} \sup_{r \leq R} \text{var}(\sqrt{n}\mathbb{E}_n [m(X_{ir}; \beta_0)]) < C_1 < \infty$ and $\inf_{R \geq 1} \inf_{r \leq R} \text{var}(\sqrt{n}\mathbb{E}_n [\mathcal{E}_{ir}(x_r; \beta_0, \theta_{0r})]) > C_0 > 0$, uniformly over the array.

Finally define function classes $\mathcal{F}_R^* := \{\sqrt{n}\mathbb{E}_n m(X_{ir}; \beta(u), u) : (\beta, u) \in \mathcal{B} \times \mathcal{U}\}$ and for $\beta', u' \in \mathcal{B} \times \mathcal{U}$, $\mathcal{H}_{R|\beta', u'} := \{\sqrt{n}\mathbb{E}_n \mathbb{E}[m(X_{ir}; \beta', u') | x_r; \beta, u] : (\beta, u) \in \mathcal{B} \times \mathcal{U}\}$. This nests the above with $\mathcal{U} = \{u\}$.

ASSUMPTION 4 (Indexed GMM).

- (1) The maps $m(X_{ir}; \beta, u)$ are measurable and is continuous at each (β, u) with probability one and $\mathbb{P}(A_r^m | x_r; \beta, u)$ is continuous at each (β, u) with probability one.
- (2) $\mathbb{E}\mathcal{E}_{ir}(\beta(u), u)$ is continuously differentiable at $\beta_0(u)$ uniformly in \mathcal{U} and $\frac{\partial}{\partial \beta'} \mathbb{E}\mathcal{E}_{ir}(\beta, u)$ is uniformly non-singular at $\beta_0(u)$ over \mathcal{U} .
- (3) $\|\mathbb{E}[m(X_{ir}; \beta, u) | x_r; \beta, u]\| \leq L_{ir}(x_r)$, $\sup_R \sup_r \mathbb{E}(\mathbb{E}_n L_{i,r}(x_r))^{2+\delta} = O(1)$ for some $\delta > 0$.
- (4) The following uniform entropy integral condition holds:

$$\int_0^\infty \sup_{Q \in \mathcal{Q}} \sqrt{\log N(\epsilon \|\bar{F}_R^*\|_{Q,2}, \mathcal{F}_R^*, L_2(Q))} d\epsilon + \sup_{(\beta', u') \in \mathcal{B} \times \mathcal{U}} \int_0^\infty \sup_{Q \in \mathcal{Q}} \sqrt{\log N(\epsilon \|\bar{H}\|_{Q,2}, \mathcal{H}_{R|\beta', u'}, L_2(Q))} d\epsilon < \infty.$$

Assumptions 2, 3, and 4 are similar, so we discuss the GMM case. Assumption 3.1 is a standard identification condition. Assumption 3.2 is standard (e.g., Andrews, 1994) and Assumption 3.3 places uniform restrictions on higher moments of the conditional moment, slope of the moment, and information matrix allowing weak laws of large numbers to be applied. Assumption 3.4 allows these convergences to be uniform over the parameter space.

Assumption 3.5 is what allows us to pass a central limit theorem to the conditional random variable if the unconditional satisfies one.³⁵ It is reasonable in practice because we use independence across graphs and simply a uniform boundedness condition within graph. This is substantially weaker than having to assume a within-graph central limit theorem for m_{ir} , which would depend on the idiosyncrasies of the network formation model and network statistics. However, it comes at the cost of requiring data from multiple networks. We make this assumption because currently the statistics of networks literature has not characterized within-graph node characteristic interdependencies (e.g., the correlation of eigenvector centrality between nodes for various random graph families). Assumption 4.4 allows for a functional central limit theorem to be applied. Intuitively, the condition on \mathcal{F}_R^* would have to be assumed for the result to be true even without missing data and $\mathcal{H}_{R|\beta', u'}$ controls how much the conditional expectation changes as (β, u) change.

4.2.2. *Asymptotic Distribution.* In this section we show that $\hat{\beta}_{ols}$, $\hat{\beta}_{gmm}$, and $\hat{\beta}(u)$ are consistent and asymptotically Gaussian. We define covariance matrices which characterize the asymptotic variance. For linear regression,

$$H_{ols} := \lim_{R \rightarrow \infty} \mathbb{E}_n R [\mathbb{E} \mathcal{E}_{ir} \mathcal{E}'_{ir}] \text{ and } V_{ols} := \lim_{R \rightarrow \infty} \mathbb{E}_R [\text{var}(\sqrt{n} \mathbb{E}_n [\mathcal{E}_{ir} \epsilon_{ir} + \mathcal{E}_{ir}(w_{ir} - \mathcal{E}_{ir})' \beta_0])],$$

and for GMM,

$$M := \lim_{R \rightarrow \infty} \mathbb{E}_n R \left[\mathbb{E} \frac{\partial}{\partial \beta'} \mathcal{E}_{ir}(x_r; \beta_0, \theta_{0r}) \right], \quad \Omega := \lim_{R \rightarrow \infty} \mathbb{E}_R [\text{var}(\sqrt{n} \mathbb{E}_n \mathcal{E}_{ir}(x_r; \beta_0, \theta_{0r}))],$$

$$H_{gmm} := M' W M \text{ and } V_{gmm} := M' W \Omega W' M.$$

³⁵Since this paper focuses on the effect of sampling on network analysis and not on regression or GMM models on graphs, we make the assumption that the underlying model satisfies reasonable regularity conditions if the full networks were observed and focus on the effect of sampling and graphical reconstruction.

THEOREM 4.1 (Asymptotic Distribution). *Under Assumptions 1,*

- (1) *Assumption 2 implies $\sqrt{nR}(\hat{\beta}_{\text{ols}} - \beta_0) \rightsquigarrow \mathcal{N}(0, H_{\text{ols}}^{-1} V_{\text{ols}} H_{\text{ols}}^{-1})$.*
- (2) *Assumption 3 implies $\sqrt{nR}(\hat{\beta}_{\text{gmm}} - \beta_0) \rightsquigarrow \mathcal{N}(0, H_{\text{gmm}}^{-1} V_{\text{gmm}} H_{\text{gmm}}^{-1})$.*
- (3) *Assumption 4 implies $\sqrt{nR}(\hat{\beta}(\cdot) - \beta_0(\cdot)) \rightsquigarrow -\dot{\Psi}_{\beta,u}(\beta_0(\cdot), \cdot)^{-1} \mathcal{Z}(\beta_0(\cdot), \cdot)$, in $\ell^\infty(\mathcal{U})$, a mean-zero Gaussian process with covariance function*

$$\Omega(u, \tilde{u}) := \lim_{R \rightarrow \infty} \mathbb{E}_R \left[n^{-1} \sum_{i,j} \mathbb{E} [\mathcal{E}_{ir}(x_r; \beta_0(u), u) \mathcal{E}_{jr}(x_r; \beta_0(\tilde{u}), \tilde{u})] \right]$$

where $\dot{\Psi}_{\beta,u}(\beta, u) = \lim_{R \rightarrow \infty} \mathbb{E}_{n,R} \mathbb{E} \frac{\partial}{\partial \beta} \mathcal{E}_{ir}(\beta, u)$.

Intuitively, if we can uniformly replace $\hat{\theta}_r$ with θ_{0r} , since conditional expectations are centered correctly and, under regularity conditions, also satisfy central limit theorems if the unconditioned random variables do, the estimator is consistent and normal. While we wrote the theorem for vertex-level analysis, similar results with modified regularity conditions extend to regressions at the graph-level, edge-level, vertex-triples, etc. Each will allow for different amounts of interdependency in the graph formation process. To be concrete, under the above normalizing assumptions, graph level regression converges at \sqrt{R} while edge level regression converges at $\sqrt{\binom{n}{2} R} = n\sqrt{R}$.

To build further intuition, we comment on what could go wrong. First, for GMM, if one estimates the conditional expectation without reweighting, unless the model was additively separable, $\hat{\beta}_{\text{gmm}}$ would be inconsistent. Second, there are several reasons why uniform estimation may fail: the size of the networks relative to the number of networks may be too small, the network formation process may have $\dim(\Theta_r)$ exploding too fast, and the level of interdependency in the random graph processes may be too high. We provide a more detailed discussion in section 4.5.

4.3. Estimation in Practice. We describe the estimation algorithm for linear regression.

ALGORITHM (Estimation of $\hat{\beta}_{\text{ols}}$).

- (1) Use (z_r, A_r^{obs}) to estimate $\hat{\theta}_r$ based on the assumed network formation model.
- (2) Estimate $\mathcal{E}_{ir}(A_r^{\text{obs}}, z_r; \theta_{0r}) := \mathbb{E} [w_{ir}(G_r) | A_r^{\text{obs}}, z_r; \theta_{0r}]$.
 - (a) Given (z_r, A_r^{obs}) , for simulations $s = 1, \dots, S$, draw $A_{r,s}^{\text{mis}^*}$ from $\mathbb{P}_{\hat{\theta}_r}(A_r^{\text{mis}} | A_r^{\text{obs}}, z_r)$.
 - (b) Construct $w_{ir}(G_{r,s}^*)$ where $G_{r,s}^* = (A_r^{\text{obs}}, A_{r,s}^{\text{mis}^*})$.
 - (c) Estimate $\hat{\mathcal{E}}_{ir}(A_r^{\text{obs}}, z_r; \hat{\theta}_r) := \frac{1}{S} \sum_{s=1}^S w_{ir}(G_{r,s}^*)$.
- (3) Estimate $\hat{\beta}_{\text{ols}}$ using data $\{(y_{ir}, \hat{\mathcal{E}}_{ir}(A_r^{\text{obs}}, z_r; \hat{\theta}_r)) : i = 1, \dots, n, r = 1, \dots, R\}$.

The GMM algorithm is similar, requiring a reweighting term. We provide an overview of standard errors and estimation methods in Appendix C. In practice, clustering at the graph level in vertex-level regressions and using heteroskedasticity robust standard errors for network-level regressions perform well, though we have explored various bootstrapping procedures (available upon request).

4.4. Formal Theory for $\hat{\theta}_r$. In this section we discuss the uniform estimation of the network formation model parameters. We are interested in the joint convergence of $\sup_r \|\hat{\theta}_r - \theta_{0r}\|$ in the sense of Assumption 1.3. The literature on consistently estimable network formation models is young and limited. Most models of network formation lack asymptotic frames (see, e.g., exponential

random graphs models (ERGMs)). That is, larger networks do not lead to tighter parameter estimates. In general there are no known results characterizing consistency for ERGMs and is currently the topic of research in [Chandrasekhar and Jackson \(2011\)](#). In addition, [Goldsmith-Pinkham and Imbens \(2011\)](#) and [Kolotilin \(2011\)](#), among others, are working papers currently developing consistently estimable random graph models. There are a few classes of models known to be consistent and we discuss several as examples below. Given how new this literature is, it is useful to reflect on a simple, checkable sufficient conditions for joint convergence so that one could check new models as they develop. After this, we discuss three common classes of network formation models and check the condition that can be used in graphical reconstruction. The examples have been chosen to provide intuition about different problems that may arise.

We have a collection of network formation models which maximize criterion functions, $\theta_{0r} = \arg \max_{\theta} Q_{(r)}(\theta_{0r})$. We estimate these parameters with a collection of empirical criterion functions, $\widehat{Q}_{(r)}(\theta_r)$, with $\widehat{\theta}_r = \arg \max_{\theta} \widehat{Q}_{(r)}(\theta_r)$. The lemma is analogous to [Hahn and Newey \(2004\)](#).

LEMMA 4.1. *Let $V_{(r)}(\theta_r) := \nabla_{\theta} Q_{(r)}(\theta_r)$ and $\widehat{V}_{(r)}(\theta_r) := \nabla_{\theta} \widehat{Q}_{(r)}(\theta_r)$. Assume the following.*

- (1) $\forall r$, $Q_{(r)}(\theta_r)$ has unique maximum θ_{0r} ; Θ_r is compact; $Q_{(r)}(\theta) \in C^2(\Theta)$; $\sup_{\theta} |\widehat{Q}_{(r)}(\theta) - Q_{(r)}(\theta)| = o_{\mathbb{P}}(1)$.
- (2) *The criterion functions uniformly converge in the sense that for some $v > 0$*

$$\mathbb{P} \left(\sup_{r \leq R} \sup_{\theta \in \Theta_r} |\widehat{Q}_{(r)}(\theta) - Q_{(r)}(\theta)| \geq \eta \right) = o(a_R^{-v}).$$

- (3) *There exists a sequence of constants (a_R) such that (i) for all r , $a_R \cdot \widehat{V}_{(r)}(\theta_{0r}) = O_{\mathbb{P}}(1)$; (ii) for some $b > 1$, $\sup_{r \leq R} \mathbb{E} \left\| a_R \cdot \widehat{V}_{(r)}(\theta_{0r}) \right\|^b < \infty$.*
- (4) $\nabla \widehat{V}_{(r)}(\theta_r)$ satisfies a Lipschitz condition with coefficient B_r , $\sup_r \|B_r\| = O_{\mathbb{P}}(1)$.
- (5) *The Hessian satisfies $\sup_r \left\| \nabla \widehat{V}_{(r)}(\theta_r) - \nabla V_{(r)}(\theta_r) \right\| = o_{\mathbb{P}}(1)$.*

Then $a_R \cdot R^{-1/b} \cdot \sup_{r \leq R} \left\| \widehat{\theta}_r - \theta_{0r} \right\| = O_{\mathbb{P}}(1)$.

This comes from a usual first order expansion argument. Condition 1 adds extra smoothness to a standard assumption for consistency. Condition 2 requires that all the criterion functions $\widehat{Q}_{(r)}(\theta)$ uniformly lie in an η -“sleeve”, $[Q_{(r)}(\theta) - \eta, Q_{(r)}(\theta) + \eta]$; in practice this is argued by applying union bounds and controlling interdependencies across summands in the objective function. Condition 3 provides a rate of convergence of the first-order term and a moment requirement. Condition 4 requires an envelope condition for the third derivative of the objective. Condition 5 requires uniform convergence of the Hessian. Below, Lemma 4.1 holds under low-level assumptions.

4.4.1. *Mixing Coefficient.* To describe interdependence in the data, we define a mixing coefficient. Let $D \subset \mathbb{Z}^d$ be an integer lattice and to each $s \in D$ we associate a random variable z_s . Then $\{z_s : s \in D\}$ is a random field and we are interested in controlling the dependence of z_s and $z_{s'}$. Let \mathcal{A}_{Ω} be the σ -algebra generated by a random field $\{z_s : s \in \Omega\}$. We define the mixing coefficient

$$\alpha_{k,l}(m) := \sup \{ |\mathbb{P}(A_1 \cap A_2) - \mathbb{P}(A_1)\mathbb{P}(A_2)| : A_i \in \mathcal{A}_{\Omega_i}, |\Omega_1| \leq k, |\Omega_2| \leq l, d(\Omega_1, \Omega_2) \geq m \}$$

where $d(\Omega_1, \Omega_2) = \min_{x,y \in \Omega_1 \times \Omega_2} \|x - y\|_1$. We will need to assume that the level of interdependence goes to zero as the distance between the two subsets goes to infinity.³⁶

4.4.2. *Classes of Models.* We discuss three common classes of network formation models that may be used in graphical reconstruction. The examples are intended to provide a broad view of the assumptions that need to be made. The first example is the most parsimonious and assumes that the probability that an edge forms is independent given covariates. The second example allows for more flexibility in describing the structure of the network by having a rapidly growing a number of parameters. Despite this, network level reconstruction meets sufficient conditions, though node level reconstruction does not. The third example allows for network effects, where the probability of an edge forming depends on the probabilities of other edges as well as covariates.

Class 1: Conditional Edge Independence Models. We begin by considering a class of models in which edges form independently, given covariates. This is the most common class of model used in the literature (see e.g., Jackson 2008, Christakis et al., 2010, Goldsmith-Pinkham and Imbens, 2011, and Santos and Barrett, 2008). Let Ξ be a set consisting of all pairs ij . Ξ is implicitly indexed by n and has $n(n - 1)/2$ elements. We denote an element $s \in \Xi$ and, when referencing explicitly which pair it corresponds to, we write $s = s_{ij}$. Let z_s denote a covariate for the pair of nodes s_{ij} . Examples include whether two villagers are of the same caste, the distance between their households, etc. The probability that an edge forms in graph r is

$$(4.2) \quad \text{P}(A_{sr} = 1 | z_{sr}; \theta_{0r}) = \Phi(z'_{sr} \theta_{0r})$$

where $\Phi(\cdot)$ is some link function. This framework allows us to consider undirected graphs, directed graphs, and models in which nodes have to agree for a link to form. The undirected case is clear. If the graph formation model is directed, then Ξ consists of all $n(n - 1)$ ordered pairs of ij . When the model is undirected but both nodes need to agree, one may use a model such as $A_{ijr} = \mathbf{1}\{z'_{ij} \theta_{0r} - \epsilon_{ijr} \geq 0\} \cdot \mathbf{1}\{z'_{ji} \theta_{0r} - \epsilon_{jir} \geq 0\}$ with link function $\Phi(z'_{sr} \theta_{0r}) := \Psi(z'_{ij} \theta_{0r}) \Psi(z'_{ji} \theta_{0r})$, where $\Psi(\cdot)$ is the cdf of ϵ .³⁷

We maximize the log-likelihood, $|\Xi|^{-1} \sum_{s \in \Xi} q(X_s; \theta_r)$, $X_s = [A_s, z'_s]$, with summand

$$q(X_s; \theta_r) = A_{sr} \log \Phi(z'_{sr} \theta_r) + (1 - A_{sr}) \log (1 - \Phi(z'_{sr} \theta_r)).$$

For joint convergence, we require that Φ is such that the following hold.

ASSUMPTION 5 (Joint Convergence). Let $Q_{(r)}(\theta_r) := \text{plim}_{n \rightarrow \infty} |\Xi|^{-1} \sum_{s \in \Xi} \text{E}q(X_s; \theta_r)$.

- (1) $\forall \eta > 0, \inf_{r \leq R} [Q_{(r)}(\theta_{0r}) - \sup_{\theta: \|\theta - \theta_{0r}\| > \eta} Q_{(r)}(\theta)] > 0$.
- (2) $D^{|v|} q(X_{sr}; \theta)$ satisfies a Lipschitz condition with $B(X_{sr})$, for some multi-index $|v| \geq 2$.
- (3) 2^{b-1} moments exist for envelope $B(z_{sr}) \geq \|\partial Q_{(r)}(\theta_r) / \partial \theta_r\|$.

³⁶The triangular array notation is cumbersome, see Jenish and Prucha (2009), but formally is $\{z_{s,R} : s \in D_R, R \in \mathbb{R}\}$ a triangular array defined on a sequence of probability spaces where D_R is a finite subset of D and

$\alpha_{k,l}(m) := \sup_{R \geq 1} \sup \{ |P^R(A_1 \cap A_2) - P^R(A_1)P^R(A_2)| : A_i \in \mathcal{A}_{\Omega_i}^R, |\Omega_1^R| \leq k, |\Omega_2| \leq l, \Omega_i \subset D_R, d(\Omega_1, \Omega_2) \geq m \}$.

³⁷A more realistic model may have ϵ_{ij} and ϵ_{jr} being jointly normal.

Condition 1 is standard for identification, 2 requires sufficient smoothness, and 3 requires that the score functions have well-behaved envelopes. Many assumptions about $\Phi(\cdot)$ ensure that Assumption 5 holds (e.g., if the link function is logistic).

To be able to apply Lemma 4.1, we need to control the interdependence in the covariates z_{sr} . We assume that the set of nodes itself has an embedding into an integer lattice, $\Lambda \subset \mathbb{Z}^d$. Let $t \in \Lambda$ denote a generic element, and when referencing the corresponding node we write $t = t_i$. To build intuition imagine that the nodes are embedded in \mathbb{Z}^2 as analogous to geographic placement; households in a village are placed on a grid on the ground and certain households are closer to others. This closeness determines the covariance of their other characteristics. Then, every node is given a random covariate z_i , $t_i \in \Lambda$. The pair level covariate z_{ij} is given by $z_{ij} = f(z_i, z_j)$ for some function $f(\cdot, \cdot)$. The interdependency in the node level covariates will translate to interdependencies among the edge level covariates which is what we will ultimately use in our argument.³⁸

We define a distance (pseudo-metric) $d_{\Xi}(\cdot, \cdot)$ over the set of pairs, given by Ξ , where two pairs ij and kl 's distance is said to be the minimum coordinate-wise distance between an element of the first pair and an element of the second pair. Specifically, for $s_{ij}, s_{kl} \in \Xi$ and $d_{\Lambda}(t_i, t_j) := \|t_i - t_j\|_1$,

$$d_{\Xi}(s_{ij}, s_{kl}) := d_{\Lambda}(t_i, t_j) \wedge d_{\Lambda}(t_i, t_l) \wedge d_{\Lambda}(t_j, t_k) \wedge d_{\Lambda}(t_k, t_l).$$

ASSUMPTION 6 (Mixing Conditions). $\forall r, z_r := \{z_{ir} : t_i \in \Lambda \subset \mathbb{Z}^d\}$ is a stationary mixing random field, $z_{ijr} := f(z_{ir}, z_{jr})$ satisfies $\sup_r \mathbb{E} \|z_{ijr}\|^{p+\delta} \lesssim \sup_r \mathbb{E} \|z_{1r}\|^{p+\delta}$, and $\sup_r \mathbb{E} \|z_{1r}\|^{p+\delta} < \infty$, with (i) $\sup_r \alpha_{2,\infty}^r(m) \leq Ca^m$ for $a \in (0, 1)$ or (ii) $\sup_r \alpha_{2,\infty}^r(m) = o(m^{-d})$.

While hard to verify, this assumption is analogous to those made in time series and spatial econometrics contexts. We require that the random fields z_r satisfy uniform mixing requirements where, as the distance between the sites of two random variables increase, the level of interdependency decays fast enough. The assumption on f is not very restrictive. The most natural example is a covariate based on the difference in characteristics of nodes i and j : $z_{ij} = \|z_i - z_j\|$. It is easy to see that $\mathbb{E} \|z_{ij}\|^k \leq 2^k \mathbb{E} \|z_1\|^k$ by the binomial theorem and stationarity. Because A_{sr} is a measurable function of z_{sr} , it will inherit stationarity and mixing properties and therefore so will X_{sr} .

PROPOSITION 4.1. *Assumptions 5 and 6 imply the conditions of Lemma 4.1.*

Since there are effectively $n(n-1)/2$ edge formation outcomes being estimated, usually $a_R = n$. Compared to slower rates, this enables us to have relatively more graphs for a given level of network size while still satisfying the rate requirement of Assumption 1.3.³⁹

Until now we have not discussed the role of random sampling. It is easy to see with random sampling of nodes (G^S or $G^{|S}$) or random sampling of pair data A_{ij} that the criterion function $Q_{(r)}(\theta) := \lim_{n \rightarrow \infty} \left| \Xi^S \right|^{-1} \sum_{s \in \Xi} \mathbb{E}[q(X_{sr}; \theta) \mathbf{1}\{s \in \Xi^S\}]$ is minimized at true parameter θ_{0r} . For instance, under the star subgraph we have $\mathbb{E}[q(X_{sr}; \theta) \mathbf{1}\{s \in \Xi^S\}] = (1 - (1 - \psi)^2) \mathbb{E}q(X_s; \theta_{0r})$ while

³⁸Notice $|\mathbb{P}(A_{ij} \cap A_{kl}) - \mathbb{P}(A_{ij})\mathbb{P}(A_{kl})| \leq Ca^m$ for pairs ij, kl where $i, j \neq \{k, l\}$ and $d(ij, kl) = m$ since $\Lambda_1 = \{i, j\}$ and $\Lambda_2 = \{k, l\}$ generate sub-algebras. The only time when we cannot control the probabilities like this is for the $2(n-1)$ terms where two distinct pairs share one index, but there are only $O(n)$ of these and $|\Xi| = O(n^2)$.

³⁹One can modify this model to change a_R (e.g., Christakis et al., 2010 and Goldsmith-Pinkham and Imbens, 2011) by introducing truncation terms which make certain nodes i and j have essentially zero probability of being linked. A simple example uses a covariate $-|i - j|$ with nodes $i \in \mathbb{N}$ yields $a_R = \sqrt{n}$.

$|\Xi^S| = (1 - (1 - \psi)^2) \binom{n}{2}$. More generally, if the sampling procedure is known, then in such a model augmenting the likelihood to account for the sampling will produce consistent estimates.

Class 2: Group Models. By allowing for an increasing number of parameters, a network formation model may be able to better and more flexibly describe the random graph process. Models of this vein are discussed in [Bickel and Chen \(2009\)](#), among others, who provide a discussion of what they call a nonparametric view of network models. Our specific example comes from [Chatterjee et al. \(2010\)](#), who study an environment in which the degree distribution is the sufficient statistic for graph formation: given (d_1, \dots, d_n) , one estimates a formation model.⁴⁰ Following [Diaconis and Freedman \(1984\)](#) they show the network is described by

$$P(A_{ij} = 1) \propto \exp(\theta_{0i} + \theta_{0j}),$$

which is a model that allows the number of parameters to grow at a $\Theta(n)$ rate.

We extend this framework to our environment and assume there are k_n categories of nodes. For instance, if graph formation depends on two characteristics, gender (male/female) and education (high/low), there are four such categories. By allowing k_n to grow rapidly with n , we can capture substantial variation in the formation of the network. We allow $k_n = \Theta(n)$. Define an equivalence class of nodes: if i and j are in the same class, then they have the same parameter, $\theta_{0i} = \theta_{0j}$. In our example, two individuals in the same category (e.g., female and high education) are governed by the same parameter $\theta_{0,female,high}$. If we have q_n characteristics with a uniformly bounded values (e.g., two genders, a bounded number of education levels), the number of categories can grow at $q_n = \Theta(\log n)$, which yields $k_n = \Theta(n)$. One can think of this model as having group fixed effects with a growing number of groups. It turns out that with probability approaching one, $\sup_{r \leq R} \|\hat{\theta}_r - \theta_{0r}\|_\infty \lesssim \sqrt{\log n/n}$ which is a very slow rate, though expected given how rapidly we are increasing the parameter dimension.

PROPOSITION 4.2. *Let the maximum coordinate value of $\theta_r \in \Theta_r$ be uniformly bounded over all r , $R = o(n \cdot \log^{-1} n)$, and $k_n = \Theta(n)$, $k_n < n$. Then, under stratified random sampling with either the induced or star subgraph and Assumptions 2, 3, or 4, the conclusion of Theorem 4.1 holds.*

This example shows that even when we are adding parameters at rate n , graphical reconstruction is possible in network level analyses. Here $a_R = \sqrt{n/\log n}$ and therefore $a_R^{-1} \cdot \sqrt{R^{1+2/b}} \rightarrow 0$. Meanwhile, the sufficient condition is not met for vertex level analysis as $a_R^{-1} \cdot \sqrt{nR^{1+2/b}} \rightarrow \infty$. This example provides an illustration of both the strengths and limitations of graphical reconstruction by testing the limits as we add dimensions at the same rate as the number of nodes. Certainly more slowly growing parameter dimension will only help.

Class 3: Models with Network Effects. Our last example studies a graph formation model with network effects where the edge-formation probability depends on other edges up and beyond through the correlated covariate effect. In general, the principle difficulty in studying network formation models is balancing the interdependency while still allowing consistent estimation of parameters. As previously discussed, most off-the-shelf network formation models lack asymptotic theory. In

⁴⁰Conditional on the degree distribution there is no information about the model from the actual network data.

this example we illustrate the consistency of graphical reconstruction with a simple model that exhibits network effects, taken from the dynamic binary choice literature (De Jong and Woutersen, 2010). Similar to Christakis et al. (2010), we assume that pairs of nodes meet in a sequence, and for simplicity we assume the sequence is known. Consequently there are $\binom{n}{2}$ periods in which pairs meet. Assume that the probability that i and j , meeting at period s , form a link is given by

$$P(A_{s,r} = 1 | z_{sr}, A_r^{s-1}; \theta_{0r}) = \Phi \left(\sum_{k=1}^l \theta_{0r,k} A_{s-k,r} + \theta'_{0r,l+1} z_{sr} \right)$$

where A_r^{s-1} is the structure of the graph that has formed through period $s-1$. The probability i and j form an edge depends on covariates as well as a finite collection of previously formed edges.

ASSUMPTION 7 (Mixing Conditions). (1) For each r , $z_r := \{z_{sr} : s = 1, \dots, \binom{n}{2}\}$ is a stationary mixing random sequence satisfying $\sup_r \mathbf{E} \|z_{1r}\|^{p+\delta} < \infty$, with (i) $\sup_r \alpha^r(m) \leq Ca^m$ for $a \in (0, 1)$ or (ii) $\sup_r \alpha^r(m) = o(m^{-d})$; (2) $\epsilon_{sr} | A_r^{s-1}, z_r^{s-1}$ are iid $\mathcal{N}(0, 1)$.

PROPOSITION 4.3. Assumption 7 implies the conditions of Lemma 4.1.

In this example, randomly sampled data such as the induced subgraph, star subgraph, or even randomly sampled edges are easy to accommodate. The key take-away here is not that this example is a particularly good model of network formation, but rather that it demonstrates how interdependencies in edge formation can be accommodated.

4.5. Discussion of Boundaries. In this section we have developed a general method to consistently estimate the economic parameter using graphical reconstruction. The method allows the researcher to estimate network effects using a general set of network statistics, such as eigenvector centrality, where no analytical corrections are available. We now discuss how the effectiveness of graphical reconstruction may vary with network formation models, statistics of interest, and misspecification.

First, we may be interested in how misspecification of the network formation model affects graphical reconstruction. In practice, we do not know the family of models which generated the empirical networks. Clearly, misspecification is problematic only to the extent of the covariance between the conditional expectation of the misspecified model and its deviation from the true model. While this is not easy to analytically characterize, it does suggest that the model one needs relates to the network statistic one is interested in studying. For instance, graphical reconstruction with Erdos-Renyi style models may be sufficient to study questions pertaining to the degree distribution, but may perform poorly if one is interested in clustering. Numerical simulations confirm precisely this intuition, suggesting that chosen models ought to be a function of the statistic of interest.

Second, if the model is consistently estimable, one may ask whether a dense versus sparse version of the model will exhibit better convergence properties. A dense graph has expected average degree growing asymptotically and a sparse graph has bounded degree. While dense graphs have more edge-decisions than the sparse counterparts, it turns out that such a statement cannot be made.⁴¹

⁴¹Consider two examples with $\text{Ed}(G) = \Theta(1)$. The first example in general is an Erdos-Renyi graph with $p_n = \delta n^{-\gamma}$, $\gamma \in [1, 2)$, which is sparse when $\gamma = 1$. The second is a block model where a node has d links in expectation, there are k_n blocks with m nodes per block, and nodes only connect to other nodes in the same block with iid probability

Third, models of conditional edge independence where we flexibly introduce covariates can capture substantial co-variation in network formation. Bickel and Chen (2009) present a discussion of the literature where any probability distribution on an infinite undirected graph can be represented as edges emerging conditionally independently given some latent variables and discuss how group based models are natural parametric models that approximate the nonparametric models which can capture any network formation process. In practice we lack these latent variables and therefore conditional edge independent models do poorly precisely for statistics such as clustering, raising the necessity for models with network effects.

5. NUMERICAL EXPERIMENTS

This section reports the results of numerical simulations that characterize the biases due to sampling as well as the behavior of the analytical and graph reconstruction estimators.

5.1. Simulation Setup. We specify a data-generating process for a set of random graphs and outcome data, and then carrying out the following steps.

ALGORITHM (Simulation).

- (1) Generation of data.
 - (a) Draw R networks the network formation families (below).
 - (b) Generate outcome data from a model with β_0 and data-generating process $(y, \epsilon)|G; \beta_0$.
 - (c) For each graph G_r construct sampled graphs $\{G_{rb}^S, G_{rb}^{lS} : b = 1, \dots, B\}$.
- (2) Estimation of $\hat{\beta}$ using $\{G_{rb}^S, G_{rb}^{lS} : b = 1, \dots, B\}$.
 - (a) Estimate $\hat{\beta}_b(G^{lS})$ and $\hat{\beta}_b(G^S)$.
 - (b) If applicable, estimate the adjusted estimator $\tilde{\beta}_b(G^{lS})$ and $\tilde{\beta}_b(G^S)$.
 - (c) Estimate the graphical reconstruction estimators.
- (3) Perform (1)-(2) for $\psi \in \{1/4, 1/3, 1/2, 2/3\}$.

We generate networks of $n = 250$ nodes using the following simple conditional edge independence model. We set parameters such that the average degree, clustering, path length, maximal eigenvalue and variance of the eigenvector centrality distribution from networks in our simulations mimics those moments in the empirical Indian networks data-set. Dividing the set of nodes into 6 approximately equally sized groups, we place those groups on a line, indexed from 1 to 6. The probability that an edge formed between two members within the same group is high. The probability that an edge formed between two members of two different groups declines in the cross-group distance, represented by the difference in the indexed location of those groups on the line. Formally, let $g(i)$ denote the group of vertex i . We set

$$P(A_{ijr}|z_{ijr}) = z'_{ijr}\theta_{0r}.$$

θ_{0r} is a $\binom{6}{2}$ -vector with elements $\theta_{0r,lm}$ with $1 \leq l < m \leq 6$ and z_{ijr} is a $\binom{6}{2}$ -vector with $z_{ijr,lm}$ with $1 \leq l < m \leq 6$. $\theta_{0r,lm}$ is the probability that a member of group l is linked to a member of group

$p = d/m$. The first model has $\hat{p} - p = O_P(n^{-1-\gamma})$ and the latter has $\hat{p} - p = O_P(n^{-1/2})$. The dense Erdos-Renyi model has $\hat{p} - p = O_P(n^{-1})$.

m . The lm -component of z_{ijr} is a dummy for whether i and j are in groups l and m respectively, $z_{ijr,lm} = \mathbf{1}\{g(i) = l\} \mathbf{1}\{g(j) = m\}$. In order to generate $\theta_{0r,lm}$ we use a simple distance function, with $\theta_{0r,lm} = p_r (1 - |l - m|/6)$ where p_r is a uniform random variable chosen such that the average degree generated mimics the average degree from the empirical application.⁴²

5.2. Regression of Outcomes on Network Characteristics. We simulate and estimate a model with heteroskedastic residuals, $y_{ir} = \alpha_0 + w_{ir}(G_r)\beta_0 + \sigma_0 \cdot u_{ir}$, where $(\alpha_0, \beta_0) = (1, 2)$ and

$$u_{ir} := \mathcal{N}\left(0, \sigma_{ir}/\sqrt{\hat{\mu}_{\sigma_{ir}}^2}\right), \quad \sigma_{ir} := 3 \frac{w_{ir} - w_{ir \min}}{w_{ir \max} - w_{ir \min}} + 0.2, \quad \text{and } \hat{\mu}_{\sigma_{ir}}^2 := \mathbb{E}_{n,R} \left[\sigma_{ir}^2 \right].$$

This formulation creates a fan-like heteroskedasticity. We then can easily set the R^2 of the regression to approximately 0.3 by defining $\sigma_0^2 := (1/R^2 - 1) \cdot \mathbb{E}_{n,R}(\tilde{y}_{ir} - \bar{\tilde{y}}_{ir})^2$ for $\hat{y}_{ir} = \alpha_0 + w_{ir}(G_r)\beta_0$.

Columns 1-5 of Tables 1 and 2 show the estimation bias, in percentages, for regression parameters when using sampled network data for a variety of network statistic regressors. Table 1 shows the biases when estimating regressions at the network level while Table 2 shows the biases when estimating regressions at the node level.

At the network level, we consider average degree, graph clustering, graph span, average path length, and λ_{\max} . In addition, we show results for the standard deviation of the eigenvector centrality distribution and the spectral gap. The eigenvector centrality represents how important a node is in information transmission (Jackson, 2008b) and the spectral gap of a graph characterizes how rapidly diffusion processes on networks spread (Chung, 1997). The latter is closely related to the expansiveness of a network that Ambrus et al. (2010) show characterizes good risk-sharing properties. Hochberg et al. (2007) use eigenvector centrality in applied work.

At the node level, we show results for the degree, clustering coefficient, and eigenvector centrality of a node. Moreover we consider two regressions which characterize how far a node i is from another node j . We select a random node j (corresponding to a randomly treated node in an experimental setting) and generate a regressor which is the path length from i to j . In addition, we partition the nodes into two subsets which communicate the most within themselves and least across the sets. We say i is on the same side of the spectral partition of j if they are in the same subset. This partition is related to the spectral gap (Chung, 1997) and therefore has implications for the Ambrus et al. (2010) approach to characterizing risk-sharing.

Overall we find that sampling the network leads to significant biases. To illustrate this, we discuss the biases at 1/3 sampling for the graph and node level. At the graph level the maximum bias is 260% (λ_{\max}), the mean is 90.9%, and the minimum is 15%.⁴³ The biases include expansion bias in the cases of degree, maximal eigenvalue, spectral gap, and graph clustering (for the star subgraph). The node level regressions exhibit a similar pattern: the maximum bias magnitude is 91%, the mean is 63%, and the minimum is 7% (same side of the spectral partition).

⁴²We have conducted simulations for alternative formation models, such as one in which covariates are generated by an autoregressive process and the edge formation probability is governed by a logistic regression. Results are qualitatively and quantitatively similar.

⁴³When we looking at maximum, mean, and minimum, we are interested in the magnitude of the biases, so our discussion focuses on the absolute value of the bias.

In columns 6-10 of Table 1, we present the results using the analytical corrections from section 3.⁴⁴ We find that the adjusted regression estimators perform uniformly better than the unadjusted estimators. The biases are usually low. Overall, the analytical corrections improve 100% of the parameter estimates in the simulations. At 1/3 sampling, the mean reduction in bias percentage when comparing the analytical correction to the raw network statistic is 69pp with a median reduction of 69pp and a maximum of 243pp. Columns 6-10 of Table 2 applies the analytical corrections which were derived for the case of graph level regression, to node level analysis; the results are not motivated by theory and of course are mixed.⁴⁵

We consider the graphical reconstruction estimators in columns 11-15 of Tables 1 and 2. It nearly uniformly outperforms the estimator using the sampled data alone. Biases are mostly very low across a number of linear and nonlinear network statistics. For illustration we discuss examples with 1/3 sampling: at the graph level the median bias is 5.7%, the minimum is 0.6%, and the reconstruction estimator reduces the bias in 54 of the 56 parameters estimated in Table 1. The mean reduction in bias is 73pp and the maximum reduction is 254pp. We find similar results at the node level. The median bias is 1.4% and graph reconstruction reduces the bias in 100% of the parameters estimated in the table. Furthermore, the median reduction in bias is 62pp with a maximum of 85pp. Not surprisingly, at a given sampling rate reconstruction with G^S performs uniformly better than with G^{IS} . The effective share of edges observed in the star subgraph is $1 - (1 - \psi)^2$ but is ψ^2 in the induced subgraph. Typically 2/3 sampling with G^{IS} (4/9 share of the edges observed) yields a reconstruction procedure which is only as good as 1/4 sampling with G^S (7/16 share of the edges observed).

In Table 4 we study the behavior of significance testing and provide evidence that graphical reconstruction may often increase t -statistics.⁴⁶ Specifically, we present the ratio of the t -statistic under graphical reconstruction to the t -statistic under the naive estimator using the sampled network statistic. We find that at the network level across 86% of the cases the t -statistic increases (48 of 56 estimated parameters) and at the node level across 96% of the cases (46 of the 48 estimated parameters) graphical reconstruction yields a higher t -statistic than the naive estimator. Moreover, we find that the average ratio of the t -statistic of reconstruction to the naive estimator is high.

5.3. Regression of Outcomes on Network Neighbors' Outcomes. Table 3 presents the results for simulations for the model of equation (2.3) with $(\alpha_0, \rho_0, \gamma_0, \delta_0) = (1, 0.5, 2, 0.5)$. We use three specifications to demonstrate the emergence of biases in peer effects regression due to two distinct causes: correlation of the instrument with the errors-in-variables problem and a weak instrument/finite sample problem induced by sampling. The table presents the mean bias percentage as well as the standard error of the bias.

We present three methods of estimating peer effects with sampled data and one correction. First, we show an estimate of the peer effect model with the network given by the induced subgraph.

⁴⁴In Table 8 of Appendix D we show an example of an analytical correction that involves estimating $\hat{\sigma}_v^2$.

⁴⁵The correction working for degree with G^S is mechanical since there is no mismeasurement for d_i with $i \in S$.

⁴⁶Note that this is a numerical result and not a theoretical one. The results may be specific to network formation models and statistics examined.

Second, we present an estimation of the model where (y, x) are known for every node, but the network used is the star subgraph. Finally, we present the same specification but only allow the researcher to have covariates for surveyed nodes. Each specification exhibits biases.

We vary the number of networks and disturbance size across three models to study how the bias varies. In Model 1, we use one network with 250 nodes per simulation drawn from the aforementioned model and set the number of simulations to 50,000. In Model 2, we use one network with 250 nodes per simulation and use 50,000 simulations, but reduce the variance of the disturbance. Model 3 presents results from 2,500 simulations of 20 networks each with 250 nodes, drawn from the model.⁴⁷

All three specifications show significant bias in the estimates of the endogenous and exogenous peer effects. Comparing Panels A and B of the first and second set of columns shows the biases are greater when there is more noise in the system. Moreover, comparing Panels A and B of columns 1-5 and 11-15 shows that increasing the number of graphs in the estimation from 1 to 20 only modestly reduces the bias due to sampling. There are still non-trivial biases which remain.

Overall, the analytical correction performs well. In the Models 2 and 3, the estimates are essentially unbiased across all sampling levels presented. Moreover, the analytical correction for Model 1 exhibits negligible bias for sampling rates of 2/3 and 1/2. However, biases emerge at very low sampling rates, 1/3 and 1/4, in the case of Model 1. Furthermore, as evidenced by the standard errors at 1/4, the estimates are extremely unstable.

To measure whether there is a weak instruments problem, in Panel C we display a generalization of the concentration parameter of the first stage, allowing for interdependence in the variance following Kleibergen (2007).⁴⁸ The intuition is that in these networks, even for the analytical correction there is measurement error in the instrument. Since the number of connections to neighbors and second neighbors in a network is low, the amount of noise in the first stage increases.⁴⁹ Panel C shows that the concentration parameter is very low for the first stage estimates in Model 1, especially at low sampling levels. Moreover, once the number nodes in the network is high enough or the amount of independent data (20 networks) is high enough, the concentration parameter is extremely high. In these cases our analytical correction removes the bias entirely while biases remain with the sampled estimators.

5.4. A Model of Diffusion. We numerically study the Jackson and Rogers (2007b) model of diffusion and present the results in Table 5. In Panel A we use the aforementioned simulated network data to generate a model with $\beta_0 = e^{-2}$.⁵⁰ Columns 1-5 present evidence of severe expansion bias in the estimates $\hat{\beta}$ when using sampled data. At 1/3 sampling, the transmission parameter is

⁴⁷The number of simulations was chosen to roughly equate the computation time, on the order of $n^4 \cdot \#$ of simulations, for each of the three specifications.

⁴⁸For a first stage $X = Z\pi + v$, we use the generalized concentration parameter $\mu^2 := \pi' \Sigma_\pi^{-1} \pi$ where $\hat{\pi} = (Z'Z)^{-1} Z'X$ and $\Sigma_\pi = \text{var}(\hat{\pi})$.

⁴⁹The extent to which this matters can be seen by noticing that the concentration parameter is 2 for $\psi = 1/4$, while if the signal to noise ratio had stayed the same in the first stage, the concentration parameter should have only decreased from 16 to 4.

⁵⁰This choice was motivated by Jackson and Rogers (2007b) who numerically show this corresponds to a 20% steady-state rate of diffusion. This matches the microfinance take-up rate in our empirical application.

overestimated by 250% when we study the induced subgraph and 85% when we turn to the star subgraph. Columns 6-10 presents the graphical reconstruction results; the procedure removes the entire bias.

5.5. Robustness to Misspecification. To investigate how well the procedure works with empirical data, where we do not know the data generating processes, we conduct numerical experiments using the networks of the [Banerjee et al. \(2011\)](#) data-set, described in greater detail in section 6. We repeat the simulation algorithm of section 5.1 with the only difference coming in step 1(a). Instead of generating networks from the aforementioned model, we take 50 independent draws with replacement from the [Banerjee et al. \(2011\)](#) data-set.⁵¹ When we fit a network formation model in step 2(c), we use the model given by (4.2). We use as covariates the GPS distance between households as well as the difference in the number of rooms, beds, roofing material type, and electricity access. Table 6 presents summary statistics from graphical reconstruction exercises analogous to those of Tables (1) and (2). We find that graphical reconstruction reduces the bias in 98% of the network statistics when using the induced subgraph and 100% when using the star subgraph. In addition, the median bias is 9% with the star subgraph when using reconstruction with a median reduction of bias of 23pp. Similarly, the median bias is 32% with the induced subgraph and the median reduction in bias is 32pp.

Panel B of Table 5 presents the results from numerical experiments done for the [Jackson and Rogers \(2007b\)](#) model of diffusion using the empirical networks instead of simulated networks. We find that at 1/3 sampling graphical reconstruction yields biases of 5% and 8% for the star and induced subgraphs, respectively. Taken together, the results of these exercises suggest that even when allowing for network formation model misspecification, graphical reconstruction typically outperforms what the researcher otherwise would have estimated.

6. EMPIRICAL APPLICATION: DIFFUSION OF MICROFINANCE

This section presents an empirical application using data from [Banerjee, Chandrasekhar, Dufflo, and Jackson \(2011\)](#), which studies how households' decisions to participate in microfinance diffuses through village networks. We use detailed demographic and social network data in 43 villages in Karnataka, India, which range from a 1.5 to 3 hour's drive from Bengaluru. The data was collected six months before a microfinance institution started its operation in those villages. The networks are randomly sampled at ~46%.

The key outcome variable is the microfinance take-up decisions of households in the network. Information about microfinance access is typically spread by members and the MFI has administrative data which allows us to observe the diffusion of membership. Identification is based on the principle that the MFI followed the same procedure in informing villagers about microfinance in each village. The MFI identified a collection of pre-set individuals in the village (*anganwadi* teachers, shop keepers, etc.), informed them about the program in a private meeting, and then

⁵¹We treat the networks as if they are fully-observed. In step 1(c) of the algorithm we sample each graph at rate ψ . The authors of [Banerjee et al. \(2011\)](#) currently are obtaining a 100% network sample in a resurvey.

asked them to invite individuals to an information session. [Banerjee et al. \(2011\)](#) contend that this scheme provides arguably exogenous variation in the centrality of those households.

To account for the partial sampling, we assume that an edge forms between a pair of households conditionally independently, given a set of covariates (GPS coordinate Euclidean distance between the two households and the difference in the number of beds, number of rooms, electricity access, and roofing material of the two households). We estimate the model separately on each village using a logistic regression in which the observed data between two households are coded as 1 (connected) and 0 (not connected).

Panel A of [Table 7](#) reports estimates of village-level regressions where the microfinance take-up rate in a village is regressed on network characteristics. Columns 1-4 presents regressions of microfinance take-up on network statistics, suggested by diffusion theory to be associated with take-up. Column 1 shows the regression of take-up rate on the average eigenvector centrality of the set of initially informed households. Diffusion theory suggests that eventual take-up of microfinance ought to be higher when the first people to be informed are more central. The increase of the average centrality in the set of nodes by 0.1 corresponds to a 16.3pp increase in take-up rate when using the sampled data; graph reconstruction places this estimate as a 24.2pp increase in take-up rate. If the initially informed households were from the 75th percentile of the centrality distribution as compared to the 25th percentile, this represents a 7.5pp increase in microfinance take-up when estimated using reconstruction as compared to a 4.5pp increase when using the sampled data. Recalling that the average take-up rate is 18.49%, this suggests that sampling the network causes significant under-estimation of the network effect. Column 2 presents the regression of take-up on the average path length. If it takes one extra step on average to traverse the graph, this corresponds to a 5.4pp decrease in take-up of microfinance, according to the sampled network, though reconstruction suggests that the estimate ought to be a 9.3pp decrease (with a t-statistic of 1.56). Furthermore, consistent with the results of [Table 4](#), the t-statistics associated with the estimates typically increase after reconstruction, suggesting that the researcher can better detect anticipated effects with reduced measurement error.

In Panel B, we turn to household-level regressions where whether a household joins microfinance is regressed on network characteristics. Column 1 reports a regression on the eigenvector centrality of a node. Graphical reconstruction only yields a modest change in this example: an increase in 0.1 of a household's eigenvector centrality corresponds to a 5.5pp increase in take-up likelihood using the sampled data and a 6.6pp increase in take-up likelihood using graphical reconstruction. Column 2 provides a more stark example in a regression of take-up on the inverse social distance of a household to the set of initially informed households. The network effect more than doubles when using graphical reconstruction. Being distance 1 versus 4 increases the probability of joining microfinance by 3.4pp under graphical reconstruction but only 1.6pp using the star graph. In addition, the point estimate is not statistically significant at conventional levels (with a t-statistic of 0.9), but graphical reconstruction establishes that zero is nearly excluded from a 90% confidence interval (with a t-statistic of 1.62).

Finally, in Panel C we consider the regression of a household's decision to join microfinance or not on the sum of its neighbors' decisions.⁵² The of ρ estimate corresponds to the impact of one neighbor joining microfinance on the probability that a household joins microfinance. Column 2 displays the parameter estimate of the effect of the exogenous covariate, γ , whether a household is initially informed about microfinance. Column 3 displays the parameter estimates of exogenous network effect, δ . This estimate describes the impact of one extra neighbor being a initially informed on a household's likelihood of joining microfinance. We focus on column 1 as the endogenous network effect is the key parameter of interest. The star subgraph data suggests that a one neighbor's take-up corresponds to a 2.7pp decrease in the likelihood of a household joining microfinance. Meanwhile, the star subgraph data where we use the microfinance data and injection point data only for sampled households suggests that a one more neighbor's take-up corresponds to a 4.7pp increase in the likelihood of a household taking up. Finally, the analytical correction shows that a one neighbor's take-up corresponds to a 7pp decrease in likelihood of take-up by a household. Therefore, the sampled data has lead to severe under-estimation and even sign-switching of the endogenous network effect of interest. In particular, partial sampling may cause the researcher to mistake a substituting peer effect for a complementary peer effect. The remainder of the table suggests that, in addition, the exogenous peer effect is also under-estimated and, in the sampled data cases, the effects seem to load on the exogenous own covariate coefficient.

7. USING THE RESULTS TO BETTER COLLECT DATA

In this section we discuss how researchers can adopt our framework to think about data collection. The question we are interested in is: given that a researcher faces a budget constraint and needs to trade off the sampling rate and the number of networks in her sample, is there a method by which she can assess the trade-off?

We suppose that a researcher is interested in estimating a coefficient in a regression of an outcome on a network statistic. Assume that the researcher has a project budget b and a pilot budget p . Each village has a fixed cost f associated with the survey as well as a variable cost c for sampling.⁵³ We assume that the cost to sample individuals is linear and therefore the cost to sample a ψ -sample of the village is $c\psi n$. Finally, let \bar{R} be the maximum number of villages available to study.

We posit that the researcher is interested in minimizing mean-squared error (MSE) in the estimation of β_0 .⁵⁴ The relevant program is⁵⁵

$$(7.1) \quad \min_{\psi \in [0,1], R \leq \bar{R}} \text{MSE}(\psi, R) \text{ s.t. } (c\psi n + f) R \leq b.$$

At the optimum $\psi = (b/R - f)/(cn)$ and therefore we may consider the concentrated objective function $\text{MSE}(R) = \text{MSE}(\psi(R), R)$. The researcher may estimate the MSE by fully sampling a

⁵²The estimating approach ignores problems raised by a discrete dependent variable, following the approach taken in this literature (e.g., Bramoulle and Kranton, 2007; Gaviria and Raphael, 2001; Sacerdote, 2001). The estimated standard errors handle the heteroskedasticity of the binary response variable.

⁵³This method can be applied to richer budgeting frameworks.

⁵⁴Researchers can replace this with an objective function of their choosing.

⁵⁵For formal asymptotics we may have to let $b = b_n$ grow such that $\frac{b_n}{Rn} \rightarrow k$ some positive constant.

small number of networks and hypothesizing β_0 and R^2 from the linear regression, in a manner analogous to performing power calculations by positing effect sizes and intra-cluster correlations before conducting a field experiment (e.g., Duflo et al., 2007). A researcher first randomly selects k of the \bar{R} graphs using the pilot budget, where $k = p/(cn + f)$. Then, using these k networks, the researcher conducts a numerical experiment, sampling them at different rates and applying graphical reconstruction to estimate the MSE. By doing this, she can select the optimal ψ and R .
 ALGORITHM (Research Design).

- (1) Pick network statistics and network-based hypotheses to test.
- (2) Hypothesize β , R^2 , and generate outcome variable.
- (3) Randomly sample $k = p/(cn + f)$ out of \bar{R} villages and obtain entire networks.
- (4) For each $R \in \{\underline{R}, \dots, \bar{R}\}$
 - (a) Randomly draw, with replacement, R villages from the collection of k networks.
 - (b) Estimate $\text{MSE}(R)$ using the sample and hypothesized parameter values.
 - (i) Sample each of the R village networks at rate $\psi(R) = (b/R - f)/(cn)$.
 - (ii) Apply graphical reconstruction to estimate $\hat{\beta}_{\text{ols}}$ using outcome variable from (2).
 - (iii) Repeat 4(a) and 4(b).i-ii for B simulations.
- (5) Pick $R^* \in \text{argmin}_R \widehat{\text{MSE}}(R)$ and pick $\psi^* = \psi(R^*)$.

The algorithm enables the researcher to estimate the trade-off she faces, given her interest in specific network effects and the distribution of graphs in her region of study. We conduct a simulation exercise to demonstrate this procedure. We set $b = \$152,400$, $f = \$1200$, $c = \$12$, $n = 200$, $\bar{R} = 150$ and assume that the networks are drawn from the family described in section 5 and the empirical Indian networks.⁵⁶ We consider a grid of $R \in \{33, 40, 50, 60, 70\}$ and $\psi \in \{1, 0.7, 0.4, 0.2, 0.1\}$.

Figure 4 displays results for two node-level statistics, eigenvector centrality and clustering, as well as a network-level statistic, the maximal eigenvalue of the adjacency matrix (λ_{max}). We repeat the exercise for both our simulated network data as well as the Indian networks. The figure shows $\text{MSE}(\psi(R), R)$ for sampled networks and graphical reconstruction. It also displays a theoretical lower bound on MSE by plotting the MSE corresponding to using R graphs sampled at 100% instead of at $\psi(R)$. Of course, we find that MSE increases greatly as we move away from 100% sampling and use the raw sampled data. Next we turn to graphical reconstruction and focus on the star subgraph. Looking at the Indian networks, eigenvector centrality has the lowest MSE at 40% sampling while clustering has an optimum at 70%. For these statistics, the simulated networks give 100% as the optimum. Meanwhile, λ_{max} has the lowest MSE at 100% sampling with the Indian networks but 20% sampling is the optimum in the simulated networks.

Taken together, the results suggest that, first, performing graphical reconstruction is very important, even with model misspecification as the researcher will not know the true families generating the empirical networks. Second, the MSE-minimizing sampling rate depends greatly on parameters, the network family, and the statistic of interest. It is difficult, if not impossible, to say *ex ante* where the optimum lies, and systematic procedures that depend on the setting may be better than

⁵⁶The numbers are motivated from Banerjee et al. (2011).

rules of thumb. Third, the results push against the prevailing habit of researchers to obtain more cluster-units (e.g., villages) at lower sampling rates when conducting cluster-level analysis. Our results suggest that, at times, just obtaining better data with fewer cluster-units may be worthwhile. Though it is not surprising that network-level statistics exhibit higher levels of MSE, as there are only R as opposed to nR observations, this says nothing about the trade-off between the sampling rate and number of villages.

8. CONCLUSION

Applied social network analysis often use graphs constructed from data collected from a partial sample of nodes. Even when nodes are selected randomly, the partial sampling induces non-classical measurement error and consequently biases estimates of regression coefficients and GMM parameters. Moreover, these biases are of unclear sign and magnitude. We analytically examine the biases in the estimation of a number of network-based regression and GMM models with applications to a variety of economic environments. To address the problem in general, we develop a consistent and asymptotically normal method to estimate the economic parameters using graphical reconstruction, while allowing for substantial heterogeneity across networks. Specifically, the method allows for every network in the sample to be generated by a different model.

We conclude that network-based applied work must proceed cautiously, paying close attention to network data quality. From an applied perspective, researchers should be careful to work either with specifications which provide conservative results when facing sampled data or implement bias correction procedures if possible. Moreover, researchers ought to address the bias problem *ex ante*, either by choosing a unit of study where more complete data is available or using graphical reconstruction to understand how mean-squared error may vary with the sampling rate. Undoubtedly, the performance of graphical reconstruction with empirical network data will only improve as the burgeoning literature on consistently estimable network formation models matures. To that end, from a theoretical perspective the lacuna in the literature is the absence of network formation models that both allow for higher-order dependencies in link formation and are also consistently estimable.

REFERENCES

- ALATAS, V., A. BANERJEE, A. G. CHANDRASEKHAR, B. OLKEN, AND R. HANNA (2011): “Identifying the Poor: Social Networks and the Aggregation of Information,” *MIT Working Paper*.
- AMBRUS, A., M. MOBIUS, A. SZEIDL, AND L. CENTER (2010): “Consumption risk-sharing in social networks,” *NBER Working Paper*.
- ANDREWS, D. (1992): “Generic uniform convergence,” *Econometric Theory*, 8, 241–257.
- (1994): “Asymptotics for semiparametric econometric models via stochastic equicontinuity,” *Econometrica: Journal of the Econometric Society*, 62, 43–72.
- ANGELUCCI, M., G. DE GIORGI, M. RANGEL, AND I. RASUL (2009): “Family networks and school enrolment: Evidence from a randomized social experiment,” *Journal of Public Economics*.

- ARAL, S. AND D. WALKER (Forthcoming): “Creating social contagion through viral product design: A randomized trial of peer influence in networks,” *Management Science*.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. JACKSON (2011): “The Diffusion of Microfinance,” MIT working paper.
- BARABASI, A. AND R. ALBERT (1999): “Emergence of scaling in random networks,” *Science*, 286, 509.
- BICKEL, P. AND A. CHEN (2009): “A nonparametric view of network models and Newman–Girvan and other modularities,” *Proceedings of the National Academy of Sciences*, 106, 21068.
- BLOCH, F., G. GENICOT, AND D. RAY (2008): “Informal insurance in social networks,” *Journal of Economic Theory*, 143, 36–58.
- BLUME, L., W. BROCK, S. DURLAUF, AND Y. IOANNIDES (2011): “Identification of social interactions,” *Handbook of Social Economics*, 1, 855–966.
- BOLLOBÁS, B., C. BORGS, J. CHAYES, AND O. RIORDAN (2010): “Percolation on dense graph sequences,” *The Annals of Probability*, 38, 150–183.
- BRAMOULLE, Y., H. DJEBBARI, AND B. FORTIN (2009): “Identification of peer effects through social networks,” *Journal of Econometrics*, 150, 41–55.
- BRAMOULLE, Y. AND R. KRANTON (2007): “Risk-sharing networks,” *Journal of Economic Behavior & Organization*, 64, 275–294.
- CHANDRASEKHAR, A. G. AND M. O. JACKSON (2011): “Tractable and consistent exponential random graph models,” *Mimeo*.
- CHATTERJEE, S., P. DIACONIS, AND A. SLY (2010): “Random graphs with a given degree sequence,” *arXiv Preprint: 1005.1136*.
- CHERNOZHUKOV, V., I. FERNANDEZ-VAL, AND B. MELLY (2009): “Inference on counterfactual distributions,” .
- CHRISTAKIS, N., J. FOWLER, G. IMBENS, AND K. KALYANARAMAN (2010): “An Empirical Model for Strategic Network Formation,” *NBER Working Paper*.
- CHUNG, F. (1997): *Spectral Graph Theory*, American Mathematical Society.
- CONLEY, T. AND C. UDRY (2010): “Learning about a new technology: Pineapple in Ghana,” *The American Economic Review*, 100, 35–69.
- DE GIORGI, G., M. PELLIZZARI, AND S. REDAELLI (2010): “Identification of social interactions through partially overlapping peer groups,” *American Economic Journal: Applied Economics*, 2, 241–275.
- DE JONG, R. AND T. WOUTERSEN (2010): “Dynamic time series binary choice,” *Econometric Theory*, 1, 1–30.
- DEMARZO, P., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, and Unidimensional Opinions,” *Quarterly Journal of Economics*, 118, 909–968.
- DIACONIS, P. AND D. FREEDMAN (1984): “Partial exchangeability and sufficiency,” *Statistics: Applications and New Directions*, 205–236.
- DUDLEY, R. (2002): *Real analysis and probability*, vol. 74, Cambridge Univ Pr.

- DUFLO, E., R. GLENNERSTER, AND M. KREMER (2007): “Using randomization in development economics research: A toolkit,” *Handbook of Development Economics*, 4, 3895–3962.
- FRANK, O. (1980): “Sampling and inference in a population graph,” *International Statistical Review/Revue Internationale de Statistique*, 33–41.
- (1981): “A survey of statistical methods for graph analysis,” *Sociological methodology*, 12, 110–155.
- GHANI, A., C. DONNELLY, AND G. GARNETT (1998): “Sampling biases and missing data in explorations of sexual partner networks for the spread of sexually transmitted diseases,” *Statistics in Medicine*, 17, 2079–2097.
- GOEREE, J., M. MCCONNELL, T. MITCHELL, T. TROMP, AND L. YARIV (2010): “The 1/d law of giving,” *American Economic Journal: Microeconomics*, 2, 183–203.
- GOLDSMITH-PINKHAM, P. AND G. IMBENS (2011): “Determinants of Social Networks: Large Sample Theory,” *Harvard Working Paper*.
- GOLUB, B. AND M. JACKSON (2009): “How homophily affects learning and diffusion in networks,” *Working Paper*.
- (2010a): “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2, 112–149.
- (2010b): “Using selection bias to explain the observed structure of Internet diffusions,” *Proceedings of the National Academy of Sciences*, 107, 10833.
- GRANOVETTER, M. S. (1973): “The Strength of Weak Ties,” *The American Journal of Sociology*, 78, 1360–1380.
- HAHN, J. AND G. KUERSTEINER (2004): “Bias reduction for dynamic nonlinear panel models with fixed effects,” *Econometric Theory*, 1, 1–40.
- HAHN, J. AND W. NEWEY (2004): “Jackknife and analytical bias reduction for nonlinear panel models,” *Econometrica*, 72, 1295–1319.
- HALL, P. AND J. HOROWITZ (1996): “Bootstrap critical values for tests based on generalized-method-of-moments estimators,” *Econometrica: Journal of the Econometric Society*, 891–916.
- HANDCOCK, M. AND K. GILE (2010): “Modeling social networks from sampled data,” *The Annals of Applied Statistics*, 4, 5–25.
- HOCHBERG, Y., A. LJUNGQVIST, AND Y. LU (2007): “Whom you know matters: Venture capital networks and investment performance,” *The Journal of Finance*, 62, 251–301.
- HUISMAN, M. (2009): “Imputation of missing network data: Some simple procedures,” *Journal of Social Structure*, 10.
- IYER, R. AND M. PURI (forthcoming): “Understanding bank runs: the importance of depositor-bank relationships and networks,” *American Economic Review*.
- JACKSON, M. (2008a): “Average distance, diameter, and clustering in social networks with homophily,” *Internet and Network Economics*, 4–11.
- (2009a): “Networks and economic behavior,” *Annu. Rev. Econ.*, 1, 489–513.
- (2009b): “An Overview of Social Networks and Economic Applications,” *The Handbook of Social Economics.*, ed. Jess Benhabib, Alberto Bisin and Matthew O. Jackson. Elsevier Press.

- forthcoming*.
- JACKSON, M. AND B. ROGERS (2007a): “Meeting strangers and friends of friends: How random are social networks?” *American Economic Review*, 97, 890–915.
- (2007b): “Relating network structure to diffusion properties through stochastic dominance,” *The BE Journal of Theoretical Economics*, 7, 1–13.
- JACKSON, M. AND L. YARIV (2007): “Diffusion of Behavior and Equilibrium Properties in Network Games,” *American Economic Review*, 97, 92–98.
- JACKSON, M. O. (2008b): *Social and Economic Networks*, Princeton University Press.
- JACKSON, M. O., T. RODRIGUEZ-BARRAQUER, AND X. TAN (Forthcoming): “Social Capital and Social Quilts: Network Patterns of Favor Exchange,” *American Economic Review*.
- JENISH, N. AND I. PRUCHA (2009): “Central limit theorems and uniform laws of large numbers for arrays of random fields,” *Journal of Econometrics*, 150, 86–98.
- KARLAN, D., M. MOBIUS, T. ROSENBLAT, AND A. SZEIDL (2009): “Trust and Social Collateral,” *Quarterly Journal Of Economics*, 124, 1307–1361.
- KINNAN, C. AND R. TOWNSEND (2011): “Kinship and Financial Networks, Formal Financial Access and Risk Reduction,” *Northwestern University and MIT Working Paper*.
- KLEIBERGEN, F. (2007): “Generalizing weak instrument robust IV statistics towards multiple parameters, unrestricted covariance matrices and identification statistics,” *Journal of Econometrics*, 139, 181–216.
- KOLOTILIN, A. (2011): “Estimation of a Scale-Free Network Formation Model,” *MIT Working Paper*.
- KOSSINETS, G. (2006): “Effects of missing data in social networks,” *Social Networks*, 28, 247–268.
- KREMER, M. AND E. MIGUEL (2007): “The Illusion of Sustainability,” *The Quarterly Journal of Economics*, 122, 1007–1065.
- LAHIRI, S. N. (1992): “Edgeworth Correction by ‘Moving Block’ Bootstrap for Stationary and Nonstationary Data,” in *Exploring the Limits of Bootstrap*, ed. by R. Lepage and L. Billard, Wiley-Interscience, vol. 270, 183–214.
- LEIDER, S., M. M. MÖBIUS, T. ROSENBLAT, AND Q.-A. DO (2009): “Directed Altruism and Enforced Reciprocity in Social Networks,” *Quarterly Journal of Economics*, 124, 1815–1851.
- LIBEN-NOWELL, D. AND J. KLEINBERG (2008): “Tracing information flow on a global scale using Internet chain-letter data,” *Proceedings of the National Academy of Sciences*, 105, 4633.
- LIN, C., Z. LIN, AND C. LU (1996): *Limit theory for mixing dependent random variables*, vol. 378, Springer.
- MACOURS, K. (2003): “Comparing a Direct with an Indirect Approach to Collecting Household Level data: Who tells the truth about what?” *Washington, DC, United States: Johns Hopkins University. Mimeographed document*.
- MANSKI, C. (1993): “Identification of endogenous social effects: The reflection problem,” *The Review of Economic Studies*, 60, 531–542.
- MÖBIUS, M. AND A. SZEIDL (2006): “Trust and cooperation in social networks,” *Harvard University and UC Berkeley Working Paper*.

- NEWMAN, M., S. STROGATZ, AND D. WATTS (2001): “Random graphs with arbitrary degree distributions and their applications,” *Physical Review E*, 64, 26118.
- ROTHENBERG, R. (1995): “Commentary: sampling in social networks,” *Connections*, 18, 104–110.
- RUBIN, D. (1976): “Inference and missing data,” *Biometrika*, 63, 581.
- SANTOS, P. AND C. BARRETT (2008): “What do we learn about social networks when we only sample individuals? Not much.” *Working Paper*.
- SNIJDERS, T. (1992): “Estimation on the basis of snowball samples: How to weight?” *Bulletin de méthodologie sociologique*, 36, 59.
- STEPHAN, F. (1945): “The expected value and variance of the reciprocal and other negative powers of a positive Bernoullian variate,” *The Annals of Mathematical Statistics*, 16, 50–61.
- STUMPF, M., C. WIUF, AND R. MAY (2005): “Subnets of scale-free networks are not scale-free: sampling properties of networks,” *Proceedings of the National Academy of Sciences of the United States of America*, 102, 4221.
- TAKASAKI, Y., T. BARHAM OLIVER, AND L. BRADFORD (2000): “Rapid rural appraisal in humid tropical forests: an asset possession-based approach and validation methods for wealth assessment among forest peasant households,” *World Development*, 28, 1961–1977.
- THOMPSON, S. (2006): “Adaptive web sampling,” *Biometrics*, 62, 1224–1234.
- TOWNSEND, R. M. (2007): “Townsend Thai Project Household Annual Resurvey, 1998-2006,” [Computer file]. 2nd Release. Chicago, IL: Social Sciences Computing Services, the University of Chicago [Producer & distributor].
- VAN DER VAART, A. AND J. WELLNER (1996): *Weak convergence and empirical processes*, Springer Verlag.
- WATTS, D. AND S. STROGATZ (1998): “Collective dynamics of small-world networks,” *Nature*, 393, 440–442.

TABLE 2. Node Level Regressions (Simulated Network Data)

Statistic	Type	Output	Raw Network Statistics					Experimental Corrections					Graphical Reconstruction					
			1/4	1/3	1/2	2/3	1	1/4	1/3	1/2	2/3	1	1/4	1/3	1/2	2/3	1	
Degree	G ^S	Bias %	-60.60%	-56.56%	-32.33%	-13.66%	0.03%	-0.47%	-0.52%	-0.62%	-0.43%	-0.43%	-0.43%	-0.85%	-0.93%	-0.39%	-0.27%	0.03%
	G ^S	Bias %	106.09%	91.04%	53.03%	32.06%	-0.06%	-48.58%	-42.37%	-23.29%	-13.99%	0.36%	0.36%	-7.58%	-6.36%	-2.00%	-1.07%	-0.06%
Clustering	G ^S	Bias %	-80.68%	-75.74%	-44.20%	-17.09%	-0.01%	-89.75%	-85.63%	-55.18%	-25.47%	-0.06%	-0.06%	-1.55%	-1.35%	-0.63%	-0.34%	-0.01%
	G ^S	Bias %	-83.85%	-79.88%	-59.12%	-39.37%	-0.39%	-83.95%	-80.12%	-59.20%	-36.04%	0.02%	0.02%	-14.40%	-9.62%	-3.75%	-2.01%	-0.39%
Path Length	G ^S	Bias %	-68.03%	-55.98%	-22.32%	-9.35%	0.40%							-3.01%	-2.94%	-1.16%	-0.26%	0.40%
	G ^S	Bias %	-100.66%	-88.82%	-51.77%	-30.64%	-0.12%							-27.64%	-20.26%	-7.30%	-2.98%	-0.12%
Eig Centrality	G ^S	Bias %	-60.74%	-54.46%	-30.00%	-14.56%	0.19%							-2.02%	-1.17%	-0.50%	-0.14%	0.19%
	G ^S	Bias %	-89.42%	-85.81%	-69.14%	-50.59%	0.32%							-4.38%	-3.81%	-3.04%	-1.23%	0.32%
Same Side	G ^S	Bias %	-9.97%	-6.89%	-2.27%	-0.88%	-0.17%							-0.23%	0.04%	-0.10%	0.22%	-0.17%
	G ^S	Bias %	-45.60%	-32.78%	-7.66%	-3.90%	0.23%							0.17%	-0.13%	0.07%	-0.19%	0.23%

Notes: Bias % denotes $(\beta - \beta_0)/\beta_0$, x 100% and coverage is available upon request. All reported mean bias % simulations have standard errors less than 0.65% with most having estimated standard errors of the mean bias % between 0.10% and 0.20% (percentage points). A complete table including simulation precision is available from the authors upon request. Estimation was performed using 50 homophily graphs with 200 nodes, 10 cliques of equal size, and expected degree of ~15, as described in the text. Each simulation was performed 100 times. To compute the conditional expectation in graphical reconstruction 40 simulations were performed. Several standard errors were computed. Coverage was computed using block bootstrap standard errors; clustered standard errors yield very similar coverage rates.

TABLE 3. Peer Effects Regressions (Simulated Graphs)

	Model 1: 1 Graph; $\sigma_e=1$; 50k Sims.					Model 2: 1 Graph; $\sigma_e=1/4$; 50k Sims.					Model 3: 20 Graphs; $\sigma_e=1$; 2,500 Sims.				
	1/4	1/3	1/2	2/3	1	1/4	1/3	1/2	2/3	1	1/4	1/3	1/2	2/3	1
<i>Panel A: 2SLS Mean Bias %</i>	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]	[11]	[12]	[13]	[14]	[15]
G^S m obs.	1606%	-146%	-56%	-40%	-1%	-107%	-78%	-56%	-37%	0%	-56%	-46%	-32%	-19%	0%
	[1682%]	[54%]	[2%]	[0%]	[0%]	[39%]	[2%]	[0%]	[0%]	[0%]	[1%]	[1%]	[0%]	[0%]	[0%]
G^S (y,x)-census, n obs.	-69%	[0%]	[0%]	[0%]	[0%]	-69%	-55%	-27%	-11%	0%	-35%	-22%	-5%	0%	0%
	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]
G^S (y,x)-sample, m obs.	553%	404%	32%	54%	-2%	-81%	900%	288%	-267%	0%	432%	382%	140%	62%	0%
	[776%]	[442%]	[328%]	[51%]	[0%]	[302%]	[1848%]	[145%]	[443%]	[0%]	[147%]	[19%]	[1%]	[1%]	[0%]
G^S Analytic Correction; (y,x)-census, m obs.	18%	-19%	-3%	-2%	-1%	-1%	-1%	0%	0%	0%	-1%	0%	0%	0%	0%
	[178%]	[9%]	[1%]	[0%]	[0%]	[1%]	[0%]	[0%]	[0%]	[0%]	[1%]	[0%]	[0%]	[0%]	[0%]
<i>Panel B: 2SLS Mean Bias %</i>															
G^S m obs.	-3692%	70%	-58%	-34%	1%	-30%	-62%	-53%	-36%	0%	-125%	-122%	-102%	-73%	0%
	[3602%]	[111%]	[5%]	[1%]	[1%]	[75%]	[3%]	[1%]	[0%]	[0%]	[1%]	[1%]	[1%]	[1%]	[0%]
G^S (y,x)-census, n obs.	-86%	[1%]	[1%]	[1%]	[1%]	-83%	-78%	-58%	-28%	0%	-150%	-143%	-104%	-51%	0%
	[1%]	[1%]	[1%]	[1%]	[1%]	[0%]	[0%]	[0%]	[0%]	[0%]	[1%]	[1%]	[0%]	[0%]	[0%]
G^S (y,x)-sample, m obs.	-900%	-905%	-113%	-113%	2%	194%	-1631%	-614%	584%	0%	-913%	-800%	-295%	-129%	0%
	[1486%]	[935%]	[701%]	[108%]	[1%]	[606%]	[3642%]	[321%]	[961%]	[0%]	[304%]	[39%]	[3%]	[1%]	[0%]
G^S Analytic Correction; (y,x)-census, m obs.	-41%	38%	5%	2%	1%	2%	2%	0%	0%	0%	2%	0%	0%	0%	0%
	[361%]	[19%]	[1%]	[1%]	[1%]	[2%]	[1%]	[0%]	[0%]	[0%]	[1%]	[1%]	[1%]	[0%]	[0%]
<i>Panel C: Concentration Parameter</i>															
G^S m obs.	1.2	1.8	3.7	6.6	15.6	15.7	26.6	60.9	111.6	272.0	17	28	51	74	111
G^S (y,x)-census, n obs.	7.0	8.2	10.6	13.1	15.6	118.2	139.6	182.1	227.3	271.2	131	136	128	119	112
G^S (y,x)-sample, m obs.	1.1	1.2	1.9	4.2	15.5	15.2	16.5	28.7	69.6	271.0	1	3	14	36	112
G^S Analytic Correction; (y,x)-census, m obs.	1.7	2.9	6.3	10.0	15.5	26.1	47.2	106.4	172.2	271.6	25	39	70	94	112

Estimators: The model is $y = \alpha + \rho_0 w(G)y + \gamma_0 x + \delta_0 w(G)x + \varepsilon$. We consider four estimators. Each estimator uses a sampled graph, a sample of (y,x), and a data sample for the 2-stage least squares (2SLS) estimation. The sampled graph is either G^S or G^S , the sample of (y,x) is either for the sampled nodes or for an entire census, and the data sample for 2SLS is either data for the m sampled nodes or all available data (any unknown or missing data is replaced with zeros) for the census of n nodes. The "Analytic Correction" uses G^S , the census of (y,x), and 2SLS estimation on only the data for the m sampled nodes. The concentration parameter c is given as follows following Kleibergen (2007). For a first stage $X = Z\pi + v$, we use $c = \pi'S^{-1}\pi$ where $S = \text{var}(p)$ and $p = (ZZ')^{-1}ZX$.

Notes: Estimation is performed using homophily graphs with 250 nodes, 6 cliques of equal size, and expected degree of ~15. The simulation model parameters are $\rho_0=0.5$ and $\delta_0=0.5$. All bias percentages are computed relative to these parameter values. Standard errors of the simulation means are shown in brackets. For comparison, OLS biases at 100% sampling are [ρ_0 ; δ_0] = [46%, -98%], [8%; -16%], and [28%; -63%], reading across the three columns.

TABLE 4. *t*-statistic Diagnosis (Simulated Network Data)

	<i># of Stats where $t_{GR} / t_{Naive} > 1$ out of tota</i>				<i>Mean of $t_{Reconstructed} / t_{Naive}$</i>			
	1/4	1/3	1/2	2/3	1/4	1/3	1/2	2/3
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
<i>Panel A: Graph Level Regressions</i>								
G^S	6/7	7/7	7/7	7/7	5.32	2.15	1.61	1.29
G^S	5/7	7/7	6/7	3/7	1.61	1.41	1.26	1.08
<i>Panel B: Node Level Regressions</i>								
G^S	5/6	5/6	6/6	6/6	4.09	2.74	2.21	1.83
G^S	6/6	6/6	6/6	6/6	3.23	2.94	2.46	1.93

TABLE 5. Bias in Estimation of β_0 in Jackson and Rogers (2007a) Model

	<i>Raw Network Data</i>				<i>Graphical Reconstruction</i>			
	1/4	1/3	1/2	2/3	1/4	1/3	1/2	2/3
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
<i>Panel A: Simulated Networks, Bias % in Estimation of β_0</i>								
G^S	329.0%	250.0%	104.0%	55.0%	1.0%	0.0%	0.0%	0.0%
G^S	117.0%	85.0%	28.0%	12.0%	0.0%	0.0%	0.0%	0.0%
<i>Panel B: Indian Networks, Bias % in Estimation of β_0</i>								
G^S	263.0%	230.0%	103.0%	53.0%	8.0%	8.0%	7.0%	6.0%
G^S	129.0%	92.0%	31.0%	14.0%	6.0%	5.0%	3.0%	1.0%

Notes: Table presents bias in estimation of β_0 the transmission parameter in the Jackson-Rogers diffusion model described in section

TABLE 6. Summary of Results from Numerical Simulations with Indian Village Networks

	Network Level Regressions				Node Level Regressions			
	1/4 [1]	1/3 [2]	1/2 [3]	2/3 [4]	1/4 [5]	1/3 [6]	1/2 [7]	2/3 [8]
<i>Panel A: G^S</i>								
Raw Network Statistic Median Bias	67.52%	48.62%	23.48%	9.88%	88.02%	70.10%	48.68%	25.33%
Reconstruction Median Bias	0.90%	0.56%	0.35%	0.10%	6.97%	4.62%	1.84%	0.73%
Reconstruction Reduces Bias	8/8	8/8	8/8	8/8	5/5	5/5	5/5	5/5
Raw Network Statistic Bias	67.52%	47.76%	23.48%	11.56%				
Analytic Correction Median Bias	14.58%	6.94%	2.33%	0.65%				
Analytic Corrections Reduce Bias	3/3	3/3	3/3	3/3				
<i>Panel B: G^{LS}</i>								
Raw Network Statistic Median Bias	88.27%	73.20%	51.36%	32.78%	90.94%	81.50%	60.25%	38.46%
Reconstruction Median Bias	13.50%	5.79%	2.28%	0.66%	28.43%	16.91%	7.98%	3.97%
Reconstruction Reduces Bias	7/8	8/8	8/8	8/8	5/5	5/5	5/5	5/5
Raw Network Statistic Bias	190.26%	132.95%	75.19%	41.41%				
Analytic Correction Median Bias	18.57%	13.62%	8.69%	5.41%				
Analytic Corrections Reduce Bias	2/2	2/2	2/2	2/2				

Notes: We ran all the regressions from Tables 1 and 2. We present the median bias across several network statistics at each sampling rate, where the network statistics are those analyzed in Tables 1 and 2. Networks are drawn from the data-set of Banerjee et al. (2011) where the graphs are treated as full networks and then sampled at rates 1/4, 1/3, 1/2, and 2/3. The network formation model used is a conditional edge-independence logistic regression model where covariates are distance in GPS coordinates, number of beds, number of rooms, number of latrines, and electricity provision. The full table is available upon request.

TABLE 7. Empirical Application: Diffusion of Microfinance

<i>Panel A: Graph Level Regressions of Microfinance Take-up Rate of Village on Network Characteristics</i>				
	Eigenvector Centrality of Injection Point [1]	Average Path Length [2]	Conductance [3]	Var. of Eig. Centrality Distribution [4]
G ^s	1.6335 [0.8715]	-0.0545 [0.0488]	0.9659 [0.2671]	58.6554 [40.5225]
Graphical Reconstruction	2.4244 [1.0680]	-0.0928 [0.0597]	1.4443 [0.3644]	87.7406 [49.4477]

<i>Panel B: Vertex Level Regressions of Microfinance Take-up Decision of a Household on Network Characteristics</i>		
	Eigenvector Centrality of Household [1]	Inverse Distance to Injection Point [2]
G ^s	0.5528 [0.1887]	0.0221 [0.0277]
Graphical Reconstruction	0.662 [0.2091]	0.0453 [0.0286]

<i>Panel C: Regression of Household's Take-up Decision on Neighbors' Take-up Decisions</i>			
	ρ [1]	γ [2]	δ [3]
G ^s , (y,x)-Census; n obs.	-0.027274 [0.0221]	0.058066 [0.0138]	0.016371 [0.0115]
G ^s , (y,x)-Sample; m obs.	0.047085 [0.0098]	0.079207 [0.0175]	-0.00090318 [0.0047]
Analytic Correction, (y,x)-Census; m obs.	-0.071297 [0.049]	0.025769 [0.0206]	0.025271 [0.0163]

Notes for Panel A: This panel provides estimates of regressions of the microfinance take-up rate of a village on its network characteristics. Each column presents a separate regression and all regressions control for number of households, fraction of households with savings access, and fraction of households that are SC/ST (scheduled caste/ scheduled tribe). The sample consists of 43 villages. The mean of the dependent variable is 0.1849. All regressions control for the number of nodes in the graph. Standard errors are heteroskedasticity robust.

Notes for Panel B: This panel provides estimates of regressions of whether a household decides to join microfinance on its network characteristics. The sample consists of 8441 households. The mean of the dependent variable is 0.1750. All regressions include village fixed effects. Standard errors are clustered by village. Inverse distance to injection point is the inverse of the minimal path length from the household in question to the set of initially informed households.

Notes for Panel C: This panel provides 2SLS estimates of regressions of whether a household joins microfinance on whether its neighbors join microfinance. The sample of (y,x) is either for just the sampled nodes or the entire census and the data sample for 2SLS is either data for the m sampled nodes or all available data for the census of n nodes. The analytic correction uses the induced subgraph, the census of (y,x), and 2SLS estimation on only the data for the m sampled nodes. Standard errors are clustered by village.

TABLE 8. Bias for Network Level Regression for Average Degree (Simulated Graphs)

Estimator for Regression on Avg. Degree	Property	Edge Sampling Rate				
		(1/4) ²	(1/3) ²	(1/2) ²	(2/3) ²	1
β Naive Estimator	% Bias	1361.2%	756.8%	291.4%	123.0%	0.0%
	Coverage	0%	0%	0%	0%	96%
β Analytic Correction	% Bias	-8.7%	-4.8%	-2.2%	-0.9%	0.0%
	Coverage	39%	61%	81%	87%	95%
β Regularized Analytic Correction (Using Analytic Variance of Correction)	% Bias	0.3%	0.3%	-0.2%	-0.1%	0.0%
	Coverage	97%	96%	96%	96%	95%
β Reg. Analytic Correction, Bootstrap BC (Using Analytic Variance of Correction)	% Bias	-0.4%	-0.1%	-0.4%	-0.1%	0.0%
	Coverage	97%	96%	96%	96%	95%

Estimators: The naive estimator uses the subgraph generated by independently sampling edges at the given sampling rates. The analytic correction estimator, as noted, has residual bias. The regularized analytic correction adjusts for the dispersion term using an analytic formula for the variance of the remaining measurement error left after performing the analytic correction, which is only correct in expectation.

Notes: The sampling rates have been chosen to be comparable with the edge count in the G^S sampling simulations. Coverage computed using bootstrapped standard errors for all estimators. R^2 was set to 0.7. There were 200 simulations done using the standard setup: homophilic networks with 6 groups, graphs with 250 nodes, and 50 graphs per regression.

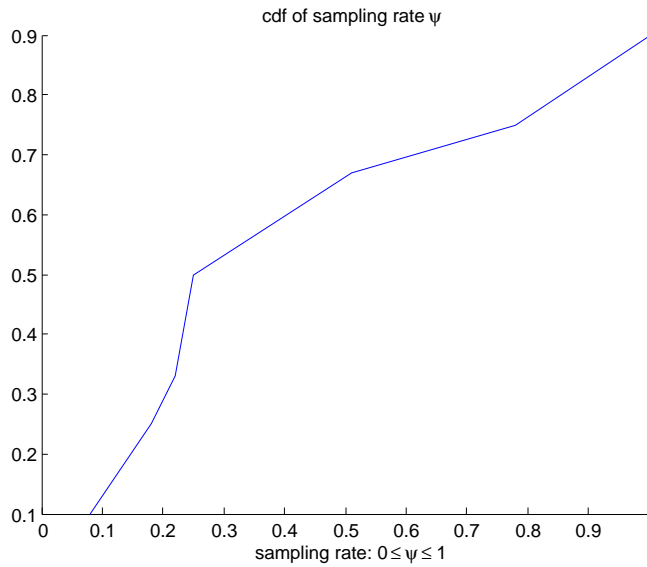
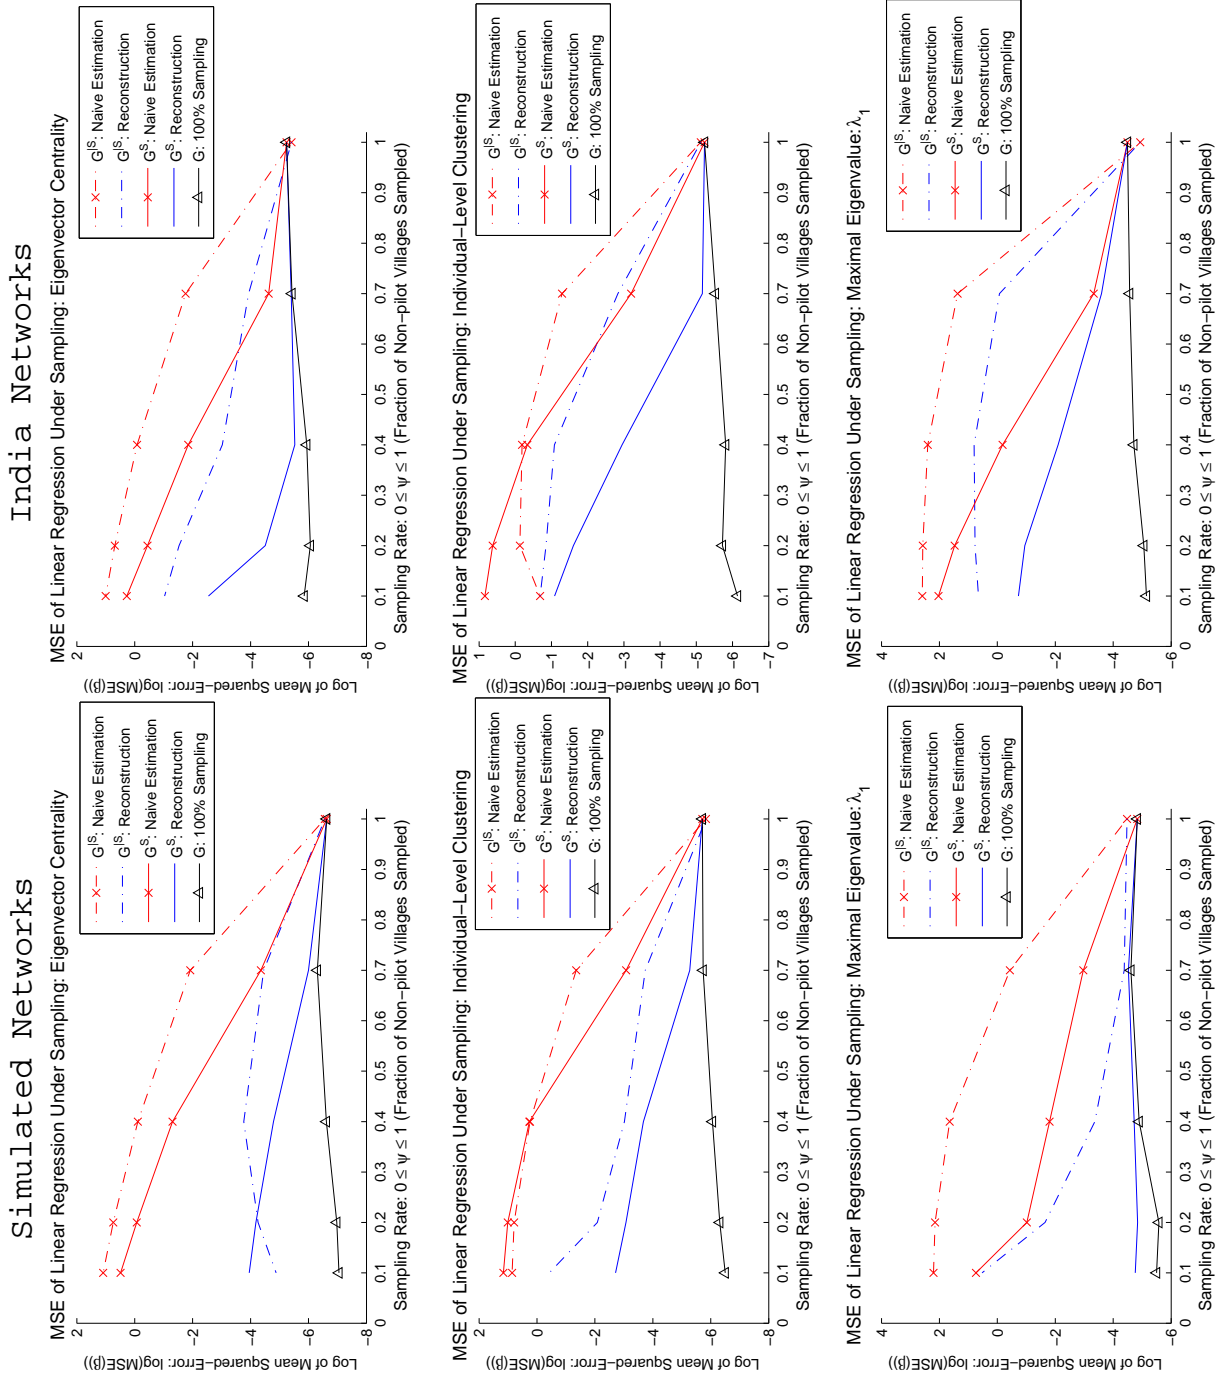


FIGURE 3. A CDF of the sampling rate in the literature. A detailed description of the literature survey is available from the author upon request.

FIGURE 4. Budget Trade-Offs



* The sampling rates used for these plots were 1, 0.7, 0.4, 0.2, and 0.1 which correspond to sampling a total of 33, 40, 50, 60, and 70 villages respectively (including 9 pilot villages).

APPENDIX A. PROOFS FOR SECTION 3

A.1. **Proofs for Section 3.1.1.** Let (a_R) be a sequence of normalizing constants.

ASSUMPTION A.1. *The mismeasured regressor is $X_{rR}^* = \pi_R X_{rR} + v_{rR}$, $\lim_{R \rightarrow \infty} \pi_R = \pi$, $\text{plim}_{R \rightarrow \infty} a_R^{-2} \mathbb{E}_R [v_r X_r] = 0$.*

ASSUMPTION A.2.

- (1) *The data consists of $(y_{rR}, X_{rR})_{r \leq R}$ with $\mathbb{E}[\epsilon_r | X_R] = 0$ and the model is $y_{rR} = \alpha + X_{rR} \beta_0 + \epsilon_{rR}$.*
- (2) *$0 < \sigma_x^2 < \infty$ where $\sigma_x^2 := \text{plim}_{R \rightarrow \infty} a_R^{-2} \mathbb{E}_R (X_r - \bar{X}_{Rr})^2$.*
- (3) *$\sigma_v^2 := \text{plim}_{R \rightarrow \infty} a_R^{-2} \mathbb{E}_R v_r^2 < \infty$.*

LEMMA A.1. *Under Assumption A.1 and A.2, $\hat{\beta} \xrightarrow{P} \pi^{-1} \beta_0 \frac{\sigma_x^2}{\sigma_x^2 + \pi^{-2} \sigma_v^2}$.*

Proof. The proof is classical and follows from $\text{plim} \hat{\beta} = \text{plim} \left(a_R^{-2} X_R^{*'} X_R^* \right)^{-1} a_R^{-2} (\pi_R X_R + v_R)' X_R \beta_0 = \frac{\beta_0}{\pi} \frac{\sigma_x^2}{\sigma_x^2 + \pi^{-2} \sigma_v^2}$. \square

Proof of Proposition 3.1. From Lemma A.2 and A.3, Assumption A.1 holds,

$$\mathbb{E}[d(G_{rR}^S) | G_{rR}] = \underbrace{(\psi + \Theta(n^{-1}))}_{\pi_R} \underbrace{d(G_r)}_{X_{rR}} \quad \text{and} \quad \mathbb{E}[d(G_r^S) | G_r] = \underbrace{(1 - (1 - \psi)^2 + \Theta(n^{-1}))}_{\pi_R} \underbrace{d(G_r)}_{X_{rR}}.$$

The assumptions of the proposition ensure Assumption A.2 holds. The result follows from Lemma A.1. \square

In what follows, we let $t(G) := \sum_{i < j < k} A_{ij} A_{jk}$ denote the number of two-stars ($A_{ij} A_{jk} = 1$), and $p(G) := \sum_{i < j < k < l} A_{ij} A_{kl}$ denote the number of disjoint pairs in the graph.

LEMMA A.2. *Under uniform random sampling of m out of n nodes, with $m/n \rightarrow \psi$ and as $n \rightarrow \infty$,*

$$\mathbb{E}[d(G^{lS}) | G] = (\psi + \Theta(n^{-1})) d(G) \quad \text{and} \quad \text{var}(d(G^{lS}) | G) = \Theta(n^{-1} d(G)) + \Theta(n^{-1} t(G)) + \Theta(n^{-2} p(G)).$$

Proof. Step 1: Let \mathcal{S} be the set of all combinations of m vertices with σ a generic element. Then

$$\mathbb{E}[d(G^{lS}) | G] = \sum_{\sigma \in \mathcal{S}} d(\sigma) P(\sigma) = \binom{n}{m}^{-1} \sum_{\sigma \in \mathcal{S}} m^{-1} \sum_{i \wedge j \in \sigma} A_{ij}.$$

Each pair ij appears $|\{\sigma \in \mathcal{S} : i \wedge j \in \sigma\}| = \binom{n-2}{m-2}$ times, it follows that $\mathbb{E}[d(G^{lS}) | G] = \frac{m-1}{n-1} \cdot \frac{1}{n} \sum_{i,j} A_{ij} = (\psi + \Theta(n^{-1})) d(G)$.

Step 2: Use $|\{\sigma \in \mathcal{S} : i_1 \wedge \dots \wedge i_k \in \sigma\}| = \binom{n-k}{m-k}$ and let $\epsilon_1 := \sum A_{ij}$, $\epsilon_2 := \sum \{A_{ij} A_{kl} : i \vee j \in \{k, l\}, \neg i \wedge j \in \{k, l\}\}$, and $\epsilon_3 := \sum \{A_{ij} A_{kl} : i \notin \{j, k, l\}, j \notin \{k, l\}, k \neq l\}$. Then $\epsilon_1 = 2|E|$, $\epsilon_2 = 8t(G)$, $\epsilon_3 = 8p(G)$, and

$$\mathbb{E}[d(G^{lS})^2 | G] = \frac{1}{m^2} \frac{m!}{(n)_m} \sum_{j=1}^3 \binom{n - (j+1)}{m - (j+1)} \epsilon_j.$$

Some algebra yields $\mathbb{E}[d(G^{lS})^2 | G] = \frac{1}{m} \frac{(m-1)}{(n-1)} \frac{2|E|}{n} + \frac{1}{m} \frac{(m-1)_2}{(n-1)_2} \frac{8t(G)}{n} + \frac{1}{m} \frac{(m-1)_3}{(n-1)_3} \frac{8p(G)}{n}$. Meanwhile, we can expand the square of the mean and obtain coefficients with the exact same summand terms $(\mathbb{E}[d(G^{lS}) | G])^2 = \frac{1}{n^2} \left(\frac{m-1}{n-1} \right)^2 \left(\sum_{i,j} A_{ij} + \sum A_{ij} A_{kl} \right)$.

We have three sets of coefficients. Working with the coefficients on $\sum_{i,j} A_{ij}$, we have that

$$\frac{1}{m} \frac{(m-1)}{(n-1)} \frac{2|E|}{n} - \frac{1}{n} \left(\frac{m-1}{n-1} \right)^2 \frac{2|E|}{n} = \frac{1}{n-1} \left(\frac{m-1}{m} - \frac{(m-1)^2}{n(n-1)} \right) d(G) = (n^{-1} (1 - \psi^2) + \Theta(n^{-2})) d(G).$$

The second term is given by

$$\left(\frac{1}{m} \frac{(m-1)_2}{(n-1)_2} - \frac{1}{n} \left(\frac{m-1}{n-1} \right)^2 \right) \frac{8t(G)}{n} = \frac{1}{n-1} \left(\frac{1}{m} \frac{(m-1)(m-2)}{(n-2)} - \frac{(m-1)^2}{n(n-1)} \right) 8t(G) = (n^{-1} \psi (1 - \psi) + \Theta(n^{-2})) 8t(G).$$

Finally, we compute the last term $\left(\frac{1}{m} \frac{(m-1)_3}{(n-1)_3} - \frac{1}{n} \left(\frac{m-1}{n-1} \right)^2 \right) \frac{8p(G)}{n} = \Theta(n^{-2}) p(G)$, which completes the proof. \square

LEMMA A.3. Under uniform random sampling of m out of n nodes, with $m/n \rightarrow \psi$ and as $n \rightarrow \infty$,

$$\mathbb{E}[d(G^S)|G] = \psi(2 - \psi + \Theta(n^{-1}))d(G) \text{ and } \text{var}(d(G^S)|G) = \Theta(n^{-1}d(G)) + \Theta(n^{-1}t(G)) + \Theta(n^{-2}p(G)).$$

Proof. Same as previous lemma, noting each pair ij appears $|\{\sigma \in \mathcal{S} : i \vee j \in \sigma\}| = 2\binom{n-1}{m-1} - \binom{n-2}{m-2}$ times in \mathcal{S} . \square

Proof of Proposition 3.2. Step 1: For G^S , the expected number of triangles is given by $\mathbb{E}[\rho(G^S)|G] = (3\psi^2(1 - \psi) + \psi^3 + o(1))\rho(G)$ since with probability $3\psi^2(1 - \psi)$ a transitive triangle can be counted with two nodes being sampled and ψ^3 with all three being sampled.⁵⁷ Meanwhile, $\mathbb{E}[\tau(G^S)|G] = (\psi(1 - \psi)^2 + 3\psi^2(1 - \psi) + \psi^3 + o(1))\tau(G)$ where the extra term comes from the fact that the exact center of the triple in question can be selected. In turn,

$$\frac{\rho(G^S)}{\tau(G^S)} = \frac{3\psi^2(1 - \psi) + \psi^3 + o(1)}{\psi(1 - \psi)^2 + 3\psi^2(1 - \psi) + \psi^3 + o(1)} \cdot \frac{\rho(G)}{\tau(G)} = \frac{\psi(3 - 2\psi)}{1 + \psi(1 - \psi)} \cdot \frac{\rho(G)}{\tau(G)} + o(1).$$

For $G^{|S}$, since the sampling probability is $\psi^3 + o(1)$ for both the numerator and the denominator, there is no bias asymptotically in the estimate of clustering.

Step 3: From steps 1 and 2, Assumption A.1 holds. The assumptions of the proposition ensure Assumption A.2 holds. The result follows from Lemma A.1. To check A.2, let ϕ_ρ and ϕ_τ be the limiting scalings above. Notice

$$\left| \frac{\phi_\rho \rho_r}{\phi_\tau \tau_r} - \frac{\mathbb{E}[\rho(\bar{G}_r)|G_r]}{\mathbb{E}[\tau(\bar{G}_r)|G_r]} \right| \leq \left| \frac{\phi_\rho \rho_r \phi_\tau \tau_r - \mathbb{E}[\rho(\bar{G}_r)|G_r] \phi_\tau \tau_r}{\phi_\tau \tau_r \mathbb{E}[\tau(\bar{G}_r)|G_r]} \right| + \left| \frac{\phi_\rho \rho_r \mathbb{E}[\tau(\bar{G}_r)|G_r] - \phi_\rho \rho_r \phi_\tau \tau_r}{\phi_\tau \tau_r \mathbb{E}[\tau(\bar{G}_r)|G_r]} \right|.$$

Then we want to show $\text{plim} \frac{1}{R} \sum_{r=1}^R \frac{\rho_r}{\tau_r} \left| \frac{\phi_\rho \rho_r - \mathbb{E}[\rho(\bar{G}_r)|G_r]}{\mathbb{E}[\tau(\bar{G}_r)|G_r]} \right| = 0$ and $\text{plim} \frac{1}{R} \sum_{r=1}^R \frac{\phi_\rho \rho_r^2}{\phi_\tau \tau_r^2} \left| \frac{\mathbb{E}[\tau(\bar{G}_r)|G_r] - \phi_\tau \tau_r}{\mathbb{E}[\tau(\bar{G}_r)|G_r]} \right| = 0$. It is easy to check each summand converges to zero, e.g. $\frac{a_R \rho_r}{a_R \tau_r} \frac{a_R |\phi_\rho \rho_r - \mathbb{E}[\rho(\bar{G}_r)|G_r]|}{a_R \mathbb{E}[\tau(\bar{G}_r)|G_r]} = O(1) \cdot \frac{o(1)}{O(1)}$ since $a_R \mathbb{E}[\tau(\bar{G}_r)|G_r] \geq o(1) + \underline{\alpha} > 0$, and similarly for the second term. Notice that each summand is dominated by $\frac{\bar{\alpha}}{\underline{\alpha}}$ and a uniform constant. The result follows. \square

A.2. Proofs for Section 3.1.2.

Proof of Lemma 3.1. We show $|d(G^S) - (1 - (1 - \psi)^2)d(G)| = o_P(1)$. The arguments for $G^{|S}$ and $d_2(\cdot)$ are analogous. We have already shown that $\mathbb{E}[d(G^S)|G] = (1 - (1 - \psi)^2 + \Theta(n^{-1}))d(G)$ in Lemma A.3. Let χ_{ij} be an indicator of $i \vee j \in S$. Condition on the sequence of events $\mathcal{E}_n := \{d(G) \in (c_1 a_{1n} \pm \epsilon_1 a_{1n}) \cap d_2(G) \in (c_2 a_{2n} \pm \epsilon_2 a_{2n})\}$. By assumption of growth rates a_{1n} and a_{2n} , we these events happen with probability approaching one,

$$\mathbb{P}\left(\left|\frac{d(G)}{a_{1n}} - c_1\right| < \epsilon_1\right) \rightarrow 1 \text{ and } \mathbb{P}\left(\left|\frac{d_2(G)}{a_{2n}} - c_2\right| < \epsilon_2\right) \rightarrow 1.$$

Notice that the probability distribution and $d(G)$ are both implicitly indexed by n . Then using Chebyshev's inequality and letting $\varphi = (1 - (1 - \psi)^2) + \Theta(n^{-1})$,

$$\mathbb{P}\left(\left|\frac{1}{n} \sum_i \sum_j A_{ij} \chi_{ij} - \varphi d(G)\right| > \epsilon \mid \mathcal{E}_n\right) = \mathbb{P}\left(\left|\frac{1}{n} \sum_i \sum_{j>i} A_{ij} (\chi_{ij} - \varphi)\right| > \epsilon/2 \mid \mathcal{E}_n\right) \leq \frac{4}{n^2 \epsilon^2} \text{var}\left(\sum_i \sum_{j>i} Z_{ij} \mid \mathcal{E}_n\right)$$

where $Z_{ij} = A_{ij}(\chi_{ij} - \varphi)$. If we condition on a graph $G \in \mathcal{E}_n$,

$$\text{var}\left(\sum_i \sum_{j>i} Z_{ij} \mid G\right) = |E(G)| \text{var}(Z_{ij} \mid G) + \sum_i \sum_{j>i} \sum_{k \neq i, j} \text{cov}(Z_{ij}, Z_{jk} \mid G),$$

since $\text{var}(Z_{ij} \mid G) = \varphi(1 - \varphi)$ and $\text{cov}(\chi_{ij}, \chi_{i'j'}) \mid G = \psi^3(1 - \psi)$, with the covariance terms only entering when ij and $i'j'$ share a vertex. Recall that $d(G) = 2|E(G)|/n$, we have

$$\mathbb{P}\left(\left|\frac{1}{n} \sum_i \sum_j A_{ij} \chi_{ij} - \varphi d(G)\right| > \epsilon \mid \mathcal{E}_n, G\right) \lesssim \frac{1}{n} \left\{d(G) \varphi(1 - \varphi) + d_2(G) \psi^3(1 - \psi)\right\}$$

⁵⁷We simply count, e.g., $\left\{3\binom{n-3}{m-2} + \binom{n-3}{m-3}\right\} \binom{n}{m}^{-1} = \frac{m(m-1)(m-2)}{n(n-1)(n-2)} = \psi^3 + o(1)$.

for every $G \in \mathcal{E}_n$ where the constant is uniform. Notice that $d(G) \in (c_1 a_{1n} \pm \epsilon_1 a_{1n})$ implies $d(G)/n \rightarrow 0$ because by assumption $a_{1n}/n \rightarrow 0$. The same is true for $a_{2n}/n \rightarrow 0$. This proves the result.

Next we look at $|d_2(G^S) - k(\psi)d_2(G)| = o_P(1)$.

$$\mathbb{E}[n^{-1} \sum_i \sum_{j \neq i} \sum_{k \neq i, j} A_{ij} A_{jk} \chi_{ij} \chi_{jk} | G] = n^{-1} \sum_i \sum_{j \neq i} \sum_{k \neq i, j} A_{ij} A_{jk} \mathbb{E}[\chi_{ij} \chi_{jk} | g] = k(\psi) d_2(G).$$

This follows from $\mathbb{E}[\chi_{ij} \chi_{jk} | G] = 1 - (1 - \psi)^2 [2 - (1 - \psi)]$. A similar argument to the above completes the proof. \square

We relate the above to the statistical physics literature. To characterize the degree distribution we define a probability generating function $\Psi_0(x) = \sum_{k=0}^{\infty} p_k x^k$ such that $\Psi_0(1) = 1$, where $p_k := P(k)$. Note that the expected degree is given by the derivative of the function evaluated at one, $\langle d \rangle = \sum_{k \in \mathbb{N}} k p_k = \Psi_0'(1)$. Assume nodes are sampled iid with probability ψ and let $\Psi_0^S(x)$ be the generating function describing G^{lS} and $\Psi_0^S(x)$ be that describing G^S .

LEMMA A.4. $\Psi_0^S(x) = \psi \Psi_0(x) + (1 - \psi) \Psi_0(1 - \psi(1 - x))$.

Proof. The proof for $\Psi_0^S(x) = \psi \Psi_0(1 - \psi(1 - x))$ comes from [Stumpf et al. \(2005\)](#). We study $\Psi_0^S(x)$. Consider a random node of degree i . With probability ψ , the node is selected and thus the degree in the subgraph is i . With probability $(1 - \psi)$, the node is not selected and thus the degree is $k \leq i$. The sampled probability is $p_{s,k}$, given by

$$p_{s,k} = \sum_{i \geq k} p_i P(k|i) = \sum_{i \geq k} p_i \left[\psi \mathbf{1}\{i = k\} + (1 - \psi) \binom{i}{k} \psi^k (1 - \psi)^{i-k} \right] = \psi p_k + (1 - \psi) \sum_{i \geq k} p_i \binom{i}{k} \psi^k (1 - \psi)^{i-k} = \psi p_k + (1 - \psi) p_k^*$$

where p_k^* is the sampled probability under G^{lS} (e.g., [Stumpf et al. \(2005\)](#)). Thus, we have $\Psi_0^S(x) = \psi \Psi_0(x) + (1 - \psi) \Psi_0^S(x) = \psi \Psi_0(x) + (1 - \psi) \Psi_0^S(1 - \psi(1 - x))$. \square

Note that the implied expected degrees under this approximation are $\langle d(G^{lS}) \rangle = \psi \langle d \rangle$ and $\langle d(G^S) \rangle = \psi(2 - \psi) \langle d \rangle$, which matches the previous derivations in lemmas [A.2](#) and [A.3](#).

LEMMA A.5. $\langle d_2(G^{lS}) \rangle = \psi^2 \langle d_2 \rangle$ and $\langle d_2(G^S) \rangle = k(\psi) \langle d_2 \rangle$.

Proof. Follows by inspection. \square

ASSUMPTION A.3 (Graph Span Conditions). Let (1) $\mathbb{E}[\epsilon|\ell] = 0$, (2) $0 < \text{plim}_{R \rightarrow \infty} \mathbb{E}_R \log^{-2}(\gamma d_{2r} d_r^{-1}) < \infty$ for $\gamma \in (0, 1]$, (3) the conditions of [3.1](#) hold with $d(G), d_2(G) \in o_P(n^{\gamma_1/2-1} \log^{\gamma_2/2} n)$ for $\gamma_1 < 2$ and γ_2 .

Proof of Proposition 3.3. Recall $\zeta_r = d_2(G_r)/d(G_r)$ and let $\log \bar{\zeta}_r := \log \zeta_r + \log \gamma$. To sign the bias we are interested in

$$\lim_{R \rightarrow \infty} \frac{R^{-1} \sum \text{cov}(\log^{-1}(d_2(\bar{G}_r) d(\bar{G}_r)^{-1}), \log^{-1} \zeta_r)}{R^{-1} \sum \text{var}(\log^{-1}(d_2(\bar{G}_r) d(\bar{G}_r)^{-1}))}.$$

First observe that

$$R^{-1} \sum \text{cov}(\log^{-1}(d_2(\bar{G}_r) d(\bar{G}_r)^{-1}), \log \zeta_r) = R^{-1} \sum \text{cov}(\log^{-1} \bar{\zeta}_r, \log^{-1} \log \zeta_r) + o_P(1).$$

This follows from $|\log^{-1}(d_2(\bar{G}_r) d(\bar{G}_r)^{-1}) - \log^{-1} \bar{\zeta}_r| = o_P(1)$ which we can see by considering the numerator of the fraction and noting

$$|\log \bar{\zeta}_r - \log(d_2(\bar{G}_r) d(\bar{G}_r)^{-1})| \leq |\log d_2(G_r) - \log d(G_r) + \log \gamma - \log d_2(\bar{G}_r) + \log d(\bar{G}_r)| = o_P(1),$$

by Lemma [3.1](#), where $\gamma = k(\psi)/\psi$ or ψ depending on G^S or G^{lS} , using the fact that $\log(\cdot)$ is Lipschitz on $\mathbb{R}_{\geq 1}$.

Therefore we are interested in $\lim_{R \rightarrow \infty} \frac{R^{-1} \sum \text{cov}(\log^{-1} \bar{\zeta}_r, \log^{-1} \zeta_r)}{R^{-1} \sum \text{var}(\log^{-1} \bar{\zeta}_r)}$.

If $\gamma > \zeta_r^{-1}$ for every r , then $\text{cov}(\log^{-1} \bar{\zeta}_r, \log^{-1} \zeta_r)$ is positive for every r by definition. In addition, for every r , since $\log \gamma < 0$,

$$\text{cov}((\log \zeta_r + \log \gamma)^{-1}, \log^{-1} \zeta_r) < \text{var}(\log^{-1} \zeta_r).$$

This shows $\lim_{R \rightarrow \infty} \frac{R^{-1} \sum \text{cov}(\log^{-1} \bar{\zeta}_r, \log^{-1} \zeta_r)}{R^{-1} \sum \text{var}(\log^{-1} \bar{\zeta}_r)} < 1$.

Next, assume that $\gamma < \zeta_r^{-1}$. Then $\text{sign}\{\text{cov}(\log^{-1} \bar{\zeta}_r, \log^{-1} \zeta_r)\}$ depends on the distribution of ζ_r ; it cannot be signed. Therefore $\text{plim}_{R \rightarrow \infty} \sum \text{cov}(\log^{-1} \bar{\zeta}_r, \log^{-1} \zeta_r)$ can take either sign. This is easily seen geometrically.

Finally, that the analytical corrections are consistent follows from the above argument in the first step, noting that we use $\log(\gamma d_2(\bar{G}_r)d(\bar{G}_r)^{-1})$ and therefore

$$\left| \log \zeta_r - \log(\gamma d_2(\bar{G}_r)d(\bar{G}_r)^{-1}) \right| \leq \left| \log \bar{\zeta}_r - \log(d_2(\bar{G}_r)d(\bar{G}_r)^{-1}) \right| + |\log \gamma - \log \gamma| = o_P(1)$$

which completes the result. \square

A.3. Proofs for Section 3.1.2.

Proof of Proposition 3.4. Step 1: It is clear that $\text{Tr}(A(G^S)^k) < \text{Tr}(A^k)$, since we can partition

$$n^{-1} \sum_{i_1, \dots, i_k \in V^k} A_{i_1 i_2} \dots A_{i_k i_1} = n^{-1} \sum_{\mathcal{A}} A_{i_1 i_2} \dots A_{i_k i_1} + \frac{1}{n} \sum_{V^k \setminus \mathcal{A}} A_{i_1 i_2} \dots A_{i_k i_1}$$

with $\mathcal{A} = \{i_1, \dots, i_k : \forall t \in [k], i_t \vee i_{t+\epsilon(t)} \in S, \epsilon(t) \in \{-1, 1\}\}$ and $\sum_{V^k \setminus \mathcal{A}} A(G^S)_{i_1 i_2} \dots A(G^S)_{i_k i_1} = 0$. Let $S_{k, \sigma}$ be the set of all k -sequences using elements from σ . Let η_j be the number of terms in the sum with j distinct nodes. Then notice $\frac{1}{n} \sum \eta_j = \mu^k$. For $A(G^S)$, we have

$$\mathbb{E}[\mu^k(G^S)|G] = m^{-1} \sum_{\sigma \in S} \sum_{i_1, \dots, i_k \in S_{k, \sigma}} A_{i_1 i_2} \dots A_{i_k i_1} P(\sigma) = m^{-1} P(\sigma) \sum_{\sigma} \sum_{j=2}^k \binom{n-j}{m-j} \eta_j = m^{-1} \sum_{j=2}^k \frac{\binom{m}{j}}{\binom{n}{j}} \eta_j = n^{-1} \sum_{j=2}^k \frac{\binom{m-1}{j}}{\binom{n-1}{j}} \eta_j.$$

Step 2: G^S is a compression of both G^S and G , so $\lambda_k(G^S) \leq \lambda_k(G)$, $\lambda_k(G^S)$ by the Cauchy's interlacing theorem. Noticing $[A(G^S)]_{ij} \leq [A(G)]_{ij}$, $\lambda_{\max}(G^S) = \sup_{\alpha \in S^{n-1}} \alpha' A(G^S) \alpha \leq \sup_{\alpha \in S^{n-1}} \alpha \alpha_j [A(G)]_{ij} = \lambda_{\max}(G)$. \square

A.4. Proofs for Section 3.2. We use T to denote the row-stochastized adjacency matrix, $T_{ij} = A_{ij}/d_i$, instead of w , where d_i is the degree of node i to follow the literature (e.g., Jackson, 2008b).

Proof of Proposition 3.5. Step 1: We show the argument for the case with G^S . The argument for G^S is similar, but omitted. Let

$$\bar{u} = M^0 (\rho(T - \bar{T})y + \delta(T - \bar{T})x + \epsilon) = M^0 u,$$

where $M^0 = I_n - \omega'/n$. The instrument is $\bar{Z} = [\iota, x, \bar{T}x, \bar{T}^2 x]$. It suffices to show $\mathbb{E}[\bar{Z}' M^0 u] \neq 0$. We can write

$$\mathbb{E}[\bar{Z}' \bar{u} | x, G] = \mathbb{E} \left[(\omega' M^0 u, x' M^0 u, x' \bar{T}' M^0 u, x' \bar{T}^2 M^0 u)' \middle| x, G \right].$$

The first two components mechanically have expectation zero. By the reduced form representation in section 3.2, we can write y as a function of x and powers of T . The third component requires considering terms of the form $x' \mathbb{E}[\bar{T}' M^0 (\bar{T} - T) | x, G] \delta x$. For generic x , this is zero if and only if $\mathbb{E}[\bar{T}' M^0 (\bar{T} - T) | x, G] = 0$. If we show $\text{Tr}(\mathbb{E}[\bar{T}' M^0 \bar{T} | G]) \neq \text{Tr}(\mathbb{E}[\bar{T}' M^0 T | G])$, the preceding equation does not hold. We pass the expectation by linearity, use a cyclic permutation and write

$$\mathbb{E}[\text{Tr}(\bar{T}' M^0 \bar{T})] = \mathbb{E}[\text{Tr}(M^0 \bar{T} \bar{T}')] = \mathbb{E}[\text{Tr}(\bar{T} \bar{T}')] - n^{-1} \mathbb{E}[\text{Tr}(\bar{T}' \omega \omega' \bar{T})].$$

Let $\langle \cdot, \cdot \rangle_F$ be the Frobenius inner product, $\langle A, B \rangle_F = \text{Tr}(AB')$ and $\|\cdot\|_F$ the Frobenius norm, $\|A\|_F^2 = \text{Tr}(AA')$. In Lemmas A.6, A.7, A.8, and A.9 we compute the following four terms $\mathbb{E}[\|\bar{T}\|_F^2]$, $\mathbb{E}[\langle T, \bar{T} \rangle_F]$, $n^{-1} \mathbb{E}[\text{Tr}(\bar{T}' \omega \omega' \bar{T})]$, and $n^{-1} \mathbb{E}[\text{Tr}(\bar{T}' \omega \omega' T)]$ which we then use to complete the argument. We find

$$\text{Tr}(\mathbb{E}[\bar{T}' M^0 T | G]) = (1-n^{-1})\{\|T\|_F^2 + \sum_i \xi_2(d_i, \psi)\} - n^{-1} \xi_4(\bar{d}, \psi), \quad \text{Tr}(\mathbb{E}[\bar{T}' M^0 \bar{T} | G]) = (1-n^{-1})\{\|T\|_F^2 + \sum_i \xi_1(d_i, \psi)\} - n^{-1} \xi_3(\bar{d}, \psi).$$

In Lemma A.10 we show that $(n-1) \sum_i \{\xi_2(d_i, \psi) - \xi_1(d_i, \psi)\} - \{\xi_4(\bar{d}, \psi) - \xi_3(\bar{d}, \psi)\} \neq 0$ for all but finitely many $\psi \in (0, 1)$, with an upper bound of $2 \cdot \max_i d_i$ points, which completes the argument.

Step 2: We now show that the restriction of the set of observations in the second stage to $i \in S$ yields $\mathbb{E}[Z_{GS} u_{GS}] = 0$. This follows from the fact that $\bar{T}y = Ty$ for all such $i \in S$, and therefore $\bar{u} = \epsilon$. The result follows from the fact that the instrument is correlated with Ty but orthogonal to ϵ , despite measurement error. \square

LEMMA A.6. $\mathbb{E}[\|\bar{T}\|_F^2] = \|T\|_F^2 + \sum_i \xi_1(d_i, \psi)$, where $\xi_1(d_i, \psi) := \frac{1}{1-(1-\psi)^{d_i}} \sum_{r=1}^{d_i} \frac{1}{r} \binom{d_i}{r} \psi^r (1-\psi)^{d_i-r+1} - (1-\psi)/d_i$.

Proof. Observe that $\|\bar{T}\|_F^2 = \sum_{i=1}^n \sum_{k=1}^n \bar{T}_{ik}^2 = \sum_{i:\bar{d}_i > 0} \bar{d}_i^{-1}$ since $\bar{T}_{ik}^2 = \bar{d}_i^{-2}$ is greater than zero exactly \bar{d}_i times. Note that $E[\bar{d}_i^{-1} | i \notin S, \bar{d}_i > 0]$ is the conditional expectation of the first negative moment of a binomial $\text{Bin}(d_i, \psi)$, namely⁵⁸

$$E[\bar{d}_i^{-1} | i \notin S, \bar{d}_i > 0] = \frac{1}{1 - (1 - \psi)^{d_i}} \sum_{r=1}^{d_i} \frac{1}{r} \binom{d_i}{r} \psi^r (1 - \psi)^{d_i - r}.$$

The result follows from the fact that

$$\psi \frac{1}{d_i} + (1 - \psi) \frac{1}{1 - (1 - \psi)^{d_i}} \sum_{r=1}^{d_i} \frac{1}{r} \binom{d_i}{r} \psi^r (1 - \psi)^{d_i - r} = \frac{1}{d_i} + (1 - \psi) \frac{1}{1 - (1 - \psi)^{d_i}} \sum_{r=1}^{d_i} \frac{1}{r} \binom{d_i}{r} \psi^r (1 - \psi)^{d_i - r} - \frac{(1 - \psi)}{d_i} = \frac{1}{d_i} + \xi_{1i}.$$

Summing over the i nodes yields the result. \square

LEMMA A.7. $E[\langle T, \bar{T} \rangle_F] = \|T\|_F^2 + \sum_i \xi_2(d_i, \psi)$, where $\xi_2(d_i, \psi) := -d_i^{-1}(1 - \psi)^{d_i+1}$.

Proof. We can write $\langle T, \bar{T} \rangle_F = \sum_{i=1}^n \sum_{k=1}^n T_{ik} \bar{T}_{ik} = \sum_{i:\bar{d}_i > 0} \bar{d}_i^{-1} d_i^{-1} \bar{d}_i = \sum_{i:\bar{d}_i > 0} d_i^{-1}$. As $P(\bar{d}_i > 0) = 1 - (1 - \psi)^{d_i+1}$, where none of the d_i neighbors nor i is sampled, $E[\langle T, \bar{T} \rangle_F] = \sum_i \frac{1 - (1 - \psi)^{d_i+1}}{d_i} = \|T\|_F^2 + \sum_i \xi_2(d_i, \psi)$. \square

LEMMA A.8. $E[\text{Tr}(\bar{T}' u' \bar{T})] = \|T\|_F^2 + \sum_i \xi_1(d_i, \psi) + \xi_3(\vec{d}, \psi)$, where $\xi_3(\vec{d}, \psi)$ is defined in the proof.

Proof. Notice that

$$\text{Tr}(\bar{T}' u' \bar{T}) = \sum_i |[l' \bar{T}]_i|^2 = \sum_i \left(\sum_k \bar{T}_{ki}^2 + 2 \sum_{l>k} \sum_k \bar{T}_{ki} \bar{T}_{li} \right) = \sum_{k:\bar{d}_k > 0} \bar{d}_k^{-1} + 2 \sum_{l>k} \sum_k \sum_i \bar{T}_{ki} \bar{T}_{li} = \|\bar{T}\|_F^2 + 2 \sum_{l>k:\bar{d}_l > 0} \sum_{k:\bar{d}_k > 0} \frac{|\bar{N}_k \cap \bar{N}_l|}{\bar{d}_k \bar{d}_l}.$$

With probability ψ^2 both $k \in S$ and $l \in S$, so $\bar{d}_k^{-1} \bar{d}_l^{-1} |\bar{N}_k \cap \bar{N}_l| = d_k^{-1} d_l^{-1} |N_k \cap N_l| =: c(k, l)$. With probability $\psi(1 - \psi)$ we have $\zeta_4(k, l)$ and with the same probability we have $\zeta_4(l, k)$, where ζ_4 is defined in Lemma A.9 below. Finally, with probability $(1 - \psi)^2$

$$\zeta_3(k, l) := \sum_{r=1}^{N_l - |N_k \cap N_l|} \sum_{t=1}^{N_k - |N_k \cap N_l|} \sum_{s=1}^{|N_k \cap N_l|} \binom{N_l - |N_k \cap N_l|}{r} \binom{N_k - |N_k \cap N_l|}{t} \binom{|N_k \cap N_l|}{s} \frac{\psi^{s+t+r} (1 - \psi)^{|N_k \cup N_l| - s - t - r}}{(t + s)(r + s)}.$$

Then $E[\text{Tr}(\bar{T}' u' \bar{T})] = \|T\|_F^2 + \sum_i \xi_1(d_i, \psi) + \xi_3(\vec{d}, \psi)$ where

$$\xi_3(\vec{d}, \psi) := 2 \sum_{l>k} \sum_k \left\{ \psi^2 c(k, l) + \psi(1 - \psi) \zeta_4(l, k) + \psi(1 - \psi) \zeta_4(k, l) + (1 - \psi)^2 \zeta_3(k, l) \right\}.$$

\square

LEMMA A.9. $E[\text{Tr}(\bar{T}' u' T)] = \|T\|_F^2 + \sum_i \xi_2(d_i, \psi) + \xi_4(\vec{d}, \psi)$, where $\xi_4(\vec{d}, \psi)$ is defined below.

Proof. We have

$$\text{Tr}(\bar{T}' u' T) = \sum_i [\bar{T}' u]_i [l' T]_i = \sum_i \left(\sum_k \bar{T}_{ki} T_{ki} + 2 \sum_{l>k} \sum_k \bar{T}_{ki} T_{li} \right) = \langle T, \bar{T} \rangle_F + 2 \sum_{l>k} \sum_{k:\bar{d}_k > 0} \sum_i \bar{T}_{ki} T_{li}.$$

The first term has already been controlled in Lemma A.7. To compute the second term, observe that with probability ψ , $k \in S$ and therefore and in this case $\sum_i \bar{T}_{ki} T_{li} = c(k, l)$. With probability $1 - \psi$, $k \notin S$, and as such the conditional expectation is given by

$$\zeta_4(k, l) := \sum_{t=1}^{N_k - |N_k \cap N_l|} \sum_{s=1}^{|N_k \cap N_l|} \binom{|N_k \cap N_l|}{s} \binom{N_k - |N_k \cap N_l|}{t} \frac{\psi^{s+t} (1 - \psi)^{N_k - s - t}}{(s + t) d_l}.$$

Therefore $E[\text{Tr}(\bar{T}' u' T)] = \|T\|_F^2 + \sum_i \xi_2(d_i, \psi) + \xi_4(\vec{d}, \psi)$ where $\xi_4(\vec{d}, \psi) = 2 \sum_{l>k} \sum_k \left\{ \psi c(k, l) + (1 - \psi) \zeta_4(k, l) \right\}$. \square

⁵⁸See, e.g., [Stephan \(1945\)](#).

LEMMA A.10. *Given a graph with non-degenerate coefficients above, for only a finite number of $\psi \in (0, 1)$ can $\text{Tr}(\mathbb{E}[\bar{T}'M^0\bar{T}|G]) \neq \text{Tr}(\mathbb{E}[\bar{T}'M^0T|G])$.*

Proof. Let $f(\psi, G) := (n-1) \sum_i \{\xi_2(d_i, \psi) - \xi_1(d_i, \psi)\} - \{\xi_4(\vec{d}, \psi) - \xi_3(\vec{d}, \psi)\}$. Showing $\text{Tr}(\mathbb{E}[\bar{T}'M^0\bar{T}|G]) \neq \text{Tr}(\mathbb{E}[\bar{T}'M^0T|G])$ is equivalent to showing $f(\psi, G) \neq 0$. By the definitions of ξ_k , $f(\psi, G)$ is a rational function in ψ , with a numerator polynomial of degree bounded by $2 \cdot \max_i d_i$ and coefficients given by G . As we assume that the graph does not yield coefficients so that the rational function is degenerate at zero, by the fundamental theorem of algebra, there are at most $2 \cdot \max_i d_i$ roots of the numerator polynomial in ψ . This bounds the number of sampling rates $\psi \in (0, 1)$ that would exactly satisfy the exclusion restriction for G . \square

A.5. **Proofs for Section 3.3.** As before let $\langle d \rangle := \text{Ed}$, $\langle d^2 \rangle := \text{Ed}^2$, and $\zeta := \langle d^2 \rangle / \langle d \rangle$.

Proof of Proposition 3.6. Step 1: Let β^* solve $\bar{\rho} = \sum_d \frac{\beta^* \sigma(\beta^*) d}{1 + \beta^* \sigma(\beta^*) d} \bar{P}(d)$ where $\bar{P}(d)$ is a sampled degree distribution. By (3.1) and that $\frac{\beta \sigma(\beta) d}{1 + \beta \sigma(\beta) d}$ is strictly increasing in d when $\beta > 0$, we have $\beta^* > \beta_0$ provided first order stochastic dominance of $P(d)$ over $\bar{P}(d)$. For G^S , for every d the count of nodes with at most degree d is weakly increasing under sampling; first order stochastic dominance follows. For G^{lS} the argument is more delicate and relies on this being true in the limit. We use $\limsup_{R \rightarrow \infty} \sup_{r \leq R} \sup_d |P_{Rr}(d) - P_{\infty r}(d)| = 0$ which we did not need for the star subgraph, which implies $\limsup_{R \rightarrow \infty} \sup_{r \leq R} \sup_d |P_{Rr}^{lS}(d) - P_{\infty r}^{lS}(d)| = 0$. Let $F_{|S}(x)$ be the CDF for P_{∞}^{lS} and $F(x)$ for P_{∞} .

$$\begin{aligned} F_{|S}(x) &= \sum_{d \leq x} \sum_{i \geq d} P(i) \binom{i}{d} \psi^d (1-\psi)^{i-d} = \sum_{i=1}^{\infty} P(i) \sum_{d \leq i \wedge x} \binom{i}{d} \psi^d (1-\psi)^{i-d} \\ &= \sum_{i=1}^x P(i) \cdot F_{\text{Bin}(i, \psi)}(i) + \sum_{i=x+1}^{\infty} P(i) \cdot F_{\text{Bin}(i, \psi)}(x) = \sum_{i=1}^x P(i) + \sum_{i=x+1}^{\infty} P(i) \cdot F_{\text{Bin}(i, \psi)}(x) \geq \sum_{d \leq x} P(d) = F(x), \end{aligned}$$

which confirms the stochastic dominance. The usual argument for GMM consistency shows $\text{plim} \hat{\beta} > \beta_0$ since in the limit β^* for every graph is greater than β_0 , proving the result.

Step 2: By Jackson and Rogers (2007b), in graph r infection can spread only if $\beta_0 > \zeta_{rR}^{-1}$. By arguments analogous to those in Lemma 3.1,

$$\begin{aligned} \mathbb{E}[d(G^{lS})] / \mathbb{E}[d^2(G^{lS})] &= \frac{\psi \text{Ed}}{\psi^2 \text{Ed}^2 + (1-\psi)\psi \text{Ed}} + o(1) = \zeta^{-1} + (1-\psi) \cdot \underbrace{\frac{1 + \zeta^{-1}}{\psi \zeta + (1-\psi)}}_{\text{Positive}} + o(1) \text{ and} \\ \mathbb{E}[d(G^S)] / \mathbb{E}[d^2(G^S)] &= \frac{\psi(2-\psi)\text{Ed}}{\psi \text{Ed}^2 + \psi(1-\psi^2)\text{Ed}} + o(1) = \zeta^{-1} + (1-\psi) \underbrace{\left\{ \frac{1 - \zeta^{-1}(1+\psi)}{\zeta + (1-\psi^2)} \right\}}_{\text{Positive if } \zeta > (1+\psi)} + o(1). \end{aligned}$$

The result follows since, by assumption on δ_{rR} , $\beta_0 < \zeta^{-1}(\bar{G}_{rR})$ w.p.a.1 for every r , for \bar{G} either G^{lS} or G^S . \square

APPENDIX B. PROOFS FOR SECTION 4

B.1. **Useful Results.** In what follows, for $p \times q$ matrix D , let $\|\cdot\| := \|\cdot\|_2$ be the matrix norm induced by vector norm $\|\cdot\|$, with $\|D\| := \max_{x \in S^{q-1}} \|Dx\|$.

LEMMA B.1. *Assume for all $\theta_r \in \Theta_r$, $\zeta_{ir}(\theta_r, u)$ are $q \times kq$ matrix (or vector) valued functions satisfying $\sup_{u \in \mathcal{U}} \sup_{i,r} \|\zeta_{ir}(\theta_r, u)\| \leq B_{ir}$ with $\limsup_{R \rightarrow \infty} \mathbb{E}_{nR}[\text{EB}_{ir}] < \infty$ and Assumption 1.3 hold. Then*

$$\sup_{u \in \mathcal{U}} (nR)^{1/2} \left\| \mathbb{E}_{n,R} \left[\zeta_{ir}(\bar{\theta}_r, u) (I_q \otimes (\hat{\theta}_r - \theta_{0r})) \right] \right\| = o_P(1)$$

for $\bar{\theta}_r$ on the line between $\hat{\theta}_r$ and θ_{0r} .⁵⁹

⁵⁹Each component of $\bar{\theta}_r$ may be at a different intermediate point.

Proof of Lemma B.1. This follows from

$$\sup_{u \in \mathcal{U}} (nR)^{1/2} \mathbb{E}_{n,R} \left[\left\| \zeta_{ir}(\bar{\theta}_r, u) \right\| \left\| I_q \otimes (\hat{\theta}_r - \theta_{0r}) \right\| \right] \leq (\mathbb{E}_{n,R}[B_{ir}] \cdot \sqrt{nR} \sup_{r \leq R} \left\| \hat{\theta}_r - \theta_{0r} \right\|) \leq \mathbb{E}_{n,R}[B_{ir}] \cdot O_P \left(a_n^{-1} \cdot \sqrt{nR^{1+2/b}} \right) = o_P(1)$$

as the mean of the envelopes converges and the rates obey the assumed relationship. \square

Proof of Lemma 4.1. For every network $a_n \cdot (\hat{\theta}_r - \theta_{0r}) = - \left(\nabla_{\theta} \hat{V}_{(r)}(\bar{\theta}_r) \right)^{-1} \cdot a_n \cdot \hat{V}(\theta_{0r})$. By the Lipschitz condition 4 of the Lemma and Lemma B.2, $\sup_r \left\| \nabla_{\theta} \hat{V}_{(r)}(\bar{\theta}_r) - \nabla_{\theta} \hat{V}_{(r)}(\theta_{0r}) \right\| = o_P(1)$, which can be seen from

$$\sup_r \left\| \nabla_{\theta} \hat{V}_{(r)}(\bar{\theta}_r) - \nabla_{\theta} \hat{V}_{(r)}(\theta_{0r}) \right\| \leq \sup_r \|B_r\| \cdot \sup_r \left\| \hat{\theta}_r - \theta_{0r} \right\| = o_P(1).$$

By condition 5 of the Lemma $\sup_r \left\| a_n \cdot (\hat{\theta}_r - \theta_{0r}) + \left(\nabla_{\theta} V_{(r)}(\theta_{0r}) \right)^{-1} \cdot a_n \cdot \hat{V}(\theta_{0r}) \right\| = o_P(1)$. By condition 3, we have that $a_n \cdot \sup_r \|\hat{V}_{(r)}(\theta_r)\| = O_P(R^{1/b})$, since $\mathbb{E}[\|\sup_r a_n \cdot \hat{V}_{(r)}(\theta_{0r})\|^{b+1}]^{1/b} \leq R^{1/b} \sup_r \|a_n \cdot \hat{V}_{(r)}(\theta_{0r})\|_{\ell_b}$ and therefore $a_n \cdot \sup_r \left\| \hat{\theta}_r - \theta_{0r} \right\| = O_P(R^{1/b})$. If instead of b moments we assume all moments exist, then $\sqrt{\log(R+1)}$ replaces $R^{1/b}$ by a standard Orlicz inequality (e.g., Van der Vaart and Wellner, 1996). \square

LEMMA B.2. *Conditions 1 and 2 of Lemma 4.1 imply $\sup_r \left\| \hat{\theta}_r - \theta_{0r} \right\| = o_P(1)$.*

Proof of Lemma B.2. Arguing along the lines of Theorem 5 of Supplementary Appendix I of Hahn and Newey (2004), among others, pick $\eta > 0$, define $\epsilon := \inf_{r \leq R} [Q_{(r)}(\theta_{0r}) - \sup_{\theta_r: \|\theta_r - \theta_{0r}\| > \eta} Q_{(r)}(\theta_r)]$, and condition on the event $\left\{ \sup_{r \leq R} \sup_{\theta} \left| \hat{Q}_{(r)}(\theta) - Q_{(r)}(\theta) \right| < \frac{\epsilon}{3} \right\}$ which has probability $1 - o(a_n^{-\nu})$ by Condition 2. If we look for $\hat{\theta}_r$ outside an η -radius ball of the true parameter, we have $\sup_{\theta_r: \|\theta_r - \theta_{0r}\| > \eta} \hat{Q}_{(r)}(\theta_r) < \sup_{\theta_r: \|\theta_r - \theta_{0r}\| > \eta} Q_{(r)}(\theta_r) + \frac{\epsilon}{3} < Q_{(r)}(\theta_{0r}) - \frac{2\epsilon}{3} < \hat{Q}_{(r)}(\theta_{0r}) - \frac{\epsilon}{3}$, contradicting $\hat{Q}_{(r)}(\hat{\theta}_r) \geq \hat{Q}_{(r)}(\theta_{0r})$, implying $\left\| \hat{\theta}_r - \theta_{0r} \right\| < \eta$ for all networks. \square

The next lemma is useful throughout and we make explicit the dependence on R in $P^R(\cdot)$ and E^R here.

LEMMA B.3 (Extended Vitali Convergence). *Let $\{Z_R : R \in \mathbb{N}\}$ be L^1_R -integrable functions on a sequence of measure spaces indexed by R . (1) $Z_R \xrightarrow{P} 0$ and (2) Z_R is uniformly integrable, $\sup_{R \geq 1} E^R [|Z_R| \cdot \mathbf{1}\{|Z_R| > c\}] \rightarrow 0$ as $c \rightarrow \infty$, imply $E^R [Z_R] \rightarrow 0$.*

Proof. The argument is analogous to the proofs of Theorem 10.3.5 and 10.3.6 in Dudley (2002). Let $Z_R \geq 0$ w.l.o.g. First, observe (2) implies that for every $\epsilon > 0$ there exists $\delta > 0$ such that for each A_R with $P^R(A_R) < \delta$, $E^R [Z_R \cdot \mathbf{1}\{A_R\}] < \epsilon$ for all R . To see this, by uniform integrability, given $\epsilon > 0$ and pick $\delta < \epsilon/(2K)$ where K is large enough such that $E^R [Z_R \cdot \mathbf{1}\{Z_R > K\}] < \epsilon/2$. Then

$$E^R [Z_R \cdot \mathbf{1}\{A_R\}] \leq E^R [Z_R \cdot \mathbf{1}\{A_R\} \cdot \mathbf{1}\{Z_R \leq K\}] + E^R [Z_R \cdot \mathbf{1}\{A_R\} \cdot \mathbf{1}\{Z_R > K\}] < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$$

as $P^R(A_R) < \delta$. Second, given $\epsilon > 0$, let $A_R := \{Z_R > \epsilon\}$. By (1), $P^R(A_R) \rightarrow 0$ as $R \rightarrow \infty$. As such, for R large enough $P^R(A_R) < \epsilon/(2K)$. Therefore $E^R [Z_R] = E^R [Z_R \cdot \mathbf{1}\{Z_R \leq \epsilon\}] + E^R [Z_R \cdot \mathbf{1}\{A_R\}] \leq 2\epsilon$. \square

B.2. Proof of Theorem 4.1. The argument is straightforward and follows by expanding around $\hat{\theta}$ uniformly and then checking that conditional expectations preserve asymptotic normality.

Proof of Theorem 4.1.1. Consistency is clear so we directly check normality. Let $u_{ir} = \epsilon_{ir} + (w_{ir}(G_r) - \mathcal{E}_{ir}(x_r; \hat{\theta}_r))\beta_0$ and $H_R(\hat{\theta}) := (nR)^{-1} \mathcal{E}(x; \hat{\theta})' \mathcal{E}(x; \hat{\theta})$. As usual $\sqrt{nR}(\hat{\beta} - \beta_0) = H_R(\hat{\theta})^{-1} \cdot (nR)^{-1/2} \mathcal{E}(x; \hat{\theta})' u$. **Step 1:** To show $H_R(\hat{\theta}) - H_{R,0}(\theta_0) = o_P(1)$, where $H_{R,0}(\theta_0) := \mathbb{E}_{nR} [E[\mathcal{E}_{ir}(x_r; \theta_{0r}) \mathcal{E}_{ir}(x_r; \theta_{0r})']]$. By a term-by-term expansion,

$$H_R(\hat{\theta}) = H_R(\theta_0) + \mathbb{E}_{nR} \left[\dot{\Psi}_{ir}(\bar{\theta}_r) \left(I_p \otimes (\hat{\theta}_r - \theta_{0r}) \right) \right]$$

where $\dot{\Psi}_{ir}(\bar{\theta}_r) = \frac{\partial}{\partial \theta'} \{ \mathcal{E}_{ir}(x_r; \bar{\theta}_r) \mathcal{E}_{ir}(x_r; \bar{\theta}_r)' \}$ is a $p \times pk$ matrix of derivatives w.r.t. θ_{rk} , which exists by Assumption 1. By Lemma B.5 the second term $o_P(1)$ since $\| \dot{\Psi}_{ir}(\bar{\theta}_r) \| \lesssim \left\| \frac{\partial}{\partial \theta'} \mathcal{E}_{ir}(x_r; \bar{\theta}_r) \right\| \| \mathcal{E}_{ir}(x_r; \bar{\theta}_r) \| \lesssim L_{i,r}(x_r)^2$, which by Assumption 2.2 follows from

$$\| \mathcal{E}_{ir}(x_r; \bar{\theta}_r) \| \leq L_{i,r}(x_r) \text{ and } \left\| \frac{\partial}{\partial \theta'} \mathcal{E}_{ir}(x_r; \bar{\theta}_r) \right\| \leq \sqrt{E[\|w_{ir}\|^2 | x_r; \bar{\theta}_r]} \cdot \sqrt{E[\| \mathcal{I}_{w_{ir}|x_r}(\bar{\theta}_r) \| | x_r; \bar{\theta}_r]} \leq L_{i,r}(x_r).$$

That $H_R(\theta_0) - H_{R,0}(\theta_0) = o_P(1)$ follows from Vitali convergence.

Step 2: We show $(nR)^{-1/2} \mathcal{E}(x; \hat{\theta})' u \rightsquigarrow \mathcal{N}(0, V)$. Let $f_{ir}(\theta_r; \beta) := \mathcal{E}_{ir}(x_r; \theta_r)(w_{ir}(G_r) - \mathcal{E}_{ir}(x_r; \theta_r))' \beta$. Then

$$\mathcal{E}(x; \hat{\theta})' u / \sqrt{nR} = f(\theta_0; \beta_0) / \sqrt{nR} + \mathcal{E}(x; \theta_0)' \epsilon / \sqrt{nR} + \sqrt{nR} \mathbb{E}_{nR} \left[\dot{\Phi}_{ir}(\bar{\theta}_r) \left(I_p \otimes (\hat{\theta}_r - \theta_{0r}) \right) \right]$$

where $\dot{\Phi}_{ir}(\bar{\theta}_r) = \frac{\partial}{\partial \theta'} \{ f_{ir}(\bar{\theta}_r; \beta_0) + \mathcal{E}_{ir}(x_r; \bar{\theta}_r) \epsilon_{ir} \}$. Clearly $E[f_{ir}(\theta_{0r}; \beta_0)] = 0$ and $E[\mathcal{E}_{ir}(x_r; \theta_{0r}) \epsilon_{ir}] = 0$.⁶⁰ Let $g_{ir}(\theta_{0r}; \beta_0) := f_{ir}(\theta_{0r}; \beta_0) + \mathcal{E}_{ir}(x_r; \theta_{0r}) \epsilon_{ir}$. That $\sqrt{nR} \mathbb{E}_{nR} g_{ir}(\theta_{0r}, \beta_0) \rightsquigarrow \mathcal{N}(0, V)$, where $V = \lim E_R \text{var}(\sqrt{n} \mathbb{E}_n g_{ir}(\theta_{0r}, \beta_0))$, follows from Lemma B.5 which can be applied by Assumption 2.3.

Finally, the $\dot{\Psi}_{ir}(\bar{\theta}_r)$ term is controlled by Lemma B.1 using $\left\| \frac{\partial}{\partial \theta'} f_{ir}(\bar{\theta}_r; \beta_0) \right\| \lesssim_P L_{i,r}(x_r)^2$, which follows from

$$\left\| \frac{\partial}{\partial \theta'} f_{ir}(\bar{\theta}_r; \beta_0) \right\| \leq \sup_{\beta \in \mathcal{B}} \| \beta \| \cdot \left\| \frac{\partial}{\partial \theta'} \mathcal{E}_{ir}(x_r; \bar{\theta}_r) \right\| \left\| w_{ir}(G_r) - \frac{\partial}{\partial \theta'} \mathcal{E}_{ir}(x_r; \bar{\theta}_r) \right\| \lesssim L_{i,r}(x_r) \|w_{ir}(G_r)\| + L_{i,r}(x_r)^2 \lesssim_P L_{i,r}(x_r)^2$$

as $\|w_{ir}(G_r)\| \lesssim_P E[\|w_{ir}(G_r)\|] \leq E[L_{i,r}(x_r)]$. \square

Next, we turn to GMM. The argument is conceptually similar and depends on uniform expansions. The main difference is control of score functions. Let us define the conditional score⁶¹ with respect to β and the conditional score with respect to θ_r as $S_{m_{ir}|x_r}(\beta; \theta_r) := \frac{\partial}{\partial \beta'} \log f(m_{ir}|x_r; \beta, \theta_r)$ and $S_{m_{ir}|x_r}(\theta_r; \beta) = \frac{\partial}{\partial \theta_r'} \log f(m_{ir}|x_r; \beta, \theta_r)$. The corresponding information matrices are

$$\mathcal{I}_{m_{ir}|x_r}(\beta; \theta_r) = E[S_{m_{ir}|x_r}(\beta; \theta_r) S_{m_{ir}|x_r}(\beta; \theta_r)' | x_r; \beta, \theta_r] \text{ and } \mathcal{I}_{m_{ir}|x_r}(\theta_r; \beta) = E[S_{m_{ir}|x_r}(\theta_r; \beta) S_{m_{ir}|x_r}(\theta_r; \beta)' | x_r; \beta, \theta_r].$$

Notice that $\| \mathcal{I}_{m_{ir}|x_r}(\beta; \theta_r) \| \leq \| \mathcal{I}_{m_{ir}|x_r}(\beta, \theta_r) \|$ and $\| \mathcal{I}_{m_{ir}|x_r}(\theta_r; \beta) \| \leq \| \mathcal{I}_{m_{ir}|x_r}(\beta, \theta_r) \|$, since each is a projection of the larger information matrix, by Cauchy's interlacing theorem. We also use a shorthand $\mathcal{E}_{ir}(\beta, \theta_r)$ for $\mathcal{E}_{ir}(x_r; \beta, \theta_r)$.

LEMMA B.4. *Under Assumptions 1 and 3, $\hat{\beta}_{\text{gmm}} \xrightarrow{P} \beta_0$.*

Proof. In four steps we check conditions 1-4 of Andrews (1994), Theorem A-1. **Step 1:** The first part is clear by Assumption 3.2, since $E[m(y_{ir}, w_{ir}; \beta) | x_r; \theta_r, \beta] = \mathcal{E}_{ir}(\beta, \theta_r)$. That $\mathcal{E}_{ir}(\beta, \theta_r)$ satisfies a uniform law of large numbers, $\sup_{(\beta, \theta) \in \mathcal{B} \times \prod_{r \in \mathcal{N}} \Theta_r} \| \mathbb{E}_{nR} \{ \mathcal{E}_{ir}(\beta, \theta_r) - E \mathcal{E}_{ir}(\beta, \theta_r) \} \| = o_P(1)$, follows from a pointwise convergence, which is clear, and stochastic equicontinuity. Stochastic equicontinuity follows from a Lipschitz condition. Define the following.

$$\epsilon_{i,r}^1(\bar{\beta}, \beta, \theta_r') := E \left[\frac{\partial}{\partial \beta'} m(y_{ir}, w_{ir}(G_r); \bar{\beta}) | z_r, A_r^o, y_r; \theta_r, \beta \right], \quad \epsilon_{i,r}^2(\beta', \bar{\beta}, \theta_r') := E[m(y_{ir}, w_{ir}(G_r); \beta') \cdot S_{m_{ir}|x_r}(\bar{\beta}; \theta_r')' | z_r, A_r^o, y_r; \theta_r, \bar{\beta}],$$

$$\text{and } \epsilon_{i,r}^3(\beta', \bar{\theta}_r, \theta_r) := E[m(y_{ir}, w_{ir}(G_r); \beta') \cdot S_{m_{ir}|x_r}(\bar{\theta}_r; \beta')' | z_r, A_r^o, y_r; \theta_r, \beta'].$$

By a Taylor expansion it follows that $\| E[m(y_{ir}, w_{ir}(G_r); \beta') | z_r, A_r^o, y_r; \theta_r', \beta'] - E[m(y_{ir}, w_{ir}(G_r); \beta) | z_r, A_r^o, y_r; \theta_r, \beta] \|$ is bounded by $\| \epsilon_{i,r}^1(\bar{\beta}, \beta, \theta_r') + \epsilon_{i,r}^2(\beta', \bar{\beta}, \theta_r') \| \| \beta' - \beta \| + \| \epsilon_{i,r}^3(\bar{\beta}, \theta_r') \| \| \theta_r' - \theta_r \|$. Then we have by Assumption 3.3,

$$\begin{aligned} \| \epsilon_{i,r}^1(\bar{\beta}, \beta, \theta_r') \| &\leq E \left[\left\| \frac{\partial}{\partial \beta'} m(y_{ir}, w_{ir}(G_r); \bar{\beta}) \right\| | x_r; \theta_r', \beta \right] \leq L_{i,r}(x_r), \\ \| \epsilon_{i,r}^2(\beta', \bar{\beta}, \theta_r') \| &\leq \sqrt{E[\|m_{ir}(\beta')\|^2 | x_r; \theta_r', \bar{\beta}]} \cdot \sqrt{\| \mathcal{I}_{m_{ir}|x_r}(\bar{\beta}; \theta_r') \|} \leq L_{i,r}(x_r), \text{ and} \\ \| \epsilon_{i,r}^3(\beta', \bar{\theta}_r, \theta_r) \| &\leq \sqrt{E[\|m_{ir}(\beta')\|^2 | x_r; \theta_r', \bar{\beta}]} \cdot \sqrt{\| \mathcal{I}_{m_{ir}|x_r}(\bar{\theta}_r; \beta') \|} \leq L_{i,r}(x_r). \end{aligned}$$

⁶⁰The former by iterated expectations and the latter by assumption on ϵ .

⁶¹The right-hand side of the expression abuses notation. Given that graphs occupy a discrete space, with continuously distributed disturbances in the model we may formally have to decompose the measure into its absolutely continuous part and its pure point part (e.g., Lebesgue decomposition theorem). We rule out the singular part by assumption.

As such, with $d_{\mathcal{B} \times \prod_{r \in \mathbb{N}} \Theta_r}((\beta, \theta), (\beta', \theta')) := \|\beta - \beta'\| \vee \sup_{r \in \mathbb{N}} \|\theta_r - \theta'_r\|$,

$$\left\| \mathbb{E}_{n,R} \left\{ \mathcal{E}_{ir}(\beta, \theta_r) - \mathcal{E}_{ir}(\beta', \theta'_r) \right\} \right\| \leq \mathbb{E}_{n,R} L_{i,r} \cdot \sup_{r \leq R} \left(\|\beta' - \beta\| + \|\theta'_r - \theta_r\| \right) \leq \mathbb{E}_{n,R} L_{i,r} \cdot d_{\mathcal{B} \times \prod_{r \in \mathbb{N}} \Theta_r}((\beta, \theta), (\beta', \theta'))$$

Since $\sup_{R \geq 1} \mathbb{E}_{n,R} [EL_{i,r}] < \infty$, by Andrews (1992) Lemma 2 it follows that stochastic equicontinuity holds and that $\mathbb{E}_{n,R} E \mathcal{E}_{ir}(\beta, \theta_r)$ is uniformly continuous in $(\beta, \theta) \in \mathcal{B} \times \prod_{r \in \mathbb{N}} \Theta_r$.

Step 2: The second and third part of the second condition are clear: by Assumption 3.1 $\widehat{W} - W = o_P(1)$ and by Assumption 1 $(\widehat{\theta}_1, \dots, \widehat{\theta}_R) \in \prod_{r=1}^R \Theta_r$. We need only show $\sup_{\beta \in \mathcal{B}} \left\| \lim_{R \rightarrow \infty} \mathbb{E}_{n,R} E \mathcal{E}_{ir}(\beta, \widehat{\theta}_r) - \lim_{R \rightarrow \infty} \mathbb{E}_{n,R} E \mathcal{E}_{ir}(\beta, \theta_{0r}) \right\| = o_P(1)$ for $\theta_0 \in \prod_{r \in \mathbb{N}} \Theta_r$. This follows from a Taylor expansion and the fact that

$$\sup_{\beta \in \mathcal{B}} \lim_{R \rightarrow \infty} \mathbb{E}_{n,R} \left\| E \frac{\partial}{\partial \theta'} \mathcal{E}_{ir}(\beta, \bar{\theta}_r) \cdot \left(I \otimes (\widehat{\theta}_r - \theta_{0r}) \right) \right\| \leq \limsup_{R \rightarrow \infty} \mathbb{E}_{n,R} E L_{i,r} \cdot \sup_r \left\| \widehat{\theta}_r - \theta_{0r} \right\| = o_P(1)$$

where $\sup_r \left\| \widehat{\theta}_r - \theta_{0r} \right\| = o_P(1)$ by Assumption 1.3 and $\limsup_{R \rightarrow \infty} \mathbb{E}_{n,R} E L_{i,r} < \infty$ by 3.3. This, in turn, results from the fact that $\left\| E \frac{\partial}{\partial \theta'} \mathcal{E}_{ir}(\beta, \bar{\theta}_r) \right\| \leq E L_{i,r}$. To see this, note $\left\| E \frac{\partial}{\partial \theta'} \mathcal{E}_{ir}(\beta, \bar{\theta}_r) \right\| \leq E \left[\left\| E \left[m_{ir}(\beta) S_{m_{ir}|x_r}(\beta; \bar{\theta}_r) | x_r; \beta, \bar{\theta}_r \right] \right\| \right]$.

By Assumption 3.3, $\left\| E \frac{\partial}{\partial \theta'} \mathcal{E}_{ir}(\beta, \bar{\theta}_r) \right\| \leq E \left[\sqrt{E \left[\|M_{i,r}(X_{ir})\|^2 | x_r; \beta, \bar{\theta}_r \right]} \cdot \sqrt{E \left[\|I_{m_{ir}|x_r}(\beta; \bar{\theta}_r)\|^2 \right]} \right] \leq E L_{i,r}(x_r)$.

Step 3: This follows from $E \|\mathcal{E}_{ir}(\beta, \theta_r)\| \leq E L_{i,r}$ and $\limsup_{n \rightarrow \infty} \mathbb{E}_{n,R} E L_{i,r} < \infty$.

Step 4: By iterated expectations $g_n(\beta) = \mathbb{E}_{n,R} E \left[E \left[m(X_{ir}; \beta) | x_r; \theta_{0r}, \beta_0 \right] \right] = \mathbb{E}_{n,R} E \eta_{ir}(\beta, \theta_{0r}, \beta_0)$ where $\eta_{ir}(\beta, \theta_r, \beta) := \mathcal{E}_{ir}(\beta, \theta_r)$. By the identification condition, Assumption 3.1

$$W \lim_{R \rightarrow \infty} \mathbb{E}_{n,R} E \eta_{ir}(\beta, \theta_{0r}, \beta_0) \neq 0 \text{ and } W \lim_{R \rightarrow \infty} \mathbb{E}_{n,R} E \mathcal{E}_{ir}(\beta, \theta_{0r}) \neq 0$$

for $\beta \neq \beta_0$, but $W \lim_{R \rightarrow \infty} \mathbb{E}_{n,R} E \mathcal{E}_{ir}(\beta_0, \theta_{0r}) = W \lim_{R \rightarrow \infty} \mathbb{E}_{n,R} E \eta_{ir}(\beta_0, \theta_{0r}, \beta_0) = 0$. By positive semi-definiteness of W , letting $K'K = W$, observe that $0 \neq Wf(\beta) = K'Kf(\beta)$ implies $Kf(\beta) \neq 0$. It follows that $Q(\beta, \theta_0, W) > Q(\beta_0, \theta_0, W)$ for any $\beta \neq \beta_0$. \square

Proof of Theorem 4.1.2. Step 1: The estimator satisfies for $\widehat{\gamma}_{n,R}(\widehat{\beta}; \widehat{\theta}) := \mathbb{E}_{n,R} \left[\mathcal{E}_{ir}(\widehat{\beta}, \widehat{\theta}_r) \right]$,

$$\left[\frac{\partial}{\partial \beta'} \widehat{\gamma}_{n,R}(\widehat{\beta}; \widehat{\theta}) \right]' \widehat{W} \sqrt{nR} \widehat{\gamma}_{n,R}(\widehat{\beta}; \widehat{\theta}) = o_P(1).$$

A term-by-term expansion yields $\sqrt{nR} \widehat{\gamma}_{n,R}(\widehat{\beta}; \widehat{\theta}) = \sqrt{nR} \widehat{\gamma}_{n,R}(\widehat{\beta}; \theta_0^R) + \sqrt{nR} \mathbb{E}_{n,R} \left[\widehat{\Psi}_{ir}(\bar{\theta}_r; \widehat{\beta}) \left(I_q \otimes (\widehat{\theta}_r - \theta_{0r}) \right) \right]$, where $\widehat{\Psi}_{ir}(\bar{\theta}_r; \widehat{\beta}) = \epsilon_{i,r}^3(\bar{\theta}_r, \bar{\theta}_r, \bar{\theta}_r)$ has been controlled in the preceding lemma. Similarly $\frac{\partial}{\partial \beta'} \widehat{\gamma}_{n,R}(\widehat{\beta}; \widehat{\theta}) = \frac{\partial}{\partial \beta'} \widehat{\gamma}_{n,R}(\widehat{\beta}; \theta_0^R) + \mathbb{E}_{n,R} \left[\widehat{\Phi}_{ir}(\bar{\theta}_r; \widehat{\beta}) \left(I_q \otimes (\widehat{\theta}_r - \theta_{0r}) \right) \right]$,

$$\widehat{\Phi}_{ir}(\bar{\theta}_r; \widehat{\beta}) = E \left[\frac{\partial}{\partial \beta'} m(y_{ir}, w_{ir}; \widehat{\beta}) S_{m_{ir}|x_r}(\bar{\theta}_r; \widehat{\beta}) | A_r^o, y_r, z_r; \bar{\theta}_r, \widehat{\beta} \right] + E \left[m(y_{ir}, w_{ir}; \widehat{\beta}) S_{m_{ir}|x_r}(\widehat{\beta}; \bar{\theta}_r) S_{m_{ir}|x_r}(\bar{\theta}_r; \widehat{\beta}) | A_r^o, y_r, z_r; \bar{\theta}_r, \widehat{\beta} \right].$$

We can bound the first term by $(E[\|\frac{\partial}{\partial \beta'} m(y_{ir}, w_{ir}; \widehat{\beta})\|^2 | x_r; \bar{\theta}_r, \widehat{\beta}]) \cdot E[\|S_{m_{ir}|x_r}(\bar{\theta}_r; \widehat{\beta})\|^2 | x_r; \bar{\theta}_r, \widehat{\beta}])^{1/2}$ and the second by

$$\sqrt{E \left[\left\| \frac{\partial}{\partial \beta'} m(y_{ir}, w_{ir}; \widehat{\beta}) \right\|^2 | x_r; \bar{\theta}_r, \widehat{\beta} \right]} \sqrt{E \left[\|S_{m_{ir}|x_r}(\widehat{\beta}; \bar{\theta}_r)\|^4 | x_r; \bar{\theta}_r, \widehat{\beta} \right]} \cdot \sqrt{E \left[\|S_{m_{ir}|x_r}(\bar{\theta}_r; \widehat{\beta})\|^4 | x_r; \bar{\theta}_r, \widehat{\beta} \right]}$$

and therefore $\left\| \widehat{\Phi}_{ir}(\bar{\theta}_r; \widehat{\beta}) \right\| \leq 2L_{i,r}(x_r)$. By Lemma B.1, $\left[\frac{\partial}{\partial \beta'} \widehat{\gamma}_{n,R}(\widehat{\beta}; \theta_0^R) \right]' \widehat{W} \sqrt{nR} \widehat{\gamma}_{n,R}(\widehat{\beta}; \theta_0^R) = o_P(1)$.

Step 2: A term-by-term expansion for the j th term yields $\sqrt{nR} \widehat{\gamma}_{n,R}^j(\widehat{\beta}; \theta_0^R) = \sqrt{nR} \widehat{\gamma}_{n,R}^j(\beta_0; \theta_0^R) + \sqrt{nR} \frac{\partial}{\partial \beta'} \widehat{\gamma}_{n,R}(\widehat{\beta}; \theta_0^R) (\widehat{\beta} - \beta_0)$, so it remains to be seen that $\left\| \frac{\partial}{\partial \beta'} \widehat{\gamma}_{n,R}(\widehat{\beta}; \theta_0^R) - M \right\| = o_P(1)$, where $M = \lim_{n \rightarrow \infty} \mathbb{E}_{n,R} \left[E \frac{\partial}{\partial \beta'} \mathcal{E}_{ir}(x_r; \beta_0, \theta_{0r}) \right]$, and that $\sqrt{nR} \widehat{\gamma}_{n,R}(\beta_0; \theta_0^R) \rightsquigarrow \mathcal{N}(0, \Omega)$, where $\Omega := \lim_{n \rightarrow \infty} \mathbb{E}_{n,R} \left[\text{var} \left(\sqrt{n} \mathbb{E}_{n,R} \mathcal{E}_{ir}(x_r; \beta_0, \theta_{0r}) \right) \right]$. For the derivative,

$$\begin{aligned} \left\| \frac{\partial}{\partial \beta'} \widehat{\gamma}_{n,R}(\widehat{\beta}; \theta_0^R) - M \right\| &\leq \left\| \frac{\partial}{\partial \beta'} \widehat{\gamma}_{n,R}(\widehat{\beta}; \theta_0^R) - \mathbb{E}_{n,R} E \frac{\partial}{\partial \beta'} \mathcal{E}_{ir}(x_r; \widehat{\beta}, \theta_{0r}) \right\| + \left\| \mathbb{E}_{n,R} E \frac{\partial}{\partial \beta'} \mathcal{E}_{ir}(x_r; \widehat{\beta}, \theta_{0r}) - \lim_{n \rightarrow \infty} \mathbb{E}_{n,R} E \frac{\partial}{\partial \beta'} \mathcal{E}_{ir}(x_r; \widehat{\beta}, \theta_{0r}) \right\| \\ &\quad + \left\| \lim_{n \rightarrow \infty} \mathbb{E}_{n,R} E \frac{\partial}{\partial \beta'} \mathcal{E}_{ir}(x_r; \widehat{\beta}, \theta_{0r}) - M \right\|. \end{aligned}$$

The first term is $o_P(1)$ by a uniform law of large numbers for the derivative which exists by Assumption 3.4 and the second term is $o_P(1)$ by as the limit exists uniformly over \mathcal{B} by Assumption 3.2. The final term is $o_P(1)$ as $\tilde{\beta} - \beta_0 = o_P(1)$ by continuity.

To show $\sqrt{nR}\widehat{\gamma}_{n,R}(\beta_0; \theta_0^R) \rightsquigarrow \mathcal{N}(0, \Omega)$, by Assumption 3.5 we may apply Lemma B.5 with $m(X_{ir}; \beta_0) = Y_{ir}$, $\mathcal{E}_{ir}(x_r; \beta_0, \theta_{0r}) = Z_{ir}$, and $x_r = X_r$ in the notation of the lemma. It is clear that $\Omega := \lim_{n \rightarrow \infty} \mathbb{E}_R \left[\text{var} \left(\sqrt{n} \mathbb{E}_n \mathcal{E}_{ir}(x_r; \beta_0, \theta_{0r}) \right) \right]$, exists under these conditions since for every r , $\lambda_{\max} \left(\text{var} \left(\sqrt{n} \mathbb{E}_n m(X_{ir}; \beta_0) \right) \right) \geq \lambda_{\max} \left(\text{var} \left(\sqrt{n} \mathbb{E}_n \mathcal{E}_{ir}(x_r; \beta_0, \theta_{0r}) \right) \right)$. \square

LEMMA B.5. *Let $Y_{ir,R}$ be a triangular array of mean-zero random variables, $Z_{ir,R} := \mathbb{E}[Y_{ir,R}|X_r]$, and define*

$$\lambda_r := \sqrt{n} \mathbb{E}_n [Y_{ir}] \text{ and } \zeta_r := \sqrt{n} \mathbb{E}_n [Z_{ir}].$$

If (1) $\sup_R \sup_r \text{var}(\lambda_r) < C_1$ and (2) $\inf_R \inf_r \text{var}(\zeta_r) > C_0$, then $\sqrt{nR} \mathbb{E}_{n,R} [Z_{i,r}] \rightsquigarrow \mathcal{N}(0, V_Z)$.

Proof. Identical to the proof of step 1 in ‘‘Proof that Lemma 11 of Chernozhukov et al. (2009) applies’’. \square

Finally, we turn to GMM with an index, which clearly nests GMM with $\mathcal{U} = \{u\}$. The argument is nearly identical to the above, so we only focus on the difference below. The key distinction is that we need to check conditions such that a functional central limit theorem applies. This setup allows the moments to be not differentiable everywhere, per se, in the parameter, though their expectations indeed will be smooth. Therefore, we need to control the complexity of the function class. The usual argument in the iid case would make use of P -Donskerity of function classes, though in our inid case we use Theorem 2.11.1 of Van der Vaart and Wellner (1996), which is analogous. Let us put \sqrt{n} -sums

$$\zeta_{Rr}(\beta, u) = \sqrt{n} \mathbb{E}_n \mathcal{E}_{ir}(\beta, u) / \sqrt{n} \text{ and } \lambda_{Rr}(\beta, u) = \sqrt{n} \mathbb{E}_n m_{ir}(\beta, u) / \sqrt{n}$$

where $\mathcal{E}_{ir}(\beta, u)$ denotes $\mathcal{E}_{ir}(\beta, u, \theta_{0r})$. Our estimator is a root of $\widehat{\Phi}(\cdot, u, \widehat{\theta})$, $\widehat{\Phi}(\widehat{\beta}(u), u, \widehat{\theta}) := \mathbb{E}_{n,R}[\mathcal{E}_{ir}(\widehat{\beta}(u), u, \widehat{\theta}_r)] = 0$. Define $\Psi(\beta, u) := \lim_{R \rightarrow \infty} \mathbb{E}_R \mathbb{E}[\zeta_{Rr}(\beta, u) / \sqrt{n}] = \lim_{R \rightarrow \infty} \mathbb{E}_n \mathbb{E}[\mathcal{E}_{ir}(\beta, u)]$ and $\widehat{\Psi}(\beta, u) = \mathbb{E}_{n,R}[\mathcal{E}_{ir}(\beta, u)]$. By definition $\Psi(\beta_0(u), u) = 0$. Under an identical argument presented above, using Lemma B.1, we can uniformly write

$$\widehat{\Phi}(\widehat{\beta}(u), u, \widehat{\theta}) = \widehat{\Psi}(\widehat{\beta}(u), u) + o_P(1).$$

Therefore, it suffices to consider the behavior of the root of $\widehat{\Psi}(\cdot, u)$, which will converge in probability to the root of $\widehat{\Phi}(\cdot, u)$. Our proof focuses on the differences in the indexed case. The empirical estimator satisfies

$$\widehat{\Psi}(\widehat{\beta}(u), u) = \mathbb{E}_{n,R}[\mathcal{E}_{ir}(\widehat{\beta}(u), u)] = o_P(1).$$

If we can apply Lemma 11 of Chernozhukov et al. (2009), we are done. It will be useful to define a function class

$$\mathcal{F}_R := \left\{ \zeta_{Rr}(\beta, u) = \sqrt{n} \mathbb{E}_n [\mathcal{E}_{ir}(x_r; \beta, u)] : (\beta, u) \in \mathcal{B} \times \mathcal{U} \right\}.$$

Proof that Lemma 11 of Chernozhukov et al. (2009) applies. We need to check the three conditions of the lemma. Conditions 1 and 2 are straightforward to verify and have been mostly assumed in Assumption 4. To see that condition 3, $\sqrt{nR} \left(\widehat{\Psi} - \Psi \right) \rightsquigarrow \mathcal{Z}$ in $\ell^\infty(\mathcal{B} \times \mathcal{U})$, holds we apply Theorem 2.11.1 of Van der Vaart and Wellner (1996) which allows a functional central limit theorem for the inid case with triangular arrays. This proceeds in three steps.

Step 1: To show $\mathbb{E}_R \mathbb{E} \left(\|\zeta_{Rr}(\beta, u)\|_{\mathcal{F}_R}^2 \cdot \mathbf{1} \left\{ \|\zeta_{Rr}(\beta, u)\|_{\mathcal{F}_R} > \eta \sqrt{R} \right\} \right) \rightarrow 0$ for every $\eta > 0$. Define

$$Z_R := \mathbb{E}_R \|\zeta_{Rr}(\beta, u)\|_{\mathcal{F}_R}^2 \cdot \mathbf{1} \left\{ \|\zeta_{Rr}(\beta, u)\|_{\mathcal{F}_R} > \eta \sqrt{R} \right\}$$

and show that $Z_R \xrightarrow{P} 0$. A summand is non-zero only if $\|\zeta_{Rr}(\beta, u)\|_{\mathcal{F}_R} > \eta \sqrt{R}$. By Chebyshev’s inequality

$$\max_{r \leq R} \mathbb{P} \left(\|\zeta_{Rr}(\beta, u)\|_{\mathcal{F}_R}^{2+\delta} > \eta^{2+\delta} R^{1+\delta/2} \right) \leq \frac{\max_{r \leq R} \mathbb{E} \|\zeta_{Rr}(\beta, u)\|_{\mathcal{F}_R}^{2+\delta}}{\eta^{2+\delta} R^{1+\delta/2}}.$$

Observe $\|\zeta_{rR}\|_{\mathcal{F}_R} \leq \sqrt{n} \mathbb{E}_n L_{i,r} =: L_{r,R}$. Under Assumption 4.3, as $\sup_{R \geq 1} \max_{r \leq R} \mathbb{E} (\sqrt{n} \mathbb{E}_n L_{i,r}(x_r))^{2+\delta} = O(1)$,

$$\sup_{r \leq R} \mathbb{P} (\|\zeta_{rR}\|_{\mathcal{F}_R} > \eta \sqrt{R}) \leq \frac{\sup_{r \leq R} \mathbb{E} |L_{r,R}|^{2+\delta}}{\eta^{2+\delta} R^{1+\delta/2}} = O(R^{-1-\delta/2}).$$

All summands are zero with probability at least $(1 - cR^{-1-\delta/2})^R \rightarrow 1$, so $Z_R \xrightarrow{P} 0$. Next, $Z_R \leq \mathbb{E}_R L_{r,R}^2 =: \bar{L}_R^2$. Since the envelope \bar{L}_R^2 is uniformly integrable, by Vitali convergence the result follows.

Step 2: Show $\sup_{\rho((\beta,u),(\beta',u')) < \delta_R} \mathbb{E}_R \mathbb{E} [\zeta_{Rr}(\beta, u) - \zeta_{Rr}(\beta', u')]^2 \rightarrow 0$ for every $\delta_R \downarrow 0$. Notice $\|(\beta, u) - (\beta', u')\| \rightarrow 0$ implies $\zeta_{Rr}(\beta, u) \xrightarrow{a.s.} \zeta_{Rr}(\beta', u')$, from Assumption 4.1. Next, apply the extended Vitali convergence theorem since $\lim(nR)^{-1} \sum_r \mathbb{E}(2 \sum_i L_{i,r}(x_r))^2 < \infty$, meaning that $\mathbb{E}_R \mathbb{E} [\zeta_{Rr}(\beta, u) - \zeta_{Rr}(\beta', u')]^2 \rightarrow 0$ as $\delta_R \rightarrow 0$.

Step 3: To show that $\int_0^{\delta_R} \sqrt{\log N(\epsilon, \mathcal{F}_R, d_n)} d\epsilon \xrightarrow{P} 0$ for every $\delta_R \downarrow 0$. First, \mathcal{F}_R consists of measurable functions by 1 so we have

$$\int_0^{\delta_R} \sqrt{\log N(\epsilon, \mathcal{F}_R, d_n)} d\epsilon \leq \int_0^\infty \sup_{Q \in \mathcal{Q}} \sqrt{\log N(\epsilon \|F_R\|_{Q,2}, \mathcal{F}_R, L_2(Q))} d\epsilon$$

which follows from⁶²

$$\int_0^{\delta_R} \sqrt{\log N(\epsilon, \mathcal{F}_R, d_n)} d\epsilon \leq \int_0^{\delta_R / \|F_R\|_{\mu_R}} \sqrt{\log N(\epsilon \|F_R\|_{Q,2}, \mathcal{F}_R, L_2(Q))} d\epsilon \|F_R\|_{\mu_R} \leq \int_0^\infty \sup_{Q \in \mathcal{Q}} \sqrt{\log N(\epsilon \|F_R\|_{Q,2}, \mathcal{F}_R, L_2(Q))} d\epsilon$$

by the standard argument (Van der Vaart and Wellner, 1996, Proof of Lemma 2.11.6). The strategy is to show that the uniform entropy integral of \mathcal{F}_R can be controlled by the uniform entropy integral of \mathcal{F}_R^* , the unconditional moment functions, added with the uniform entropy integral of $\sup_{\beta', u'} \mathcal{H}_{R|\beta', u'}$, functions which characterize how the probability distribution changes as (β, u) change.

To control \mathcal{F}_R^* , we begin by noting

$$\begin{aligned} \|\zeta_{Rr}(\beta, u) - \zeta_{Rr}(\beta', u')\|_{Q,2} &= \sqrt{n} \left\| \mathbb{E}_n \mathbb{E} [m(X_{ir}; \beta, u) | x_r; \beta, u] - \mathbb{E}_n \mathbb{E} [m(X_{ir}; \beta', u') | x_r; \beta', u'] \right\|_{Q,2} \\ &\leq_{(i)} \sqrt{n} \left\| \mathbb{E} \left[\mathbb{E}_n \left\{ m(X_{ir}; \beta, u) - m(X_{ir}; \beta', u') \right\} | x_r; \beta, u \right] \right\|_{Q,2} \\ &\quad + \sqrt{n} \left\| \mathbb{E}_n \mathbb{E} [m(X_{ir}; \beta', u') | x; \beta, u] - \mathbb{E}_n \mathbb{E} [m(X_{ir}; \beta', u') | x; \beta', u'] \right\|_{Q,2} \end{aligned}$$

by the triangle inequality. Applying Jensen's inequality to the first term,

$$\left\| \mathbb{E} \left[\lambda_r(\beta, u) - \lambda_r(\beta', u') | x_r; \beta, u \right] \right\|_{Q,2} \leq \mathbb{E} \left[\left\| \lambda_r(\beta, u) - \lambda_r(\beta', u') \right\|_{Q,2} | x_r; \beta, u \right] \leq \left\| \lambda_r(\beta, u) - \lambda_r(\beta', u') \right\|_{Q,2}$$

since Q is a measure over X . By Assumption 4.4, $\int_0^\infty \sup_{Q \in \mathcal{Q}} \sqrt{\log N(\epsilon \|F_R^*\|_{Q,2}, \mathcal{F}_R^*, L_2(Q))} d\epsilon < \infty$. Next, we control the second term in the triangle inequality (i) above. These are functions of the form $h(x_r; \beta, u)$, members of $\mathcal{H}_{R|\beta', u'}$ and it has envelope $\bar{H}(x_r) = \sqrt{n} \mathbb{E}_n L_{ir}(x_r)$ by Assumption 4.3. By Assumption 4.4, we have $\int_0^\infty \sup_{Q \in \mathcal{Q}} \sqrt{\log N(\epsilon \|\bar{H}\|_{Q,2}, \mathcal{H}_R, L_2(Q))} d\epsilon < \infty$. Combining the above steps, the result follows as we can bound

$\int_0^\infty \sup_{Q \in \mathcal{Q}} \sqrt{\log N(\epsilon \|F_R\|_{Q,2}, \mathcal{F}_R, L_2(Q))} d\epsilon$ by

$$\int_0^\infty \sup_{Q \in \mathcal{Q}} \sqrt{\log N(\epsilon \|F_R^*\|_{Q,2}, \mathcal{F}_R^*, L_2(Q))} d\epsilon + \sup_{(\beta', u') \in \mathcal{B} \times \mathcal{U}} \int_0^\infty \sup_{Q \in \mathcal{Q}} \sqrt{\log N(\epsilon \|\bar{H}\|_{Q,2}, \mathcal{H}_{R|\beta', u'}, L_2(Q))} d\epsilon < \infty.$$

□

B.3. Example 1: Conditional Edge Independent Models.

Proof of Proposition 4.1. We check the conditions of Lemma 4.1. The first condition is clear by definition and Assumption 5. Condition 2 is shown in Lemma B.7. We show directly $\sup_r \left| |\Xi|^{-1/2} \mathbb{E}_\Xi [v(X_{rs}; \theta_{0r})] \right| = O_P(R^{1/b})$ in Lemma B.6, which is what condition (2) is used for in the proof of Lemma 4.1. The Lipschitz condition follows from a secondary expansion and the assumption of the existence of an envelope function. That is,

$$|\Xi|^{-1} \sum_{s \in \Xi} \{ \nabla_\theta v(X_{rs}; \theta_r^*) - \nabla_\theta v(X_{rs}; \theta_{0r}) \} = |\Xi|^{-1} \sum_{s \in \Xi} F_{rs} (I \otimes (\theta_r^* - \theta_{0r}))$$

⁶²We have $\|F_R\|_{\mu_R} > \delta_R$ if the covering number is $1/\epsilon$, so the integrand is defined as zero for $\epsilon > 1$.

where F_{rs} is a conformable matrix of derivatives evaluated at an intermediate value and $F_{rs} \leq B$, which can be done by Assumption 5.2. The Hessian condition hold by the envelope condition and Chebyshev's inequality. \square

LEMMA B.6. *Under Assumption 5.3, $\sup_r \left| \frac{1}{\sqrt{|\Xi|}} \sum_{s \in \Xi} v(X_{rs}; \theta_{0r}) \right| = O_P(R^{1/b})$.*

Proof. Let $T := |\Xi|$ and define $y_{rT} := \sum v(X_{rs}; \theta_{0r})/\sqrt{T}$. We assume y_{rT} has b^{th} moments, $\sup_{r \leq R} E \|y_{rT}\|^b < \infty$, which will be shown below. That $(E \|\sup_r y_{rT}\|^b)^{1/b} \leq R^{1/b} \sup_r (E \|y_{rT}\|^b)^{1/b}$ implies $\sup_r y_{rT} = O_P(R^{1/b})$. Notice $E [\|y_{rT}\|^b] = E \left[\left\| \sum_s v(X_{rs}; \theta_{0r}) \right\|^b \right] T^{-b/2}$. We need only control $E \left[\left\| \sum_s v(X_{rs}; \theta_{0r}) \right\|^b \right] \leq E \left[\left(\sum_s B(X_{rs}) \right)^b \right]$. Observe $E \left[\left(\sum_{s \in \Xi} B(X_{rs}) \right)^b \right] = E \left[\sum B(X_{rs})^b \right] + E \left[\sum \prod_j B(X_{rs_j})^{\gamma_j} \right]$ where $\sum \gamma_j = b$. If 2^{b-1} moments exist for the envelopes, then this can be majorized into terms of $\prod E[B(X_{rs_j})^{\delta_j}]$ where $\delta_j \leq 2^{b-1}$ by repeated application of Holder inequalities. \square

LEMMA B.7. *Under Assumptions 5 and 6, $P \left(\sup_r \sup_{\theta \in \Theta} \left| \widehat{Q}_{(r)}(\theta) - Q_{(r)}(\theta) \right| \geq \eta \right) = o(|\Xi|^{-1})$.*

Proof. The argument is along the lines of Lemma 2 of Hall and Horowitz (1996) and Lemma 3 of Supplementary Appendix I of Hahn and Newey (2004). We use Lemma B.8. First we use a union bound over the R graphs and focus on $\sum_{r=1}^R P \left(\sup_{\theta \in \Theta} \left| \widehat{Q}_{(r)}(\theta) - Q_{(r)}(\theta) \right| \geq \eta \right)$. Next, considering a given graph, we choose $\epsilon > 0$ such that $2\epsilon \cdot \sup_r E[B(X_{rs})] < \frac{\eta}{3}$. Divide Θ into subsets $\Theta_1, \dots, \Theta_{M(\epsilon)}$ such that $\|\theta - \theta'\| \leq \epsilon$ when θ and θ' are in the same subset. A second union bound gives us $\sum_{j=1}^{M(\epsilon)} P \left(\sup_{\theta \in \Theta_j} \left| \widehat{Q}_{(r)}(\theta) - Q_{(r)}(\theta) \right| \geq \eta \right)$.

Let θ_j denote a point in Θ_j . Noticing that,

$$\begin{aligned} \left| \widehat{Q}_{(r)}(\theta) - Q_{(r)}(\theta) \right| &\leq \left| \widehat{Q}_{(r)}(\theta_j) - Q_{(r)}(\theta_j) \right| + \left| \widehat{Q}_{(r)}(\theta) - \widehat{Q}_{(r)}(\theta_j) - Q_{(r)}(\theta) + Q_{(r)}(\theta_j) \right| \\ &\leq \left| \widehat{Q}_{(r)}(\theta_j) - Q_{(r)}(\theta_j) \right| + \frac{\epsilon}{|\Xi|} \left| \sum B(X_{rs}) - E[B(X_{rs})] \right| + 2\epsilon E[B(X_{rs})], \text{ and} \end{aligned}$$

$$P \left(\sup_{\theta \in \Theta_j} \left| \widehat{Q}_{(r)}(\theta) - Q_{(r)}(\theta) \right| \geq \eta \right) \leq P \left(\left| \widehat{Q}_{(r)}(\theta_j) - Q_{(r)}(\theta_j) \right| \geq \frac{\eta}{3} \right) + P \left(\left| \Xi \right|^{-1} \left| \sum_{s \in \Xi} (B(X_{sr}) - EB(X_{sr})) \right| \geq \frac{\eta}{3\epsilon} \right) = o(|\Xi|^{-k})$$

by Lemma B.8, the result holds as $R = o(|\Xi|^k)$. \square

LEMMA B.8. *For each r , suppose $\{X_{sr} : s \in \Xi\}$ be covariates satisfying Assumption 6. Let $R = O(|\Xi|^h)$, and let h, k, p, γ , with k, p, γ defined below, satisfy $h + 1 \leq k < p/2 - \gamma pd$. Then $\forall \eta > 0$,*

$$\max_r P \left(\left| |\Xi|^{-1} \sum_{s \in \Xi} X_{sr} \right| > \eta \right) = o(|\Xi|^{-k}) \text{ and } P \left(\max_r \left| |\Xi|^{-1} \sum_{s \in \Xi} X_{sr} \right| > \eta \right) = o(|\Xi|^{-1}).$$

Proof. The argument follows Lemma 1 of Hall and Horowitz (1996) and Lemma 2 of Supplementary Appendix I of Hahn and Newey (2004).

Step 1: By Chebyshev's inequality

$$P \left(\left| |\Xi|^{-1} \sum_{s \in \Xi} X_{sr} \right| > \eta \right) \leq \frac{C \cdot E \|X_{1r}\|^{p+\delta} \left(|\Xi|^{p/2} m^{dp} + |\Xi|^p \alpha_{2,\infty}^r(m)^{\frac{\delta}{p+\delta}} \right)}{\eta^p |\Xi|^p},$$

for $1 \leq m \leq C(p) |\Xi|$ where the second inequality follows from Lemma B.9, which we can write under Assumption 6.⁶³ We can bound the right hand side by $\eta^{-p} C \cdot E \|X_{1r}\|^{p+\delta} m^{dp} |\Xi|^{-p/2} + \eta^{-p} C \cdot E \|X_{1r}\|^{p+\delta} \alpha_{2,\infty}^r(m)^{\frac{\delta}{p+\delta}}$ and for $m = |\Xi|^\gamma$ for some γ with $0 < \gamma \leq 1$, using the bound $\sup_r \alpha_{\infty,1}^r(m) \leq Ca^m$ on the mixing coefficient,

$$\max_r |\Xi|^k P \left(\left| |\Xi|^{-1} \sum_{s \in \Xi} X_{sr} \right| > \eta \right) \leq \eta^{-p} C \cdot \max_r E \|X_{1r}\|^{p+\delta} \left(|\Xi|^{\gamma pd+k-p/2} + |\Xi|^k a^{\frac{\delta}{p+\delta} |\Xi|^\gamma} \right).$$

⁶³We write the proof for $\sup_r \alpha_{\infty,2}^r(m) \leq Ca^m$, though the extension to $\sup_r \alpha_{\infty,2}^r(m) = o(m^{-d})$ is straightforward and will merely result in more stringent requirements on r, k, γ, d . For Assumption 6(ii) we have

$$\max_r |\Lambda|^k P \left(\left| \frac{1}{|\Xi|} \sum_{s \in \Xi} X_{sr} \right| > \eta \right) = O \left(|\Xi|^{d\gamma p+k-p/2} + |\Xi|^{k-\gamma d-\gamma\epsilon} \right)$$

for $\alpha_{1,\infty}^r(m) = O(m^{-d-\epsilon}) = o(m^{-d})$. Then the requirement is $k < (p/2 + \gamma pd) \vee (d + \epsilon)$.

We have shown that for Assumption 6, $\max_r |\Xi|^k \mathbb{P} \left(\left| \frac{1}{|\Xi|} \sum_{s \in \Xi} X_{sr} \right| > \eta \right) = O \left(|\Xi|^{\gamma p d + k - p/2} \right) = o(1)$ if $\gamma p d + k < p/2$ which is the first result.

Step 2: By a union bound $|\Xi| \mathbb{P} \left(\max_r \left| \frac{1}{|\Xi|} \sum_{s \in \Xi} X_{sr} \right| > \eta \right) \leq |\Xi| R \cdot \max_r \mathbb{P} \left(\left| \frac{1}{|\Xi|} \sum_{s \in \Xi} X_{sr} \right| > \eta \right)$. If $R = O(|\Xi|^h)$, it follows that $|\Xi| \mathbb{P} \left(\max_r \left| \frac{1}{|\Xi|} \sum_{s \in \Xi} X_{sr} \right| > \eta \right) \leq O(|\Xi|^{h+1}) o(|\Xi|^{-k})$. Since $h+1 \leq k$, $O(|\Xi|^{h+1}) o(|\Xi|^{-k}) = o(1)$, $\mathbb{P} \left(\max_r \left| \frac{1}{|\Xi|} \sum_{s \in \Xi} X_{sr} \right| > \eta \right) = o(|\Xi|^{-1})$ which proves the result. \square

LEMMA B.9. *Let $\{X_i : t_i \in \Lambda \subset \mathbb{Z}^d\}$ be a mean zero stationary random field satisfying Assumption 6 and $\{X_{ij} : ij \in \Xi\}$ covariates. Then for any positive integer r and for $1 \leq m < C(k) \cdot |\Xi|$, we have $\mathbb{E} \left[\left(\sum_{ij \in \Xi} X_{ij} \right)^k \right] \leq C(r) \mathbb{E} \|X_1\|^{k+\delta} \left(|\Xi|^{k/2} m^{kd} + |\Xi|^k \alpha_{2,\infty}(m)^{\frac{\delta}{k+\delta}} \right)$.*

Proof. The proof builds on Lahiri (1992), with two differences: the first is an extension to mixing random fields,⁶⁴ and the second is that we are interested over moments of random variables on Ξ as opposed to Λ . The Lahiri (1992) style of argument proceeds in four parts; we include the entire argument for completeness though the key differences are in the last two steps. First, we can control the first $k/2$ terms via a standard result. This will enable us to bound this part of the sum by a $|\Xi|^{k/2}$ rate. Second, for the remaining terms, we will divide the space into a set of all pieces with a well separated point τ on the lattice whose random variable X_τ has power 1 and is at least of distance m from any other point in the collection, and into its complement. Third, we will control this set using the mixing coefficient and fourth, by a counting argument we create an upper bound on the number of points in the complement. It is useful to note that for $ij \in \Xi$, for $1 \in \Lambda$, $\mathbb{E} \|X_{ij}\|^k \lesssim \mathbb{E} \|X_1\|^k$ by the assumption and stationarity. Moreover, we defined the pseudo-metric $d_\Xi(ij, kl) := d_\Lambda(i, k) \wedge d_\Lambda(i, l) \wedge d_\Lambda(j, k) \wedge d_\Lambda(j, l)$ where $d_\Lambda(x, y) := \|x - y\|_1$.

Step 1: For $k = 2h$, we can expand the term into a polynomial,

$$\left(\sum_{s \in \Xi} X_s \right)^k = \sum_{j=1}^k \sum_{\alpha_1, \dots, \alpha_j} c(\alpha_1, \dots, \alpha_j) \sum_{s_1, \dots, s_j} \prod_{t=1}^j X_{s_t}^{\alpha_t}$$

where $t = 1, \dots, j$ is an arbitrary index of a j -tuple $(s_1, \dots, s_j) \subset \Xi$ and $c(\cdot)$ are coefficients. We can control the first $k/2$ -tuples by a standard argument, e.g., Bhattacharya and Rao (1986), making use of $\mathbb{E} \|X_{ij}\|^k \lesssim \mathbb{E} \|X_1\|^k$,

$$\left| \sum_{j=1}^{k/2} \sum_{\alpha_1, \dots, \alpha_j} c(\alpha_1, \dots, \alpha_j) \sum_{s_1, \dots, s_j} \mathbb{E} \left[\prod_{t=1}^j X_{s_t}^{\alpha_t} \right] \right| \leq C(k) |\Xi|^{k/2} \mathbb{E} \|X_1\|^k.$$

In what follows it suffices to show for fixed $j > k/2$ and $(\alpha_1, \dots, \alpha_j)$,

$$\left| \sum_{s_1, \dots, s_j} \mathbb{E} \left[\prod_{t=1}^j X_{s_t}^{\alpha_t} \right] \right| \leq C(k) \mathbb{E} \|X_1\|^{k+\delta} \left(|\Xi|^{k/2} m^{kd} + |\Xi|^k \alpha_{2,\infty}(m)^{\frac{\delta}{k+\delta}} \right).$$

Step 2: Next, we create a set that counts the sites s_τ where X_{s_τ} has power $\alpha_\tau = 1$ and s_τ is sufficiently far from the other s_t in the j -tuple. Let $u := j - k/2$. We put $A := \{t : \alpha_t = 1\}$ as the set of all points that have coefficient 1. Then let $\beta_0 = |A|$. This is a set that counts the number of indices that show up exactly one time. We want to show that this set is non-empty. Note that $1 \leq u \leq k/2$. Also, since $k = \sum_{t=1}^j \alpha_t \geq \beta_0 + 2(j - \beta_0)$, this means $2u \leq \beta_0 \leq k$ and therefore $(j - k/2) \leq \beta_0 \leq k/2$.

Then we partition the set of all j -tuples into B_m and B_m^c , with In Lahiri's notation, $\sum_2 \equiv \sum_{s_1, \dots, s_j} = \sum_3 + \sum_4$. To define the sets, put

$$B_m := \{(s_1, \dots, s_j) : \inf_{l \neq t} d_\Xi(s_l, s_t) = d_\Lambda(i_{s_l}, k_{s_t}) \wedge d_\Lambda(j_{s_l}, k_{s_t}) \wedge d_\Lambda(i_{s_l}, l_{s_t}) \wedge d_\Lambda(i_{s_l}, l_{s_t}) > m \text{ for some } t \in A\}.$$

Now \sum_3 sums over the terms in B_m and \sum_4 over the terms in B_m^c .

Step 3: We want to control B_m . Fix $\tau \in A$. Then decompose $\prod_{t=1}^j X_{s_t}^{\alpha_t} = X_a X_b$, with $X_b = X_{s_\tau}$ and $X_a = \prod_{t \neq \tau} X_{s_t}^{\alpha_t}$. Then X_a is a random field with respect to $\sigma(X_{s_1}, \dots, \widehat{X}_{s_\tau}, \dots, X_{s_j})$, where we use notation to indicate the omission of a term, $(a, \widehat{b}, c) := (a, c)$, and X_b with respect to $\sigma(X_{s_\tau})$. These are of size $j - 1$ and 1,

⁶⁴Hahn and Kuersteiner (2004) extend the Lahiri argument to the time series setting.

respectively. By definition the distance between these two sets is at least m , so by applying Lin et al. (1996) and using the fact that $\mathbb{E}[X_{ij}X_{kl}] \leq \|X_{ij}\|_{L^p(\mathbb{P})} \cdot \|X_{kl}\|_{L^q(\mathbb{P})} \alpha_{2,2}^{1-q^{-1}-p^{-1}}(d_{\Xi}(ij,kl))$, we have

$$\left| \mathbb{E} \left[\prod_{t=1}^j X_{s_t}^{\alpha_t} \right] \right| \leq |\mathbb{E}[X_a X_b]| \leq C \cdot \|X_a\|_{L^p(\mathbb{P})} \cdot \|X_b\|_{L^q(\mathbb{P})} \alpha_{2(j-1),2}^{1-q^{-1}-p^{-1}}(m).$$

Taking $p = k + \delta$ and $\frac{1}{q} = \frac{k-1}{k+\delta}$, using stationarity and repeatedly applying the Holder inequality,

$$\left| \mathbb{E} \left[\prod_{t=1}^j X_{s_t}^{\alpha_t} \right] \right| \leq C \cdot \|X_a\|_{L^{k+\delta}(\mathbb{P})} \cdot \|X_b\|_{L^{(k-1)/(k+\delta)}(\mathbb{P})} \alpha_{2(j-1),2}^{\delta/(k+\delta)}(m) \leq C \cdot \mathbb{E} \|X_1\|^{k+\delta} \alpha_{2(j-1),2}^{\delta/(k+\delta)}(m).$$

Since $\alpha_{j,1}(m) \leq \alpha_{j+1,1}(m)$, $\left| \mathbb{E} \left[\sum_{B_m} \prod_{t=1}^j X_{s_t}^{\alpha_t} \right] \right| \leq C \cdot \mathbb{E} \|X_1\|^{k+\delta} |\Xi|^k \alpha_{2,\infty}(m)^{\frac{\delta}{k+\delta}}$.

Step 4: Finally, we control \bar{B}_m^c . To get a (coarse) upper bound, first notice the maximum number of powers of 1 that can be placed is $2u$ ⁶⁵. Construct a set $\Gamma \subset \{1, \dots, j\}$, $|\Gamma| = 2u$ which will include all powers of one and perhaps some residual copies of terms with higher power. Then we will simply count

$$\bar{B}_m^c := \left\{ (s_1, \dots, s_j) : \inf_{l \neq t} d_{\Xi}(s_t, s_l) = d_{\Lambda}(i_{s_l}, k_{s_t}) \wedge d_{\Lambda}(j_{s_l}, k_{s_t}) \wedge d_{\Lambda}(i_{s_l}, l_{s_t}) \wedge d_{\Lambda}(i_{s_l}, l_{s_t}) \leq m \ \forall t \in \Gamma \right\}.$$

It will help us to define the *partners* of s_t as $\epsilon(t) \in \arg \inf_l d_{\Xi}(s_t, s_l)$. Then define $\mathcal{P}(\Gamma)$ as the set of partners of $t \in \Gamma$ and put $v = |\Gamma \cup \mathcal{P}(\Gamma)|$. This is the number of distinct elements of Ξ with power one or power greater than one whose copies form the residual members of Γ and their partners. By definition $2u \leq v \leq 4u \wedge j$ since $v \geq |\Gamma|$ and $v \leq 2|\Gamma|$ if each had a distinct partner.

We define an *m-unbroken set* as a collection of points in $\Gamma \cup \mathcal{P}(\Gamma)$ for which each member is within an m -distance under pseudo-metric d_{Ξ} . Put q as the number of m -unbroken sets. First, observe that there are less than $|\Xi|^q$ initial sites to place a seed for each of the unbroken sets.⁶⁶ Next, we have to place each of the $v - q$ terms. For a given ij there are less than $2(2m + 1)^d$ elements of Ξ within an m -distance since $d_{\Lambda}(k, i) \wedge d_{\Lambda}(l, i) \wedge d_{\Lambda}(k, j) \wedge d_{\Lambda}(l, j) \leq m$. An upper bound is $O(m^{(v-q)d})$. Finally, we can arbitrarily place the remaining $j - v$ elements yielding $|\Xi|^{j-v}$. This gives us $|\bar{B}_m^c| \lesssim |\Xi|^{q+j-v} m^{(v-q)d}$. The same counting exercise as in Lahiri (1992) gives us $j + q - v \leq h$ and $v - q \leq k$ which yields $|\Xi|^{q+j-v} m^{(v-q)d} \leq |\Xi|^h m^{kd}$. This concludes the proof. \square

B.4. Groups Model. The proof of Proposition 4.2 is a corollary to the following lemma. We need to replicate the arguments from section B.2 replacing Lemma B.1 by Lemma B.10. The model satisfies Assumption (2) below by Chatterjee et al. (2010).

LEMMA B.10. Assume (1) for all $\theta_r \in \Theta_r$, $\zeta_r(\bar{\theta}_r, u)$ are $q \times k_n q$ matrix (or vector) valued functions, $\sup_{u \in \mathcal{U}} \sup_r \|\zeta_r(\theta_r, u)\| \leq B_r$ with $\text{plim} \mathbb{E}_{nR}[B_{ir}] < \infty$, (2) for each r , $\sup_{i \leq k_n} \|\hat{\theta}_{ir} - \theta_{0ir}\| \leq C(L) \sqrt{\frac{\log n}{n}}$ with probability at least $1 - cn^{-2}$, where L, c depend only on $\Theta_r \subset \mathbb{R}^{k_n}$ and ψ , (3) $\Theta_r \subset [a, b]^{k_n}$ for all r , and (4) $R = O(n^\gamma)$ with $\gamma < 1$. Then

$$\sup_{u \in \mathcal{U}} R^{1/2} \left\| \mathbb{E}_R \zeta_r(\bar{\theta}_r, u) (I_q \otimes (\hat{\theta}_r - \theta_{0r})) \right\|_2 \lesssim \sqrt{\frac{R \cdot \log n}{n}}$$

with probability approaching one. Specifically the probability is at least $(1 - cn^{-2})^{R_n}$.

Proof of Lemma B.10. Step 1: This follows from $\|\cdot\|_2 \leq \|\cdot\|_1$ and

$$\begin{aligned} \sup_{u \in \mathcal{U}} R^{1/2} \left\| \mathbb{E}_R \zeta_r(\bar{\theta}_r, u) (I_q \otimes (\hat{\theta}_r - \theta_{0r})) \right\|_1 &\leq \sup_{u \in \mathcal{U}} R^{1/2} \mathbb{E}_R \left[\left\| \zeta_r(\bar{\theta}_r, u) \right\|_1 \left\| (I_q \otimes (\hat{\theta}_r - \theta_{0r})) \right\|_\infty \right] \\ &\leq \sup_{u \in \mathcal{U}} R^{1/2} \mathbb{E}_R \left[\left\| \zeta_r(\bar{\theta}_r, u) \right\|_1 \left\| \hat{\theta}_r - \theta_{0r} \right\|_\infty \right] \\ &\leq \sup_{u \in \mathcal{U}} \mathbb{E}_R \left\| \zeta_r(\bar{\theta}_r, u) \right\|_1 \cdot \sqrt{R} \sup_{r \leq R} \left\| \hat{\theta}_r - \theta_{0r} \right\|_\infty. \end{aligned}$$

We have that $\sup_{u \in \mathcal{U}} \frac{1}{\sqrt{q}} \left\| \zeta_r(\bar{\theta}_r, u) \right\|_1 \leq \sup_{u \in \mathcal{U}} \left\| \zeta_r(\bar{\theta}_r, u) \right\|_2 \leq B_r$ and therefore

$$\sup_{u \in \mathcal{U}} R^{1/2} \left\| \mathbb{E}_R \zeta_r(\bar{\theta}_r, u) (I_q \otimes (\hat{\theta}_r - \theta_{0r})) \right\|_1 \lesssim \sqrt{R} \sup_{r \leq R} \left\| \hat{\theta}_r - \theta_{0r} \right\|_\infty + o_{\mathbb{P}}(1)$$

⁶⁵Let x be the maximal number. Then $k = x + 2(j - x)$. Solving this yields the result.

⁶⁶We leave ourselves open to certain double-counting; the bound is coarse.

by assumption on the envelope functions. It suffices to control $\sup_{r \leq R} \left\| \widehat{\theta}_r - \theta_{0r} \right\|_{\infty}$.

Step 2: Observe that $\sup_{r \leq R} \left\| \widehat{\theta}_r - \theta_{0r} \right\|_{\infty} \leq C(L) \sqrt{\frac{\log n}{n}}$ with probability at least $(1 - cn^{-2})^{Rn}$, since the bound holds for each of the R (independent) terms with probability at least $1 - cn^{-2}$. This approaches one for $\gamma < 1$. \square

B.5. DeJong and Woutersen (2011) Model.

Proof of Proposition 4.3. We check that the conditions of Lemma 4.1 hold. First, by example 2 of Hahn and Kuersteiner (2004), using Theorem 2 of De Jong and Woutersen (2010) the model satisfies Conditions 1-7 of Hahn and Kuersteiner (2004). Notice $X_{s,r} = [A_{s,r}, A_{s-1,r}, \dots, A_{s-k,r}, z_{sr}]$ satisfies condition 3 of Hahn and Kuersteiner (2004). It is straightforward to verify conditions 1-5 of Lemma 4.1 hold. Condition 1 follows from Theorem 3 of De Jong and Woutersen (2010) and checking that the probit model is twice continuously differentiable in θ . By an argument identical to Lemma B.7, condition 2 holds. The advantage here is that we can directly apply Lemma 1 of Hahn and Kuersteiner (2004) instead of Lemma B.9, since we essentially are dealing with time-series type data. Since

$$q(z_{sr}, A_r^{s-1}; \theta_{0r}) = A_s \log \Phi \left(\sum \theta_{0r,k} A_{s-k,r} + \theta'_{0r,l+1} z_{sr} \right) + (1 - A_s) \log \left(1 - \Phi \left(\sum \theta_{0r,k} A_{s-k,r} + \theta'_{0r,l+1} z_{sr} \right) \right),$$

we have

$$v(z_{sr}, A_r^{s-1}; \theta_{0r}) = A_s \frac{\phi \left(\sum \theta_{0r,k} A_{s-k,r} + \theta'_{0r,l+1} z_{sr} \right)}{\Phi \left(\sum \theta_{0r,k} A_{s-k,r} + \theta'_{0r,l+1} z_{sr} \right)} - (1 - A_s) \frac{\phi \left(\sum \theta_{0r,k} A_{s-k,r} + \theta'_{0r,l+1} z_{sr} \right)}{1 - \Phi \left(\sum \theta_{0r,k} A_{s-k,r} + \theta'_{0r,l+1} z_{sr} \right)}$$

which is bounded in order by $1 + \left| \sum \theta_{0r,k} A_{s-k,r} + \theta'_{0r,l+1} z_{sr} \right|$ using Feller's inequality. Thus the envelope depends on X_{sr} only, which satisfies the mixing condition 3 of Hahn and Kuersteiner (2004). By compactness of Θ_r and the moment assumption on z , we can apply the argument of Lemma B.6, with $|\Xi| = \binom{n}{2}$ and $v(X_{sr}, A_r^{s-1}; \theta_r)$ as above to show condition 3 holds. Conditions 4 and 5, which are a Lipschitz condition on the Hessian and a pointwise LLN (uniform across R) on the Hessian hold because for a probit model all finite order of derivatives of $\log \Phi(\cdot)$ exist and are continuous, on a compact parameter space they satisfy the Lipschitz condition and are uniformly bounded. \square

APPENDIX C. OVERVIEW OF ESTIMATION ALGORITHM AND STANDARD ERRORS

We present a non-technical overview of the standard errors and estimation procedures used. A theoretical discussion of standard error estimation is beyond the scope of this paper. Here, we discuss standard errors and finite-sample simulation bias. A lengthier formal, technical discussion and simulations results are available from the authors upon request. Several approaches can be used for inference: heteroskedasticity-robust, clustered, block bootstrapped, first-stage bootstrapped, and importance sampled standard errors. Network-level inference can be performed using heteroskedasticity-robust standard errors under cross-network independence, even in the presence of sampling. Individual-level inference under sampling requires attention to within-network, cross-individual autocorrelation exacerbated by the reconstructed regressor's mismeasurement, what we call the reconstruction error. When an edge is missing in the computation of one individual's network statistic, that edge is also missing for all other individuals' network statistic computation. Clustered standard errors and their nonparametric block bootstrap analog both account for heteroskedasticity and within-graph autocorrelation. Additionally, a parametric bootstrap which simulates the autocorrelation from the reconstruction error can estimate the contribution of the sampling-induced autocorrelation to the reconstruction estimator's variance. The parametric structure of this bootstrap should allow these standard errors to be estimated on a smaller collection of networks than necessary for clustered standard errors.

While the graphical reconstruction estimator is formally a "two-step" estimator, super-consistency allows the researcher to ignore the first-stage estimation uncertainty. That said, first-stage uncertainty can be understood using a first-stage parametric bootstrap and importance sampling. The first-stage parametric bootstrap estimates the distribution of the first-stage parameters and encompasses the following effects: sample from the observed collection of networks with replacement, estimate each network's first-stage parameters using the observed data, simulate a new collection of observed graphs using the first-stage estimates, and estimate the first-stage parameters using the simulated graphs. Repeat this process to obtain the bootstrapped variance of the first stage estimators to determine

whether super-consistency is justified. Importance sampling can reduce the computational burden of bootstrapping the first-stage estimates by using the first-stage model's likelihood to rebalance a collection of reconstructed network statistics and greatly reduces the cost of accounting for any first-stage variability.

With or without first-stage superconsistency, the reconstruction estimator still uses a finite number of simulations, leaving another error term which we call the simulation error. A regularization to estimate and subtract out the variance in the regressors due to this noise can be employed.

APPENDIX D. DISCUSSION OF ANALYTICAL CORRECTIONS WITH REGULARIZATION

D.1. A Regularized Estimator. We show a simple example where our analytical correction is consistent and does not require us to assume $\sigma_v^2 \rightarrow 0$. In section 3, for degree and graph clustering, we have shown that $\tilde{\beta} = \Phi \hat{\beta} \xrightarrow{P} \beta_0 \frac{\sigma_{(x)}^2}{\sigma_{(x)}^2 + \Phi^{-2} \sigma_v^2}$ where Φ is a deterministic function of sampling rate ψ . We now show we can construct estimates $\hat{\sigma}_{(x)}^2$ and $\hat{\sigma}_v^2$ and therefore develop a consistent analytical correction β^* . We present a simple example to illustrate the argument. Our example is average degree and for simplicity assume the researcher samples every edge independently with probability ψ . It is easy to see $E[d(\bar{G})|G] = \psi d(G)$, where \bar{G} is the sampled graph. Putting $v = d(\bar{G}) - \psi d(G)$, we can analytically compute $\text{var}(v|G) = \frac{2}{n} \psi (1 - \psi) d(G)$.

LEMMA D.1. $\text{var}(v|G) = 2n^{-1} \psi (1 - \psi) d(G)$.

Proof. Let $T = \binom{n}{2}$ and e index edges going from $e = 1, \dots, \binom{n}{2} = T$. Notice average degree is $d(G) = \frac{2}{n} \sum_e A_e$. Then the $v = \frac{2}{n} \sum_e \chi_e A_e - \psi \frac{2}{n} \sum_e A_e$. Notice

$$\begin{aligned} E \left[\frac{2}{n} \sum_e \chi_e A_e \right]^2 &= E \left[4n^{-2} \left\{ \sum_e \chi_e A_e + 2 \sum_{e < e'} \chi_e \chi_{e'} A_e A_{e'} \right\} \right] = 4n^{-2} \left\{ \psi \sum_e A_e + \psi^2 2 \sum_{e < e'} A_e A_{e'} \right\} \\ &= 2\psi n^{-1} d(G) + 4n^{-2} \psi^2 2 \sum_{e < e'} A_e A_{e'}. \end{aligned}$$

Meanwhile, average degree squared is $\left(\frac{2}{n} \sum_e A_e \right)^2 = 4n^{-2} \left\{ \sum_e A_e + 2 \sum_{e < e'} A_e A_{e'} \right\}$, which is useful since $(E[v|G])^2 = \left(\psi \frac{2}{n} \sum_e A_e \right)^2 = \psi^2 d(G)^2$. Therefore,

$$E[v^2|G] = 2\psi n^{-1} d(G) + \psi^2 \left\{ 4n^{-2} 2 \sum_{e < e'} A_e A_{e'} \right\} = 2n^{-1} \psi (1 - \psi) d(G) + \psi^2 d(G)^2.$$

It follows that $\text{var}(v|G) = 2n^{-1} \psi (1 - \psi) d(G) + \psi^2 d(G)^2 - \psi^2 d(G)^2 = 2n^{-1} \psi (1 - \psi) d(G)$. \square

With the analytical formula for the variance of v , we can compute β^* . Let $X := \psi \cdot (d(G_1), \dots, d(G_R))'$ denote the scaled vector of true (unobserved) average degrees, $Z := (d(\bar{G}_1), \dots, d(\bar{G}_R))$ the observed vector of sampled degrees, and $V := (v_1, \dots, v_R) = Z - X$. Then it is clear that $\beta^* := (Z'Z - V'V)^{-1} Z'y$ is consistent for β_0 . By estimating $\Sigma_v := \text{plim } V'V/R$, we have $\beta^* := \left(Z'Z - R\hat{\Sigma}_v \right)^{-1} Z'y \xrightarrow{P} \beta_0$. Under mild regularity conditions on the growth of average degree, we may therefore estimate β^* . Estimation can be improved by performing a bootstrap bias correction.

D.2. Numerical Evidence. Table 8 displays simulation results for network level regression of average degree with simulated outcomes, precisely of the form presented in Table 1. We choose ψ to make the expected edge count comparable to that of the induced subgraph. The results confirm the fact that the naive estimator exhibits significant biases, the analytical correction vastly reduces the biases but may still retain some residual bias which can further be mitigated by applying regularization.