My friend far, far away: a random field approach to exponential random graph models

VINCENT BOUCHER[†] AND ISMAEL MOURIFIÉ[‡]

[†]*Department of Economics, Université Laval, 1025, Avenue des Sciences-Humaines, Québec City, Québec, G1V 0A6, Canada.*

E-mail: vincent.boucher@ecn.ulaval.ca

[‡]*Department of Economics, University of Toronto, 150 St George Street, Toronto, Ontario, M5S 3G7, Canada.*

E-mail: ismael.mourifie@utoronto.ca

First version received: May 2016; final version accepted: May 2017

Summary We explore the asymptotic properties of strategic models of network formation in very large populations. Specifically, we focus on (undirected) exponential random graph models. We want to recover a set of parameters from the individuals' utility functions using the observation of a single, but large, social network. We show that, under some conditions, a simple logit-based estimator is coherent, consistent and asymptotically normally distributed under a weak version of homophily. The approach is compelling as the computing time is minimal and the estimator can be easily implemented using pre-programmed estimators available in most statistical packages. We provide an application of our method using the Add Health database.

Keywords: *Homophily, Network formation, Random fields, Spatial econometrics.*

1. INTRODUCTION

How do social networks form? Specifically, how can we measure the influence of an individual's socio-economic characteristics on the identity of his peers?

We know that many social networks exhibit strong racial or religious segregation; see, for instance, Echenique and Fryer (2007), Watts (2007) and Mele (2017). This raises many interesting questions regarding the cause of this segregation. For instance, we would like to be able to distinguish between the impact of individuals' characteristics (e.g. race) and the impact of individuals' positions in the networks (e.g. popularity).

The shapes of social networks also have measurable effects on individuals' choices. Many studies show that an individual's peers can have a significant influence on his actions, ranging from unhealthy consumption choices – see, e.g. Fortin and Yazbeck (2015) and references therein – to labour-force participation – see, e.g. Toomet et al. (2013), Patacchini and Zenou (2012) and Mourifié and Siow (2014).¹ However, as most social networks are endogenously

¹ See de Paula et al. (2017a) and Boucher and Fortin (2016) for recent reviews.

formed, the estimated influence of peers is likely to be biased.² An understanding of how the networks are formed could allow us to control for this endogeneity, and it might suggest policy instruments that would help to influence network formation processes.

In this paper, we use the network formation model of Mele (2017), for which the joint distribution is an exponential random graph model (ERGM). The estimation of ERGMs is typically challenging because of the intractable denominator of the joint distribution. We use random field theory to provide a simple pseudo maximum likelihood estimator (MLE), as in Besag (1975) and Strauss and Ikeda (1990), based on the product of the conditional distributions. To our knowledge, this paper is the first to apply random field theory to empirical models of network formation. The approach is promising and has been recently employed by Leung (2014) in a similar context. Our approach only requires the observation of a single, potentially very large, network.

Our model allows for a large set of admissible preferences, which are characterized by intuitive conditions. Specifically, we show that our estimator is consistent and asymptotically normally distributed, provided that individuals' preferences exhibit a weak version of homophily. Homophily is one of the most robust empirical characteristics of social networks. It formalizes the observation that similar individuals are more likely to interact with each other. As homophily is featured in both theoretical – see, e.g. Boucher (2015), Bramoullé et al. (2012) and Currarini et al. (2009, 2010) – and empirical – see, e.g. Mele (2017) and Christakis et al. (2010) – models of network formation, our methodology is applicable to many existing models of network formation. We apply this new methodology to the formation of friendship networks by US teenagers.

A fundamental challenge in estimating a network formation process is the highly dependent nature of most socio-economic relationships. Consider a friendship: the probability that Alice and Bob are friends depends on their individual characteristics. However, it can also depend on Bob's friendship with Charlotte (who perhaps does not like Alice). The probability that Alice and Bob are friends might then depend on Charlotte's individual characteristics (as well as on her other friendships).

If individuals have homophilic preferences, then the probability that Alice and Bob are friends should be primarily influenced by individuals similar to them. If Alice and Bob are high-school teenagers, for instance, then the probability that they become friends increases if they go to the same school, or if they attend the same classes. Accordingly, if Bob and Charlotte are friends, then there is a greater probability that they go to the same school, or at least live in the same country. Donald, an elderly man, living in a different country (hence having very different individual characteristics from Alice, Bob and Charlotte) probably does not influence the probability that Alice and Bob become friends. We generalize this argument, and we show that homophily implies a generalization of the ϕ -mixing property used in time-series and spatial econometric models. This fact allows us to define a consistent estimation strategy based on a pseudo-MLE.

This paper contributes to the empirical literature on strategic network formation.³ Some papers require the observation of many (mostly independent) social networks; see, e.g. Boucher (2015), Currarini et al. (2009, 2010) and Sheng (2014). However, one limitation of these approaches is that they require the observation of many (an asymptotically infinite number of)

² The literature on peer effects has recently explicitly considered the endogeneity of social networks; see, e.g. Badev (2013), Boucher (2016), Goldsmith-Pinkham and Imbens (2013) and Hsieh and Lee (2016).

³ See Chandrasekhar (2016) for an excellent recent review.

independent social networks, which is not always available in existing databases. As the data requirement is large, the computational burden is also important.

Accordingly, many recent approaches only require the observation of a single network, at one point in time; see, e.g. Graham (2017) Leung (2014, 2015) and de Paula et al. (2017a). In particular, a significant part of that literature focuses on ERGMs, which we formally define in the next section. Chandrasekhar and Jackson (2014, 2016), as well as Mele (2017), followed by Badev (2013), present strategic models of networks formation leading to ERGMs, and they focus on Bayesian inference.⁴

Chandrasekhar and Jackson (2014, 2016) focus on asymptotic theory and present conditions on the structure of the network, which allow for a consistent estimation of the parameters. By focusing on subgraphs, and under the assumption that the network is sparse, they develop a simple estimator, based on the iterative counting of subgraphs.

Alternatively, Graham (2016, 2017) focus on logit-based models allowing for individual fixed unobserved characteristics. An innovative feature in Graham (2016) is the focus on a dynamic setting where links can be dependent on the previous state of the network. He shows that conditional likelihood, for links that are stable across time, has a tractable denominator.

We contribute to this body of literature by providing an explicit, easy-to-implement pseudo-MLE that requires the observation of only one social network, at one point in time. We introduce a weakened notion of homophily, and we show that it implies that our pseudo-MLE is consistent and asymptotically normally distributed. To do so, we use the law of large numbers and the central limit theorem, following Jenish and Prucha (2009).

The rest of the paper is organized as follows. In Section 2, we present our micro-economic and econometric framework. In Section 3, we review some elements of random fields theory. In Section 4, we derive the asymptotic distribution of our estimator, and in Section 5, we define a class of network formation models suited to our econometric framework. In Section 6, we discuss the identification of the model. In Section 7, we provide an application using the Add Health database. We conclude in Section 8.

2. THE ECONOMY

Let $N = \{1, \dots, n\}$ be the set of individuals. Each individual is characterized by a non-stochastic vector of quantitative characteristics $x_i = (x_i^1, \dots, x_i^T) \in \mathcal{X} \subset \mathbb{R}^T$, where $T \geq 1$. We define the distance between two individuals as $d(i, j) = d(x_i, x_j)$, where d is a distance on $\mathbb{R}^T$.⁵

In general, the choice of this distance function will be context-dependent (see Example 2.1 and our empirical application in Section 7). In particular, the distance can represent spatial preferences of individuals.⁶ We use $X = (x_1, \dots, x_n) \in \mathcal{X}^n$ to denote the matrix of individual characteristics.

We assume that individuals interact in an undirected social network. Let $m = (n(n-1))/2$ be the number of possible unordered pairs of individuals (i, j) for $i \neq j$ in the economy. The set of all pairs will be denoted by S_m , with typical elements $s, r \in S_m$. For any pair (i, j) , we let

⁴ See also Mele and Zhu (2016).

⁵ The model can be easily extended to include stochastic and qualitative characteristics, as long as they do not affect the distance. This is discussed in Section 7.

⁶ See Henry and Mourifié (2013) for a discussion of spatial preferences in Euclidean space.

$w_{ij} = 1$ if $i \in N$ and $j \in N$ are linked by a socio-economic relationship (e.g. friendship), and $w_{ij} = 0$ otherwise.

The collection of all relationships, the network, is denoted $W_m = \{w_s\}_{s \in S_m}$. For any subset of pairs $\hat{S} \subset S_m$, we denote the collection of relationships in $\hat{S}$ as $W_{\hat{S}} = \{w_s\}_{s \in \hat{S}}$. Finally, the set of relationships in $S_m \setminus \{s\}$ is denoted by $W_{m,-s}$.

We assume that the network W_m is endogenous and determined as a function of individuals' utilities. An individual has preferences over the set of characteristics and the network structure in the economy: $u_i(W_m; \theta)$, where $\theta \in (\theta^1, \dots, \theta^K) \in \Theta$ is the set of parameters to be estimated. We assume that Θ is a compact subset of $\mathbb{R}^K$, for $K \geq 1$. Here, the utility function is best interpreted in non-transferable utility setting in the sense that we do not explicitly model transfers between individuals. Assumption 2.1 resumes the structure of an individual's preferences.

ASSUMPTION 2.1. (PREFERENCES) *For any $i, j \in N$, there exists a function $h_{ij}(W_{m,-ij}; \theta)$, three times continuously differentiable in θ , such that*

$$\begin{aligned} u_i(1, W_{m,-ij}; \theta) - u_i(0, W_{m,-ij}; \theta) &= u_j(1, W_{m,-ij}; \theta) - u_j(0, W_{m,-ij}; \theta) \\ &= h_{ij}(W_{m,-ij}; \theta). \end{aligned}$$

Note that the function $h_{ij}(W_{m,-ij}; \theta)$ can be interpreted as the value of a socio-economic relationship between i and j , and it is assumed to be symmetric with respect to its indices. The payoff h_{ij} can be best interpreted as the payoff of a local public good, equally valued by i and j . This assumption is imposed for identifiability issues as we focus on an undirected network.⁷

We also make the following assumption.

ASSUMPTION 2.2. (PREFERENCES) *For all pairs (i, j) , we assume that there exist $\chi_{ij}(\theta)$, $\{\tilde{\chi}_{ij,ik}(\theta)\}_k$ and $\{\tilde{\chi}_{ij,jk}(\theta)\}_k$, symmetric with respect to their indices, such that*

$$h_{ij}(W_{m,-ij}, \theta) = \chi_{ij}(\theta) + \sum_{k \neq i, j} \tilde{\chi}_{ij,ik}(\theta) w_{ik} + \tilde{\chi}_{ij,jk}(\theta) w_{jk},$$

where $\tilde{\chi}_{ij,ik}(\theta) = 0$ whenever $\max\{d(i, j), d(i, k)\} > \bar{d}$, where $\bar{d} > 0$ is an arbitrarily large (but finite) number, and similarly for $\tilde{\chi}_{ij,jk}(\theta)$.

The assumption that $\tilde{\chi}_{ij,ik}(\theta) = 0$ whenever $\max\{d(i, j), d(i, k)\} > \bar{d}$ implies that the individuals cannot receive positive payoffs from an infinite number of links, and it ensures that $h_{ij}(W_{m,-ij}, \theta) < \infty$.⁸ Note that this is weaker than assuming that individuals only have a finite number of links, as for instance in de Paula et al. (2017a). Indeed, at this point, individuals can have an infinite number of links; however, only a finite number of those links can affect the payoff of any particular link. We discuss the implications of this assumption in detail in Section 6.2.

The specific structure in Assumption 2.2 serves two purposes. First, following Mele (2017), it guarantees the existence of a potential function (Monderer and Shapley, 1996), that is, a function $Q(W_m)$ such that for all $i, j \in N$, $Q(1, W_{m,-ij}; \theta) - Q(0, W_{m,-ij}; \theta) = h_{ij}(W_{m,-ij}; \theta)$. The potential function is developed explicitly in Proposition 3.1 in Section 3.

⁷ For undirected networks, one only observes if a link (i.e. between i and j) is created or not. If no link is observed, then it is therefore impossible to identify if i or j (or both) 'refused the link'. Then, if the payoffs were non-symmetric, one can only identify $\min\{h_{ij}, h_{ji}\}$.

⁸ This is true under Assumption 4.3, presented in Section 4. We thank an anonymous referee for this pointing this out to us.

Secondly, Assumption 2.2 allows for the definition of an estimator based on the product of the conditional distributions in the spirit of Besag (1975) and Strauss and Ikeda (1990). A formal treatment is provided in Section 3. Example 2.1 illustrates.

EXAMPLE 2.1. We use the following as a running example:

$$h_{ij}(W_{m,-ij}; \theta) = -\beta d(i, j) + \sum_{k \neq i, j} \alpha_{ik} w_{ik} + \alpha_{jk} w_{jk}.$$

Here, $\alpha_{ik} = \alpha$ if $d(i, k) < \bar{d}$ and 0 otherwise, and $\theta \equiv (\alpha, \beta)$. Note that we assume that x_i is uni-dimensional so that $d(i, j) = |x_i - x_j|$. We can also write

$$h_{ij}(W_{m,-ij}; \theta) = \{\alpha n_i(W_{m,-ij}) + \alpha n_j(W_{m,-ij}) - \beta |x_i - x_j|\},$$

where $n_i(W_{m,-ij})$ and $n_j(W_{m,-ij})$ are bounded and represent the number of links of i and j in the network, excluding the potential link between i and j . Then, if $\alpha > 0$, as an individual creates links, it increases the incentive to create additional links. We can verify that the implied utility function is

$$u_i(W_m; \theta) = \sum_{j \neq i} -w_{ij} \beta |x_i - x_j| + \sum_{j \neq i} \sum_{k < j, k \neq i} \alpha_{ik} w_{ij} w_{ik} + \alpha_{jk} w_{ij} w_{jk},$$

because $\alpha_{ij} = \alpha_{ik}$ for all $j, k \neq i$ under Assumption 2.2. Then, i receives the direct cost (if $\beta > 0$) of linking equal to $\beta |x_i - x_j|$, but benefits (if $\alpha > 0$) from positive spillover from his other links (i.e. $\alpha w_{ij} w_{ik}$), as well as from j 's other links (i.e. $\alpha w_{ij} w_{jk}$). In other words, in this example, the value of a link between i and j decreases with the distance between them, but increases with the number of friends they have, and it can be interpreted as preferences for popularity (see Mele, 2017). Moreover, note that Assumption 2.2 holds, as we can define $\chi_{ij}(\theta) = -\beta |x_i - x_j|$, $\tilde{\chi}_{ij,ik}(\theta) = \alpha_{ik}$ and $\tilde{\chi}_{ij,jk}(\theta) = \alpha_{jk}$.

The link formation process follows directly from Mele (2017), who presents an intuitive stochastic meeting process and assumes that the data are drawn from their stationary distribution.⁹ Starting from any network W_m^0 , the meeting process goes as follows for $t = 1, \dots, \infty$.¹⁰

STEP 1. Two individuals i and j meet with probability $\lambda(ij; W_{m,-ij}^{t-1}) > 0$.

STEP 2. Given such a meeting, preferences are affected by a temporary, idiosyncratic shock ε_{ij}^t that affects the value of a link between i and j . The shock ε_{ij}^t is assumed to follow a standardized logistic distribution.¹¹

STEP 3. Individuals i and j decide to update the status of their socio-economic relationship, letting $w_{ij}^t = 1$ iff $h_{ij}(W_{m,-ij}^{t-1}; \theta) + \varepsilon_{ij}^t \geq 0$.

Note that because preferences and shocks are symmetric (i.e. $h_{ij} = h_{ji}$ and $\varepsilon_{ij}^t = \varepsilon_{ji}^t$), both individuals must agree in Step 3. Here, ε_{ij}^t are temporary idiosyncratic shocks on the value of

⁹ See also Badev (2013) for an equilibrium interpretation.

¹⁰ The meeting process is, in fact, slightly adapted to our setting and notation, assuming a non-directed network, under Assumption 2.1. Note also that the process can be seen as a logit response game (Blume, 1993) on a game where each pair acts as a player.

¹¹ Note that this is equivalent to assuming that there are i.i.d. shocks $(\varepsilon_{ij}^0, \varepsilon_{ij}^1)$ distributed according to a type 1 extreme value distribution, such that $u_i(W_m; \theta) + w_{ij} \varepsilon_{ij}^1 + (1 - w_{ij}) \varepsilon_{ij}^0$, and similarly for j .

friendships, and they can be interpreted as temporary variations of the value of the local public good (i.e. friendship) produced by i and j .

Mele (2017) showed that this process converges to a stationary distribution, which can be written as

$$\mathbb{P}(W_m|\theta) = \frac{\exp\{Q(W_m;\theta)\}}{\sum_{W'_m} \exp\{Q(W'_m;\theta)\}}, \quad (2.1)$$

where $Q(W_m)$ is the potential function.

The denominator in (2.1) is intractable as it sums over the set of all possible networks, which has a cardinality of 2^m . This challenge has been the focus of most of the recent literature on empirical network formation models; see Chandrasekhar (2016) for a discussion. We discuss how we will deal with this intractability in the next section.

3. (MARKOV) RANDOM FIELDS

In this section, we briefly introduce some basic definitions and results from the theory of random fields.¹²

DEFINITION 3.1. *The random field $\{w_{s,m}; x_s \in \mathcal{X}_{S_m}, m \in \mathbb{N}\}$ is defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where $\Omega = \{0, 1\}^m$, $\mathcal{F}$ is a σ -algebra on Ω and $\mathbb{P}$ is a probability measure on Ω .*

In essence, a random field is a list of random numbers whose indices are mapped on to a space (here, of T dimensions). It can also be viewed as a generalization of a stochastic process. For a stochastic process, observations are located in ‘time’. Here, observations (i.e. pairs) are located in $\mathcal{X}$ (see Besag, 1974).¹³

In order to describe the dependence structure of the random field, it is convenient to define a neighbourhood, for each pair. Formally, a neighbourhood in a random field is defined as follows.

DEFINITION 3.2. *A neighbourhood $N_s \subset S_m \setminus \{s\}$ of a pair $s \in S_m$ is defined as the smallest collection of pairs such that*

$$\mathbb{P}(w_s|W_{m,-s}) = \mathbb{P}(w_s|W_{N_s}).$$

Formally, Assumption 2.2 allows for a natural definition of neighbourhoods on $\{w_{s,m}; x_s \in \mathcal{X}_{S_m}, m \in \mathbb{N}\}$. Define neighboured pairs as those that share an individual, or mathematically, $N_s = \{r \in S_m \setminus \{s\} : r_1 \in s \text{ or } r_2 \in s\}$. This is illustrated in Figure 1, where we consider the set of all pairs: $S_m = \{(i, j), (i, k), (i, l), (j, k), (j, l), (k, l)\}$. The neighbours of the pair (i, j) are the pairs $\{(i, k), (j, k), (i, l), (j, l)\}$, but not the pair (k, l) as (i, j) and (k, l) have no individuals in common. Then, $N_{(i,j)} = \{(i, k), (j, k), (i, l), (j, l)\}$.¹⁴

Then, under Assumption 2.2, we can therefore have

$$\mathbb{P}(w_{ij}|W_{m,-ij}) = \mathbb{P}(w_{ij}|W_{N_{ij}}),$$

where $N_{ij} = \{ik, jk\}_{k \neq i, j}$.

¹² Most of the results presented are rarely treated in econometrics textbooks, but are standard in statistical mechanics and spatial statistics. We refer the interested reader to Cressie (2015) and Koller and Friedman (2009) for additional reading.

¹³ The position of the pairs in $\mathcal{X}$ is formally described in the next section.

¹⁴ This is also the definition used by Graham (2016).

Figure 1. A pair's neighbours. [Colour figure can be viewed at wileyonlinelibrary.com]

A random field, joint with the collection of neighbourhoods, is called a (local) Markov random field.

Now that we have laid down the basic definitions, we can go back to the intractability of the denominator in (2.1). This intractability makes use of the MLE approach – based on the maximization of $\ln \mathbb{P}(W_m | \theta)$ with respect to θ – very difficult.

In this paper, we follow the approach of Besag (1975) and Strauss and Ikeda (1990), and we present an estimator based on the conditional probabilities (i.e. $\mathbb{P}(w_s | W_{m,-s}; \theta)$), which would be more tractable.

Indeed, Hammersley and Clifford (1971) and Besag (1974) give necessary and sufficient conditions on (2.1) under which all the conditional distributions are (i) valid, in the sense of being compatible with the joint distribution $\mathbb{P}(W_m | \theta)$ and (ii) tractable, in the sense that they do not depend on the intractable denominator presents in (2.1).

The following result is a direct application of Hammersley and Clifford (1971) and Besag (1974) recast into our framework.

PROPOSITION 3.1. *Suppose that Assumptions 2.1 and 2.2 hold. Then, the potential function can be written as*

$$Q(W_m; \theta) = \sum_s \chi_s(\theta) w_s + \sum_s \sum_{r < s: r \in N_s} \chi_{r,s}(\theta) w_r w_s, \quad (3.1)$$

where $r < s$ is for an arbitrary ordering of the pairs, and the conditional probabilities are such that

$$\mathbb{P}(w_s | W_{m,-s}; \theta) = \frac{\exp\{w_s h_s(W_{m,-s}, \theta)\}}{1 + \exp\{h_s(W_{m,-s}, \theta)\}}. \quad (3.2)$$

See also Lee et al. (2001) for a more recent treatment. Example 3.1 illustrates.

EXAMPLE 3.1. (EXAMPLE 2.1 CONTINUED) We have

$$h_s(W_{m,-s}; \theta) = -\beta |x_{s_1} - x_{s_2}| + \sum_{r \in N_s} \alpha w_r.$$

The potential function is therefore

$$Q(W_m; \theta) = \sum_s -w_s \beta |x_{s_1} - x_{s_2}| + \sum_s \sum_{r < s: r \in N_s} \alpha w_s w_r.$$

Proposition 3.1 gives us the true conditional probabilities, which do not depend on the intractable denominator.

We can therefore define the following pseudo-likelihood (PL) function, in the spirit of Besag (1975) and Strauss and Ikeda (1990):

$$PL = \sum_s (w_s h_s(W_{m,-s}; \theta) - \ln(1 + \exp\{h_s(W_{m,-s}; \theta)\})), \quad (3.3)$$

which is not the true log-likelihood function $\ln \mathbb{P}(W_m | \theta)$, unless we have full independence over the random field.

Because the conditional probabilities are well specified, we can expect the maximum of the PL to be a consistent estimator of the true parameter θ_0 given that it is unique. Besag (1975) provides a heuristic argument in that direction. However, a formal proof of the general case is far from trivial as it requires the use of limits theorems that apply to dependent random fields.

In what follows, we use recent results from Jenish and Prucha (2009), specifically central limit theorems and uniform laws of large numbers, in order to show under which conditions the consistency and asymptotic normality of the estimator are obtained by maximizing the PL in (3.3). Following Jenish and Prucha (2009), we assume an increasing domain asymptotic framework.

4. LIMITED DEPENDENCE THEOREMS

In this section, we present two theorems for dependent observations. We show that under ϕ -mixing, the true value $\theta_0 \in \Theta$ can be consistently estimated using the simple PL in (3.3). These theorems are useful because, as we show in Section 4, there exist simple conditions on h_{ij} that imply ϕ -mixing.¹⁵

To clarify the exposition, we use the following simplifying notation:

$$q_{s,m}(w_{s,m} | W_{m,-s}, \theta) = w_{s,m} h_s(W_{m,-s}; \theta) - \ln(1 + \exp\{h_s(W_{m,-s}; \theta)\}).$$

We also use $q_{s,m}(\theta) = q_{s,m}(w_{s,m} | W_{m,-s}, \theta)$ when there is no ambiguity. Maximizing the PL is equivalent to maximizing

$$\mathcal{L}_m(\theta) = \kappa(m)^{-1} \sum_{s \in S_m} q_{s,m}(w_{s,m} | W_{m,-s}, \theta), \quad (4.1)$$

where $\kappa(m)$ is a scaling parameter that determines the rate of convergence of our estimator. In general, $\kappa(m)$ will be related to the structure of the network, which depends on the specification of the payoff functions h_{ij} . Note that $\kappa(m)$ should be carefully chosen to avoid the case that $\lim_{m \rightarrow \infty} E[\mathcal{L}_m(\theta)]$ is infinite (i.e. $+/ - \infty$) or a trivial function of θ , e.g. $\lim_{m \rightarrow \infty} E[\mathcal{L}_m(\theta)] = 0$.

¹⁵ Our results can easily be adapted to other mixing definitions, such as α -mixing, as has been done by Leung (2014).

Otherwise, our identification condition (see Assumption 4.2) will be trivially violated.¹⁶ In Section 5, we provide conditions on h_{ij} such that $\kappa(m) = O(\sqrt{m})$.

We define the position of a pair s in $\mathcal{X}$ as the average point between the positions of s_1 and s_2 in $\mathcal{X}$, i.e. $x_s \in \mathcal{X}$ such that $x_s = (x_{s_1} + x_{s_2})/2$.¹⁷ Using this definition, the distance between two pairs r and s is equal to

$$d(s, r) = d(x_r, x_s) = d\left(\frac{x_{s_1} + x_{s_2}}{2}, \frac{x_{r_1} + x_{r_2}}{2}\right).$$

We also define $\mathcal{X}_{S_m} \equiv \{x_s : s \in S_m\}$ as the set of pair positions in $\mathcal{X}$.

We now turn to the dependence structure of the random field $\{q_{s,m}(\theta); x_s \in \mathcal{X}_{S_m}, m \in \mathbb{N}\}$, defined on the probability space $(\Omega_q, \mathcal{F}_q, \mathbb{P}_q)$, where $\Omega_q = \mathbb{R}$, with associated σ -algebra $\mathcal{F}_q$ and probability measure $\mathbb{P}_q$. For any two events $A \in \mathcal{A}$ and $B \in \mathcal{B}$, where $\mathcal{A}, \mathcal{B}$ are sub- σ -algebras of $\mathcal{F}$, the ϕ -mixing coefficient is given by

$$\phi(\mathcal{A}, \mathcal{B}) = \sup\{|\mathbb{P}(A|B) - \mathbb{P}(A)|, A \in \mathcal{A}, B \in \mathcal{B}, \mathbb{P}(B) > 0\}.$$

As discussed in Section 3, this is analogous to standard time-series models. In a stochastic model, the estimation is consistent if $\lim_{r \rightarrow \infty} \sup_t \phi(\mathcal{F}_{-\infty}^t, \mathcal{F}_{t+r}^\infty) = 0$, where $\mathcal{F}_{t_1}^{t_2}$ is the σ -algebra for the realizations from time t_1 to time t_2 .¹⁸ We want to apply the same basic approach when the dependence between A and B goes through $\mathcal{X}$. Because the dependence in $\mathcal{X}$ is more complex than time dependence, the asymptotic convergence of the ϕ -mixing coefficient is not sufficient. In order to show the consistency and asymptotic normality of $\hat{\theta} = \arg \max_{\theta} \mathcal{L}_m(\theta)$, we use the law of large numbers and the central limit theorem for dependent observations on random fields developed by Jenish and Prucha (2009) (Theorems 1, 2 and 3). We introduce the following definition.

DEFINITION 4.1. For $U \subset \mathcal{X}_{S_m}$ and $V \subset \mathcal{X}_{S_m}$, let $\mathcal{A}_m^U \equiv \sigma(q_{s,m} : x_s \in U)$ and $\mathcal{B}_m^V \equiv \sigma(q_{s,m} : x_s \in V)$, the corresponding σ -algebra. Also, let $|U|$ and $|V|$ denote the numbers of pairs located in U and V . We define the ϕ -mixing for the random field $\{q_{s,m}(\theta); s \in S_m, m \in \mathbb{N}\}$ with the function

$$\bar{\phi}_{k,l}(d) = \sup_m (\phi(\mathcal{A}_m^U, \mathcal{B}_m^V), |U| \leq k, |V| \leq l, d(U, V) \geq d),$$

where $d(U, V) = \inf\{d(s, r) : x_s \in U \text{ and } x_r \in V\}$.

We show that a sufficient condition for the consistency and the asymptotic normality of $\hat{\theta} = \arg \max_{\theta} \mathcal{L}_m(\theta)$ is the following.

ASSUMPTION 4.1. (ϕ -MIXING) (a) $\bar{\phi}_{1,1}^{1/2}(d) = O(d^{-T+1-\epsilon})$; (b) $\bar{\phi}_{k,l}(d) = O(d^{-T+1-\epsilon})$ for $k + l \leq 4$; (c) $\bar{\phi}_{1,\infty}(d) = O(d^{-T-\epsilon})$ for some $\epsilon > 0$.

Recall that $T \geq 1$ is the dimension of $\mathcal{X}$. In words, Assumption 4.1 implies that $\bar{\phi}_{k,l}(d)$ has to converge fast enough to 0. In Section 5, we provide sufficient conditions on the primitive equations of the model under which Assumption 4.1 holds. For now, we show that the estimation

¹⁶ We thank an anonymous referee for pointing out this important issue to us.

¹⁷ This is done without loss of generality. The method is robust to other definitions of a pair's position in $\mathcal{X}$, as long as x_s is located in a given neighbourhood of x_{s_1} and x_{s_2} .

¹⁸ For example, see White (2000).

technique is valid as long as ϕ -mixing is respected. The first theorem concerns the consistency of $\hat{\theta} = \arg \max_{\theta \in \Theta} \mathcal{L}_m(\theta)$. First, we need some additional assumptions.

ASSUMPTION 4.2. (IDENTIFICATION) *There exists a unique $\theta_0 \in \text{int } \Theta$ maximizing $\lim_{m \rightarrow \infty} E[\mathcal{L}_m(\theta)]$.*

This identification condition is discussed in detail in Pötscher and Prucha (1991). It will depend greatly on the specification of h_i^j whether or not this assumption holds; researchers should make sure that it holds in practice. We further discuss this assumption in Section 6.

The next assumption describes the asymptotic behaviour of the pairs in $\mathcal{X}$.

ASSUMPTION 4.3. (INCREASING DOMAIN) *For all $s_1, s_2 \in N$, $d(s_1, s_2) \geq d_0$ for some $d_0 > 0$ (wlog $d_0 = 1$)*

This implies that there exists a minimal distance (in $\mathcal{X}$) between pairs of individuals (and, therefore, between individuals). In essence, Assumption 4.3 ensures that the distance goes to infinity as the number of individuals goes to infinity. Given the existence of a minimal distance d_0 , the subspace of $\mathcal{X}$ that contains all the pairs of individuals (i.e. $\mathcal{X}_{S_m}$) has to expand as the number of individuals (and therefore pairs) increases. It is worth noting that we require only one observable characteristic for which there exists a minimum distance. For example, one could be interested in friendships between residential neighbours. In that case, there must exist a minimum distance between any neighbouring houses. Accordingly, if the number of households increases, it implies that the residential area also increases. We return to this assumption in the context of our empirical application in Section 7.

Finally, we also need the following standard moment conditions for the payoff functions.

ASSUMPTION 4.4. (TECHNICAL CONDITIONS I) *(a) $\sup_m \sup_s E[\sup_{\theta \in \Theta} |q_{s,m}(\theta)|^{(1+\eta)}] < \infty$ for some $\eta > 0$; (b) $\sup_m \sup_s E[\sup_{\theta \in \Theta} |(\partial q_{s,m}(\theta)/\partial \theta)|] < \infty$.*

Given the previous assumptions, we have the following.

THEOREM 4.1. (CONSISTENCY) *Suppose that Assumptions 4.2, 4.3 and 4.4 hold, and that Assumption 4.1(b) is respected for $k = l = 1$. Then, the estimator $\hat{\theta} = \arg \max_{\theta \in \Theta} \mathcal{L}_m(\theta)$ converges to θ_0 as $m \rightarrow \infty$.*

We still need to derive the asymptotic distribution of $\hat{\theta}$. We define the following matrices:

$$D_0(\theta_0) = \lim_{m \rightarrow \infty} E \left[\frac{\partial^2 \mathcal{L}_m(\theta_0)}{\partial \theta \partial \theta'} \right]$$

$$B_0(\theta_0) = \lim_{m \rightarrow \infty} \kappa(m) E \left[\frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} \left(\frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} \right)' \right].$$

Because the asymptotic normality of the estimator requires more structure than the one needed for consistency, we need Assumption 4.1, as well as the following additional technical conditions.

ASSUMPTION 4.5. (TECHNICAL CONDITIONS II) *(a) $B_0(\theta_0) > 0$; (b) $D_0(\theta_0)$ is invertible; (c) $\sup_m \sup_s E[\sup_{\theta \in \Theta} \|D_{s,m}(\theta)\|^{1+\eta}] < \infty$ for some $\eta > 0$; (d) $\sup_m \sup_s E[\sup_{\theta \in \Theta} \|(\partial D_{s,m}(\theta)/\partial \theta)\|] < \infty$; (e) $\sup_m \sup_s E[\sup_{\theta \in \Theta} |(\partial q_{s,m}(\theta)/\partial \theta)|^2] < \infty$, where $D_{m,s}(\theta) = (\partial^2 q_{s,m}(\theta)/\partial \theta \partial \theta')$.*

These assumptions are standard and are sufficient to show the asymptotic normality of our estimator.¹⁹

THEOREM 4.2. (ASYMPTOTIC NORMALITY) *Let $m \rightarrow \infty$. Under Assumptions 4.1, 4.2, 4.3, 4.4 and 4.5, the estimator $\hat{\theta} = \arg \max_{\theta \in \Theta} \mathcal{L}_m(\theta)$ is normally distributed with its variance–covariance matrix given by $D_0^{-1} B_0 D_0^{-1} / \kappa(m)$.*

The variance–covariance matrix in our setting is analogous to the heteroscedasticity and autocorrelation consistent (HAC) variance–covariance matrix in a time-series setting. The estimation of these variances is not straightforward. The estimation of $D_0(\theta_0)$ follows from Theorem 4.1, as $D_0(\theta)$ has the same dependence structure as $\lim_{m \rightarrow \infty} E[\mathcal{L}_m(\theta)]$. A consistent estimator then is $D_m(\hat{\theta}) = \kappa(m)^{-1} \sum_{s=1}^m D_{s,m}(\hat{\theta})$.

Defining a consistent estimator for $B_0(\theta_0)$ is more challenging. We suggest two approaches to estimate $B_0(\theta_0)$. The first is based on a generalization of standard HAC estimators (Conley, 1999). Although valid, the estimator proposed by Conley (1999) can be very computationally intensive when the number of dimensions of $\mathcal{X}$ increases (i.e. $T \geq 4$). Bester et al. (2011) propose an alternative approach using the well-known variance cluster (VC) estimator (formally described in the Appendix). Although the estimator is not consistent under weak dependence, they show that the estimator converges to a well-defined random variable and that the standard t -test is still valid (provided it is rescaled). In other words, under mixing conditions, inference using the VC estimator is valid, even if the estimator itself is not consistent. This estimator has the advantage of requiring little computational time, and it is simple to implement.

In this section, we have shown that with ϕ -mixing and some technical conditions, $\theta_0 \in \Theta$ can be recovered using (3.3). In the next section, we show that an asymptotic version of the homophily principle is a sufficient condition for ϕ -mixing, as defined in Assumption 4.1, as well as for $\kappa(m) = O(\sqrt{m})$.

5. THE ROLE OF HOMOPHILY

We now turn to economic models of network formation. Recall that Assumption 2.2 imposes a specific dependence structure. Under those restrictions, we show that a weak version of the homophily principle is sufficient to achieve ϕ -mixing.

Homophily is a prominent feature of social networks. It characterizes the empirical fact that similar individuals have a higher probability of being linked.²⁰ We assume the following weak version of homophily.

ASSUMPTION 5.1. (ASYMPTOTIC HOMOPHILY) *We assume that, for all $ij \in S_m$:*

$$d(i, j)^{-\lambda} \sup_m \sup_{W_{m,-ij}} h_{ij}(W_{m,-ij}; \theta) \rightarrow -\infty \text{ as } d(i, j) \rightarrow \infty$$

for some $\lambda \geq 1$ and all $\theta \in \Theta$.

¹⁹ Formally, the proof of Theorem 4.2 derives the limit distribution for $\sqrt{\kappa(m)}(\hat{\theta} - \theta_0)$. We report the asymptotic distribution of $\hat{\theta}$ for presentation purposes.

²⁰ Many definitions of homophily exist in the economic literature; see, e.g. Currarini et al. (2009, 2010) and Bramoullé et al. (2012). In particular, some papers explicitly define homophily using a distance function on the space of individual characteristics; see, e.g. Boucher (2015), Johnson and Gilles (2003), Marmaros and Sacerdote (2006), and Iijima and Kamada (2010).

Assumption 5.1 then requires that the probability of a link between two infinitely distant individuals is zero, and that the rate at which the value of the relationship decreases with the distance is sufficiently fast. Note that Assumption 5.1 only requires that homophily holds asymptotically, hence allowing for a wide range of locally non-homophilic preferences.

We show that Assumption 5.1 has powerful implications.

PROPOSITION 5.1. *If Assumptions 2.1, 2.2, 4.3 and 5.1 hold, then Assumptions 4.1 and 4.4(a) also hold. Moreover, a necessary condition for the identification condition (Assumption 4.2) to hold under Assumption 5.1 is $\kappa(m) = O(\sqrt{m})$.*

Assumptions 2.1, 2.2 and 5.1 are assumptions on the preferences (i.e. h_{ij}). Assumption 4.3 is the increasing domain assumption. Those assumptions imply that Assumptions 4.1 and 4.4(a), which are high-level assumptions on the random-field $\{q_s(\theta)\}$, also hold. The main benefit of Proposition 5.1 is therefore that the conditions for ϕ -mixing can be directly verified using the primitive equation of the structural model, as shown in Example 5.1.

EXAMPLE 5.1. (EXAMPLE 2.1 CONTINUED) Recall that

$$h_{ij}(W_{m,-ij}; \theta) = \alpha n_i(W_{m,-ij}) + \alpha n_j(W_{m,-ij}) - \beta |x_i - x_j|$$

and that $x_i, x_j \in \mathcal{X} = \mathbb{R}$ so $d(i, j) = |x_i - x_j|$. Then, Assumption 5.1 holds as long as $\beta > 0$.

Note also that Proposition 5.1 provides the order of the convergence rate, which is a necessary condition for identification. In the next section, we discuss the identification of the model in detail.

6. IDENTIFICATION

We now discuss further the implications of the identification condition stated in Assumption 4.2. In order to discuss the intuition, as well as the implications of the assumptions, it helps to concentrate on the special case where $h_{ij}(\theta)$ is linear.

6.1. Linearity in θ

We now show that, for the special case where $h_s(W_{m,-s}, \theta)$ is linear in θ , i.e. $h_s(W_{m,-s}, \theta) = \Gamma'_{s,m} \theta$ for some vector $\Gamma_{s,m}$, given the adequate choice of $\kappa(m)$, Assumption 4.2 holds whenever $\mathbb{E}[\Gamma_{s,m} \Gamma'_{s,m}]$ is finite and non-singular. Formally, we have the following.

PROPOSITION 6.1. *If $h_s(\cdot)$ is linear in θ , i.e. $h_s(W_{m,-s}; \theta) = \Gamma'_{s,m} \theta$ such that $\liminf_{m \rightarrow \infty} |\lambda_m| > 0$ where λ_m is the smallest eigenvalue of $\mathbb{E}[\Gamma_{s,m} \Gamma'_{s,m}]$, and $\limsup_{m \rightarrow \infty} \mathbb{E}[\Gamma_{s,m} \Gamma'_{s,m}] < \infty$, then the limit of the PL, $\lim_{m \rightarrow \infty} E[\mathcal{L}_m(\theta)]$, is uniquely maximized at the true value θ_0 .*

Proposition 6.1 can be seen as a corollary to Proposition 3.1, in the sense that limit of the PL is uniquely maximized at the true value θ_0 mainly because the conditional distributions are well specified. It also provides more classical conditions on the payoff function, as opposed to the generic formulation in Assumption 4.2.²¹

²¹ For instance, similar conditions are presented by Qu and Lee (2015).

However, Proposition 6.1 still requires somewhat high-level conditions for identification as $\Gamma_{s,m}$ is a function of $W_{m,-s}$. In general, finding sufficient primitive conditions for identification is very hard, and this is left for future research.

In what follows, we discuss the implication of the identification condition for some special cases. In particular, in the next section, we discuss the link between the identification and the sparsity of the network, as well as the role of $\kappa(m)$ for the identification of the model.

6.2. Sparsity

As discussed by de Paula et al. (2017b), the identification condition is closely linked to the sparsity of the network. We say that a network is sparse if it is not dense, where dense means that the number of links is $O(m)$.

Recall that Assumptions 2.2 and 4.3 rule out the possibility that individuals value an infinite number of links. Together with Assumption 5.1, it implies that the network is sparse.²² However, we argue that we can still expect the network to be sparse in some cases, without necessarily assuming that $\chi_{ij,ik} = 0$ whenever $\max\{d(i, j), d(i, k)\} > \bar{d}$. Example 6.1 illustrates.

EXAMPLE 6.1. (EXAMPLE 2.1 MODIFIED) Let i and j be the two farthest individuals in the population. We have

$$h_{ij}(W_{m,-ij}, \theta) = -\beta d(i, j) + \alpha \sum_{k \neq i, j} w_{ik} + w_{jk}.$$

Note that here we assume that, $\alpha_{ik} = \alpha_{jk} = \alpha$ for all i, j, k , irrespective of the distances between them.

Let $d(i, j) = |x_i - x_j|$, since i and j are the two farthest individuals, we know, under Assumption 4.3, that there may be at most $d(i, j) + 1$ individuals in the population,²³ so:

$$h_{ij} \leq -\beta d(i, j) + 2\alpha(d(i, j) + 1) = (2\alpha - \beta)d(i, j) + 2\alpha.$$

This upper bound goes to $-\infty$ if $\beta > 2\alpha$, so the two farthest individuals (i and j) will not be linked with probability approaching 1 as $m \rightarrow \infty$.

This is confirmed by numerical simulations. Let $\alpha = 1$ and consider the following two specifications:

$$\text{Specification 1, } h_{ij} = n_i + n_j - 3|x_i - x_j|;$$

$$\text{Specification 2, } h_{ij} = n_i + n_j - \exp\{|x_i - x_j|\}.$$

We simulate 1,000 draws from the joint distribution (2.1) using a Gibbs–Sampling algorithm. Figure 2 shows $E[n_i + n_j]^2$, for each specification, as n grows. Specifically, for each n , $E[n_i + n_j]^2$ is computed as the average (over the 1,000 networks simulated), of the average (over the $n(n-1)/2$ pairs) value of $(n_i + n_j)^2$. Overall, the result is independent of the size of the population n , irrespective of the specification.

²² This can be seen as a corollary to Proposition 5.1.

²³ Recall here that individuals are located on the line. A similar argument applies for $T \geq 1$.

Figure 2. $E[n_i + n_j]^2$ as a function of the number of individuals for each specification. [Colour figure can be viewed at wileyonlinelibrary.com]

Example 6.1 shows that sparsity can be achieved even without arbitrarily bounding the number of links, although the conditions are stronger than those of Assumption 5.1. In general, we expect the sufficient conditions on $h_{ij}(W_{m,-ij}, \theta)$ to be similar to Assumption 5.1, but where λ is bounded below by a function of T .

We now discuss the intuition for $\kappa(m) = O(\sqrt{m})$ under Assumption 5.1. Example 6.2 illustrates.

EXAMPLE 6.2. (EXAMPLE 2.1 MODIFIED) Consider a modification of Example 2.1, where

$$h_{ij}(W_{m,-ij}, \theta) = -\beta \mathbf{1}\{|x_i - x_j| = 1\} - \infty \mathbf{1}\{|x_i - x_j| > 1\} + \sum_{k \neq i, j} \alpha_{ik} w_{ik} + \alpha_{jk} w_{jk}.$$

In this extreme example, individuals located at a distance greater than 1 create a link with a probability equal to 0. In other words, individuals can only create links with their direct (left and right) neighbours. Because, for a population of $m = n(n-1)/2$ pairs, there are at most $n-1$ pairs located at a distance less than 1 (under Assumption 4.3), the PL (4.1) can be written as

$$\mathcal{L}_m(\theta) = \kappa(m)^{-1} \left[\sum_{s: |s_1 - s_2| = 1} q_{s,m}(w_{s,m} | W_{m,-ij}, \theta) \right].$$

Note that we need only to sum over pairs s such that $|s_1 - s_2| = 1$ as $q_{s,m}(\theta) = 0$ for all the other pairs, with probability 1 under Assumption 4.4(a). Because there are only $n-1$ pairs s such that $|s_1 - s_2| = 1$, we immediately obtain $\kappa(m) = n-1$.

Proposition 5.1 shows that the intuition in Example 6.2 extends more generally (i.e. whenever Assumption 5.1 holds). We now turn to our empirical application.

7. FRIENDSHIP NETWORKS

We are interested in the determinants of friendship formation among high-school students; see, e.g. Boucher (2015), Currarini et al. (2009, 2010) and Mele (2017). We use the Add Health database, which provides information on the friendship networks of US high-school students. We concentrate on the ‘saturated sample’, which provides information on 3,449 teenagers from 16 schools. In the database, we know if student i nominated j as a friend, and vice versa. Because we focus on undirected networks, we assume that i and j are friends either if i nominated j , or if j nominated i .

We consider two sets of variables. The first set ($x_{ij} \in \mathcal{X} \subset \mathbb{R}^2$) represents the (normalized) location of teenagers’ homes.²⁴ As discussed in Section 4, the asymptotics are well defined on this variable as there must exist a minimal distance between any two individuals’ homes. As the number of individuals (and therefore pairs) increases, the size of the residential neighbourhoods must also increase. The second set of variables (z_{ij}) contains other relevant characteristics of the pair (i, j) . We use the following specification:

$$z_{ij} = [Gender_{ij}, White_{ij}, Black_{ij}, Hisp_{ij}, Work_{ij}, Grade_{ij}].$$

Here, $Gender_{ij}$ is a binary variable that equals 1 if i and j are of the same gender, $White_{ij}$ is a binary variable that equals 1 if i and j are both white, $Black_{ij}$ is a binary variable that equals 1 if i and j are both black, $Hisp_{ij}$ is a binary variable that equals 1 if i and j are both Hispanic, $Grade_{ij}$ represents the absolute value of the difference between the grade levels of i and j and $Work_{ij}$ gives the sum of the weekly hours worked by i and j . The intuition is that the hours spent working could act as a substitute to (in-school) friendship relations.

We assume the following utility specification,

$$h_{ij}(W_m; \theta) = \theta_0 + [n_i(W_{m,-ij}) + n_j(W_{m,-ij})] \theta_1 + \theta_2 d(x_i, x_j) + z_{ij} \beta, \quad (7.1)$$

where $\beta = [\theta_3, \theta_4, \theta_5, \theta_6, \theta_7, \theta_8]'$. We can show that Assumption 2.2 holds because $h_{ij}(W_m; \theta)$ can also be written as

$$h_{ij}(W_m; \theta) = \theta_0 + z_{ij} \beta + \theta_2 d(x_i, x_j) + \sum_{k \neq i, j} [w_{ik} \theta_1 + w_{jk} \theta_1].$$

Then, we can define $\chi_{ij}(\theta) = \theta_0 + z_{ij} \beta + \theta_2 d(x_i, x_j)$, $\tilde{\chi}_{ij,ik} = \theta_1$ and $\tilde{\chi}_{ij,jk} = \theta_1$.

The model is therefore a special case of the general model developed in the previous sections. Tables 1 and 2 give a summary of the variables used for individuals, and for pairs of individuals.²⁵ We now proceed with the estimation.

Following Bester et al. (2011), we use cluster-robust standard errors. Marginal effects are reported in Table 3. The number of friends two individuals have in common has a positive

²⁴ This is constructed using normalized GPS coordinates. See the Appendix for the construction of the distance function as well as details on the scaling factor $\kappa(m)$.

²⁵ Note that the racial variables are not necessarily exclusive (e.g. a biracial individual could identify as both black and white). We have omitted the racial categories Asian, Native and Other.

Table 1. Summary statistics (pairs).

Variable	Mean	Std	Min	Max
w_{ij}	0.004	0.066	0	1
$n_i + n_j$	8.454	4.855	0	40
<i>Grade</i>	0.992	0.849	0	5
<i>Gender</i>	0.500	0.500	0	1
<i>White</i>	0.255	0.436	0	1
<i>Black</i>	0.049	0.216	0	1
<i>Hisp</i>	0.117	0.322	0	1
<i>Work</i>	17.201	17.370	0	170
$d(i, j)$	13.868	20.940	0	379.215
Number of pairs	1,624,408			

Table 2. Summary statistics (individuals).

Variable	Mean	Std	Min	Max
n_i	4.242	3.310	0	21
<i>Female</i>	0.495	0.500	0	1
<i>Grade level</i>	10.194	1.493	7	12
<i>White</i>	0.586	0.493	0	1
<i>Black</i>	0.159	0.366	0	1
<i>Hispanic</i>	0.202	0.401	0	1
Number of individuals	3,449			
Number of residential communities	16			
Number of high schools	16			

Table 3. Maximum likelihood estimates (logit), marginal effects ($\times 1,000$).

Variable	Estimate	Std error
$n_i + n_j$	0.341**	0.116
<i>Grade</i>	-5.061**	0.861
<i>Gender</i>	1.530**	0.289
<i>White</i>	6.399**	2.342
<i>Black</i>	7.012**	2.014
<i>Hisp</i>	1.890**	0.610
<i>Work</i>	-0.024	0.023
$d(i, j)$	-0.121*	0.059

Note: ** and * denote significance at the 1% and 5% levels, respectively. Standard errors are computed using the clustered-variance estimator by residential communities. See the Appendix for details.

impact on the strength of their friendship. We also observe homophily with respect to grade levels, gender and racial variables. Geographic distance has a significant negative impact, which is coherent with Assumption 5.1. Finally, we find a small negative effect of the number of hours worked by friends on the value of the link (this effect is not statistically significant).

8. CONCLUSION

In this paper, we have used the theory of Markov random fields in order to develop a simple estimator for an interesting class of ERGMs. We have shown that an asymptotic version of homophily is sufficient for ϕ -mixing, which implies that the estimation of the underlying preference parameters can be achieved using a simple pseudo-MLE. The methodology is appealing because it is simple and flexible. We have also provided an empirical application using data about the formation of US teenagers' friendship networks. We find that an individual's number of friends has a positive influence on link formation, and we also find evidence of homophily on gender, grade levels, race and geographic location.

ACKNOWLEDGEMENTS

This paper benefited from many highly interesting discussions. In particular, we want to thank Angelo Mele for his insightful comments on earlier versions of the paper. We also want to thank Yann Bramoullé, Diego Cerdeiro, Arun Chandrasekhar, Tim Conley, Sílvia Gonçalves, Bryan S. Graham, Paul Grieco, Marc Henry, Ilze Kalnina, David Karp, Matthew Lindquist, Yao Luo, Benoit Perron, Joris Pinkse and Eirini Tatsi.

In particular, we would like to thank Jaap Abbring, Aureo de Paula and two anonymous referees for their numerous and extremely insightful comments, and for suggestions, which have helped us to improve the paper in many important respects.

We would also like to thank the participants at many conferences and seminars, including the SCSE2012, the economics department at Penn State, the economics department at Université de Montreal, the CeMMAP-CIREQ-SciencesPo-X workshop on 'The Estimation of Complementarities in Matching and Social Networks' and the SOFI at Stockholm University. An earlier version of the paper circulated under the title 'My Friend Far Far Away: Asymptotic Properties of Pairwise Stable Networks', and it provided an interpretation of ERGMs in terms of pairwise stability.

Vincent Boucher thanks the economics faculty and the Cambridge-INET at the University of Cambridge for their support and for many interesting discussions, and is also grateful for financial support from FRQSC and SSHRC. A significant part of this work was done while Ismael Mourifié was visiting the economics department at Penn State. Many thanks for their support and hospitality.

This research uses data from Add Health, a programme directed by Kathleen Mullan Harris and designed by K. Harris, J. Richard Udry and Peter S. Bearman at the University of North Carolina at Chapel Hill, and funded by grant P01-HD31921 from the Eunice Kennedy Shriver National Institute of Child Health and Human Development, with cooperative funding from 23 other federal agencies and foundations. Special acknowledgment is due to Ronald R. Rindfuss and Barbara Entwisle for their assistance in the original design. Information on how to obtain the Add Health data files is available on the Add Health web site (www.cpc.unc.edu/addhealth). No direct support was received from grant P01-HD31921 for this analysis.

REFERENCES

- Badev, A. (2013). Discrete games in endogenous networks: theory and policy. Working Paper, Federal Reserve Board.

- Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B* 36, 192–236.
- Besag, J. (1975). Statistical analysis of non-lattice data. *Journal of the Royal Statistical Society, Series D* 24, 179–95.
- Bester, C. A., T. G. Conley and C. B. Hansen (2011). Inference with dependent data using cluster covariance estimators. *Journal of Econometrics* 165, 137–51.
- Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5, 387–424.
- Boucher, V. (2015). Structural homophily. *International Economic Review* 56, 235–64.
- Boucher, V. (2016). Conformism and self-selection in social networks. *Journal of Public Economics* 136, 30–44.
- Boucher, V. and B. Fortin (2016). Some challenges in the empirics of the effects of networks. In Y. Bramoullé, A. Galeotti and B. W. Rogers (Eds.), *The Oxford Handbook on the Economics of Networks*, Chapter 12, 277–302. Oxford: Oxford University Press.
- Bramoullé, Y., S. Currarini, M. O. Jackson, P. Pin and B. W. Rogers (2012). Homophily and long-run integration in social networks. *Journal of Economic Theory* 147, 1754–86.
- Chandrasekhar, A. G. (2016). Econometrics of network formation. In Y. Bramoullé, A. Galeotti and B. W. Rogers (Eds.), *The Oxford Handbook on the Economics of Networks*, Chapter 13, 303–57. Oxford: Oxford University Press.
- Chandrasekhar, A. G. and M. O. Jackson (2014). Tractable and consistent random graph models. Working Paper 20276, National Bureau of Economic Research.
- Chandrasekhar, A. G. and M. O. Jackson (2016). A network formation model based on subgraphs. Working Paper, Stanford University.
- Christakis, N. A., J. H. Fowler, G. W. Imbens and K. Kalyanaraman (2010). An empirical model for strategic network formation. Working Paper 16039, National Bureau of Economic Research.
- Conley, T. G. (1999). Gmm estimation with cross sectional dependence. *Journal of Econometrics* 92, 1–45.
- Cressie, N. (2015). *Statistics for Spatial Data*. New York, NY: Wiley.
- Currarini, S., M. O. Jackson and P. Pin (2009). An economic model of friendship: homophily, minorities, and segregation. *Econometrica* 77, 1003–45.
- Currarini, S., M. O. Jackson and P. Pin (2010). Identifying the roles of race-based choice and chance in high school friendship network formation. *Proceedings of the National Academy of Sciences* 107, 4857–61.
- de Paula, A. et al. (2017a). Econometrics of network models. Forthcoming in B. Honoré, A. Pakes, M. Piazzesi and L. Samuelson (Eds.), *Advances in Economics and Econometrics: Theory and Applications, Eleventh World Congress*. Cambridge: Cambridge University Press.
- de Paula, A., S. Richards-Shubik and E. Tamer (2017b). Identifying preferences in networks with bounded degree. Forthcoming in *Econometrica*.
- Echenique, F. and R. G. Fryer Jrs (2007). A measure of segregation based on social interactions. *Quarterly Journal of Economics* 122, 441–85.
- Fortin, B. and M. Yazbeck (2015). Peer effects, fast food consumption and adolescent weight gain. *Journal of Health Economics* 42, 125–38.
- Goldsmith-Pinkham, P. and G. W. Imbens (2013). Social networks and the identification of peer effects. *Journal of Business and Economic Statistics* 31, 253–64.
- Graham, B. S. (2016). Homophily and transitivity in dynamic network formation. Working Paper 22186, National Bureau of Economic Research.

- Graham, B. S. (2017). An empirical model of network formation: detecting homophily when agents are heterogeneous. Forthcoming in *Econometrica*.
- Hammersley, J. M. and P. Clifford (1971). Markov fields on finite graphs and lattices. Unpublished.
- Henry, M. and I. Mourifié (2013). Euclidean revealed preferences: testing the spatial voting model. *Journal of Applied Econometrics* 28, 650–66.
- Hsieh, C.-S. and L. F. Lee (2016). A social interactions model with endogenous friendship formation and selectivity. *Journal of Applied Econometrics* 31, 301–19.
- Iijima, R. and Y. Kamada (2010). Social distance and network structures. Working Paper, Harvard University.
- Janish, N. and I. R. Prucha (2009). Central limit theorems and uniform laws of large numbers for arrays of random fields. *Journal of Econometrics* 150, 86–98.
- Johnson, C. and R. P. Gilles (2003). Spatial social networks. In M. O. Jackson and B. Dutta (Eds.), *Networks and Groups: Models of Strategic Formation*, 51–77. Berlin: Springer.
- Koller, D. and N. Friedman (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge MA.
- Lee, J., M. S. Kaiser and N. Cressie (2001). Multiway dependence in exponential family conditional distributions. *Journal of Multivariate Analysis* 79, 171–90.
- Leung, M. (2014). A random-field approach to inference in large models of network formation. Working Paper, University of Southern California.
- Leung, M. (2015). A weak law for moments of pairwise-stable networks. Working Paper, University of Southern California.
- Marmaros, D. and B. Sacerdote (2006). How do friendships form? *Quarterly Journal of Economics* 121, 79–119.
- Mele, A. (2017). A structural model of segregation in social networks. Forthcoming in *Econometrica*.
- Mele, A. and L. Zhu (2016). Variational approximations for large strategic ERGMs. Working Paper, Carey Business School.
- Monderer, D. and L. S. Shapley (1996). Potential games. *Games and Economic Behavior* 14, 124–43.
- Mourifié, I. and A. Siow (2014). Cohabitation versus marriage: marriage matching with peer effects. Working Paper, University of Toronto.
- Patacchini, E. and Y. Zenou (2012). Ethnic networks and employment outcomes. *Regional Science and Urban Economics* 42, 938–49.
- Pötscher, B. M. and I. R. Prucha (1991). Basic structure of the asymptotic theory in dynamic nonlinear econometric models, part I: consistency and approximation concepts. *Econometric Reviews* 10, 125–216.
- Qu, X. and L.-f. Lee (2015). Estimating a spatial autoregressive model with an endogenous spatial weight matrix. *Journal of Econometrics* 184, 209–32.
- Sheng, S. (2014). A structural econometric analysis of network formation games. Working Paper, University of California at Los Angeles.
- Strauss, D. and M. Ikeda (1990). Pseudolikelihood estimation for social networks. *Journal of the American Statistical Association* 85, 204–12.
- Toomet, O., M. Van Der Leij and M. Rolfe (2013). Social networks and labor market inequality between ethnicities and races. *Network Science* 1, 321–52.
- Watts, A. (2007). Formation of segregated and integrated groups. *International Journal of Game Theory* 35, 505–19.
- White, H. (2000). *Asymptotic Theory for Econometricians*. New York, NY: Academic Press.

APPENDIX A: PROOFS

LEMMA A.1. *Assumption 4.1 implies that (a) $\sum_{d=1}^{\infty} d^{T-1} \bar{\phi}_{1,1}^{1/2}(d) < \infty$; (b) $\sum_{d=1}^{\infty} d^{T-1} \bar{\phi}_{k,l}(d) < \infty$, for $k+l \leq 4$; (c) $\bar{\phi}_{1,\infty}(d) = O(d^{-T-\epsilon})$ for some $\epsilon > 0$.*

Proof: The proof is immediate. Indeed, there exists $M > 0$ such that

$$\begin{aligned} \sum_{d=1}^{\infty} d^{T-1} \bar{\phi}_{1,1}^{1/2}(d) &\leq M \sum_{d=1}^{\infty} d^{T-1} d^{-T+1-\epsilon} \\ &\leq M \sum_{d=1}^{\infty} d^{-\epsilon} < \infty. \end{aligned}$$

□

LEMMA A.2. *Condition 4 is implied by Assumption 4.4.*

Proof: From the mean value theorem, we can write

$$q_{s,m}(\theta) = q_{s,m}(\theta') + \frac{\partial q_{s,m}(\tilde{\theta})}{\partial \theta} (\theta - \theta').$$

Thus,

$$\begin{aligned} |q_{s,m}(\theta) - q_{s,m}(\theta')| &\leq \left| \frac{\partial q_{s,m}(\tilde{\theta})}{\partial \theta} \right| \|\theta - \theta'\| \\ &\leq \sup_{\theta \in \Theta} \left| \frac{\partial q_{s,m}(\theta)}{\partial \theta} \right| \|\theta - \theta'\|. \end{aligned}$$

According to Proposition 1 of Jenish and Prucha (2009), $q_{s,m}(\theta)$ is L_0 stochastically equicontinuous on Θ if the following Cesàro sum is finite, i.e

$$\limsup_m \frac{1}{|S_m|} \sum_{s=1}^m E \left[\sup_{\theta \in \Theta} \left| \frac{\partial q_{s,m}(\theta)}{\partial \theta} \right| \right] < \infty.$$

However, under Assumption 4.4, each term of the Cesàro sum is finite, in the sense that $\sup_m \sup_{s \in S_m} E[\sup_{\theta \in \Theta} |\partial q_{s,m}(\theta)/\partial \theta|] < \infty$. This fact completes the proof. □

LEMMA A.3. *Under Assumption 4.2, $\sigma_m^2/\kappa(m) \rightarrow B_0(\theta_0)$.*

Proof:

$$\begin{aligned} \frac{1}{\kappa(m)} \sigma_m^2 &= \frac{1}{\kappa(m)} \text{Var} \left(\kappa(m) \frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} \right) \\ &= \kappa(m) E \left[\frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} \frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta'} \right] + \kappa(m) E \left[\frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} \right] E \left[\frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta'} \right] \\ &= \kappa(m) E \left[\frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} \left(\frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} \right)' \right]. \end{aligned}$$

where the last inequality holds since $E[\partial \mathcal{L}_m(\theta_0)/\partial \theta] = 0$, as θ_0 maximizes $E[\mathcal{L}_m(\theta)]$ (Assumption 4.2). Hence, $\sigma_m^2/\kappa(m) \rightarrow B_0(\theta_0)$. □

LEMMA A.4. *Under Assumption 4.5, $\sigma_m^{-1} Q_m \Rightarrow N(0, I)$*

Proof: It is sufficient to show that the conditions for Theorem 1 from Jenish and Prucha (2009) hold. Specifically, (a) $d(r, s) > d_0 > 0$ for any $r, s \in S_m$; (b) ϕ -mixing on random fields;

(c) $\sup_m \sup_{s \in S_m} E[\sup_{\theta \in \Theta} |\partial q_{s,m}(\theta)/\partial \theta|^2] < \infty$; (d) $\liminf_{m \rightarrow \infty} (\sigma_m^2/\kappa(m)) > 0$. Condition 1 is implied by Assumption 4.3. Condition 3 holds from Assumption 4.5, and Condition 4 is implied by Lemma A.3. $\square$

LEMMA A.5. $\partial^2 \mathcal{L}_m(\bar{\theta}_m)/\partial \theta \theta' \xrightarrow{P} D_0(\theta_0)$.

Proof: The proof is identical to the proof for the consistency of $\hat{\theta}$, replacing $q_{s,m}(\theta)$ by $D_{s,m}(\theta)$, and using Assumptions 4.5 instead of Assumptions 4.4. $\square$

CLAIM A.1. *Assumption 4.4(a) holds. Recall that $q_{s,m}(\theta) = q_{s,m}(w_{s,m} | W_{m,-s}; \theta)$. We have*

$$\begin{aligned}
 & \sup_m \sup_s E[\sup_{\theta} |q_{s,m}(\theta)|^{1+\eta}] \\
 &= \sup_m \sup_s E\left[E[\sup_{\theta} |q_{s,m}(\theta)|^{1+\eta} | W_{m,-s}]\right] \\
 &= \sup_m \sup_s E\left[P(w_{s,m} = 1 | W_{m,-s}; \theta) \sup_{\theta} |q_{s,m}(w_{s,m} = 1 | W_{m,-s})|^{1+\eta}\right. \\
 &\quad \left.+ P(w_{s,m} = 0 | W_{m,-s}; \theta) \sup_{\theta} |q_{s,m}(w_{s,m} = 0 | W_{m,-s})|^{1+\eta}\right] \\
 &= \sup_m \sup_s E\left[\frac{\exp\{h_s(W_{m,-s}; \theta)\}}{1 + \exp\{h_s(W_{m,-s}; \theta)\}} \sup_{\theta} |h_s(W_{m,-s}; \theta) - \ln[1 + \exp\{h_s(W_{m,-s}; \theta)\}]|^{1+\eta}\right. \\
 &\quad \left.+ \frac{1}{1 + \exp\{h_s(W_{m,-s}; \theta)\}} \sup_{\theta} \ln[1 + \exp\{h_s(W_{m,-s}; \theta)\}]^{1+\eta}\right], \\
 &= \sup_m \sup_s E\left[\frac{\exp\{h_s(W_{m,-s}; \theta)\}}{1 + \exp\{h_s(W_{m,-s}; \theta)\}} |h_s(W_{m,-s}; \theta^*) - \ln[1 + \exp\{h_s(W_{m,-s}; \theta^*)\}]|^{1+\eta}\right. \\
 &\quad \left.+ \frac{1}{1 + \exp\{h_s(W_{m,-s}; \theta)\}} \ln[1 + \exp\{h_s(W_{m,-s}; \tilde{\theta})\}]^{1+\eta}\right],
 \end{aligned}$$

where $\theta^* = \arg \max\{|h_s(W_{m,-s}; \theta) - \ln[1 + \exp\{h_s(W_{m,-s}; \theta)\}]\}$ and $\tilde{\theta} = \arg \max\{\ln[1 + \exp\{h_s(W_{m,-s}; \tilde{\theta})\}]\}$. Under Assumptions 2.2 and 4.3, we know that $h_s(W_{m,-s}; \theta)$ is uniformly bounded from above with respect to θ , $W_{m,-s}$, m and s ; then the terms $1/(1 + \exp\{h_s(W_{m,-s}; \theta)\}) \ln[1 + \exp\{h_s(W_{m,-s}; \tilde{\theta})\}]^{1+\eta}$ and $(\exp\{h_s(W_{m,-s}; \theta)\})/(1 + \exp\{h_s(W_{m,-s}; \theta)\}) |\ln[1 + \exp\{h_s(W_{m,-s}; \theta^*)\}]|^{1+\eta}$ are also bounded. The remaining term to study is

$$\begin{aligned}
 & \frac{\exp\{h_s(W_{m,-s}; \theta)\}}{1 + \exp\{h_s(W_{m,-s}; \theta)\}} |h_s(W_{m,-s}; \theta^*)|^{1+\eta} \\
 &= \frac{\exp\{h_s(W_{m,-s}; \theta)\}}{1 + \exp\{h_s(W_{m,-s}; \theta)\}} |h_s(W_{m,-s}; \theta)|^{1+\eta} \left| \frac{h_s(W_{m,-s}; \theta^*)}{h_s(W_{m,-s}; \theta)} \right|^{1+\eta} \\
 &\leq \frac{\exp\{h_s(W_{m,-s}; \theta)\}}{1 + \exp\{h_s(W_{m,-s}; \theta)\}} |h_s(W_{m,-s}; \theta)|^{1+\eta},
 \end{aligned}$$

where the last inequality holds as we have

$$\begin{aligned}
 \theta^* &= \arg \max \left\{ \left| \ln \frac{\exp\{h_s(W_{m,-s}; \theta)\}}{1 + \exp\{h_s(W_{m,-s}; \theta)\}} \right| \right\} \\
 &= \arg \min \frac{\exp\{h_s(W_{m,-s}; \theta)\}}{1 + \exp\{h_s(W_{m,-s}; \theta)\}} \\
 &= \arg \min h_s(W_{m,-s}; \theta).
 \end{aligned}$$

Finally, we have

$$\lim_{h_s(W_{m,-s};\theta) \rightarrow \infty} \frac{\exp\{h_s(W_{m,-s};\theta)\}}{1 + \exp\{h_s(W_{m,-s};\theta)\}} |h_s(W_{m,-s};\theta)|^{1+\eta} \rightarrow 0$$

uniformly over θ , $W_{m,-s}$, m , and s .

CLAIM A.2. $\kappa(m) = O(n)$.

Let us consider

$$\kappa(m)\mathcal{L}_m(\theta) = \sum_s^m \left(w_s h_s(W_{m,-s};\theta) - \ln[1 + \exp\{h_s(W_{m,-s};\theta)\}] \right).$$

For a fixed individual i , let us order the rest of the individuals $i_2, \dots, i_n$ in the lattice such that $d(i, i_2) \leq d(i, i_3) \leq \dots \leq d(i, i_n)$. Note that we can have multiple individuals at the same distance $\bar{d}$ of i . Therefore, we can partition $\{i_2, \dots, i_n\}$ as $\{i_2, \dots, i_n\} = \{B_1\} \cup \dots \cup \{B_{K_i}\}$, such that $\forall(l, l') \in B_k$ $d(i, \tilde{l}) < d(i, l) = d(i, l') \equiv d_k < d(i, \hat{l})$ for $\tilde{l} \in B_{k-1}$ and $\hat{l} \in B_{k+1}$.

In such a case, $\kappa(m)\mathcal{L}_m(\theta)$ can be rewritten as

$$\kappa(m)\mathcal{L}_m(\theta) = \frac{1}{2} \sum_i^n \sum_{j=i_2}^{i_n} \left(w_{ij} h_{ij}(W_{m,-ij};\theta) - \ln[1 + \exp\{h_{ij}(W_{m,-ij};\theta)\}] \right)$$

or equivalently as

$$\kappa(m)\mathcal{L}_m(\theta) = \frac{1}{2} \sum_i^n \sum_{k=1}^{K_i} \sum_{j \in B_k} \left(w_{ij} h_{ij}(W_{m,-ij};\theta) - \ln[1 + \exp\{h_{ij}(W_{m,-ij};\theta)\}] \right).$$

Consider the series:

$$\sum_{j=i_2}^{i_n} \left(E[w_{ij} h_{ij}(W_{m,-ij};\theta)] - \ln[1 + \exp\{h_{ij}(W_{m,-ij};\theta)\}] \right).$$

Because

$$0 \leq \ln[1 + \exp\{h_{ij}(W_{m,-ij};\theta)\}] \leq \exp\{h_{ij}(W_{m,-ij};\theta)\},$$

we have

$$\begin{aligned} \sum_{j=i_2}^{i_n} \ln[1 + \exp\{h_{ij}(W_{m,-ij};\theta)\}] &\leq \sum_{k=1}^{K_i} \sum_{j \in B_k} \exp\{\bar{h}_i(d_k)\} \\ &= \sum_{k=1}^{K_i} |B_k| \exp\{\bar{h}_i(d_k)\}, \end{aligned}$$

where $\bar{h}_i(d_k) = \sup_{\theta} \sup_{ij: j \in B_k} \sup_{W_m \in \mathbb{W}_m} h_{ij}(W_{m,-ij};\theta)$.

From Lemma A.1(iii) of Jenish and Prucha (2009), there are at most Cd^{T-1} (where $C > 0$ is a positive constant) individuals located at a distance $d_k \in [d, d+1]$ from i , so we have

$$\begin{aligned} \sum_{j=i_2}^{i_n} \ln[1 + \exp\{h_{ij}(W_{m,-ij};\theta)\}] &\leq C \sum_{d=1}^{K_i} d^{T-1} \exp\{\bar{h}_i(d+1)\} \\ &\leq C \sum_{d=1}^{\infty} d^{T-1} \exp\{\bar{h}_i(d+1)\}. \end{aligned}$$

Using the Cauchy criteria for convergent series, we can see that the latter summation converges if $\lim_{d \rightarrow \infty} [(T-1)(\ln d/d) + (\bar{h}_i(d+1)/d)] < 0$, which is true for $\lambda = 1$ in Assumption 5.1.

Consider now the series

$$\sum_{j=i_2}^{i_n} E[w_{ij} h_{ij}(W_{m,-ij}; \theta)] \leq \sum_{j=i_2}^{i_n} E[w_{ij} |h_{ij}(W_{m,-ij}; \theta)|].$$

We have

$$\begin{aligned} \sum_{j=i_2}^{i_n} E[w_{ij} |h_{ij}(W_{m,-ij}; \theta)|] &= \sum_{j=i_2}^{i_n} E[|h_{ij}(W_{m,-ij}; \theta)| E[w_{ij} | W_{m,-ij}]], \\ &\leq \sum_{j=i_2}^{i_n} |\bar{h}_{ij}| E[P(w_{ij} = 1 | W_{m,-ij})], \\ &\leq \sum_{j=i_2}^{i_n} |\bar{h}_{ij}| \sup_{W_m \in \mathbb{W}_m} P(w_{ij} = 1 | W_{m,-ij}), \\ &= \sum_{j=i_2}^{i_n} |\bar{h}_{ij}| \sup_{W_m \in \mathbb{W}_m} \frac{\exp\{h_{ij}(\cdot)\}}{1 + \exp\{h_{ij}(\cdot)\}}, \\ &\leq \sum_{j=i_2}^{i_n} |\bar{h}_{ij}| \frac{\exp\{\bar{h}_{ij}\}}{1 + \exp\{\bar{h}_{ij}\}}, \end{aligned}$$

where $\bar{h}_{ij} = \sup_{W_m \in \mathbb{W}_m} h_{ij}(W_{m,-ij}; \theta)$. The last inequality holds because $\exp\{x\}/(1 + \exp\{x\})$ is monotonic in x . We have

$$\begin{aligned} \sum_{j=i_2}^{i_n} |\bar{h}_{ij}| \frac{\exp\{\bar{h}_{ij}\}}{1 + \exp\{\bar{h}_{ij}\}} &= \sum_{k=1}^{K_i} \sum_{j \in B_k} |\bar{h}_{ij}| \frac{\exp\{\bar{h}_{ij}\}}{1 + \exp\{\bar{h}_{ij}\}} \\ &\leq \sum_{k=1}^{K_i} |B_k| |\tilde{h}_i(d_k)| \frac{\exp\{\tilde{h}_i(d_k)\}}{1 + \exp\{\tilde{h}_i(d_k)\}}, \end{aligned}$$

where

$$\tilde{h}_i(d_k) = \arg \max_{\{\bar{h}_{ij}: d(i,j)=d_k\}} |\bar{h}_{ij}| \frac{\exp\{\bar{h}_{ij}\}}{1 + \exp\{\bar{h}_{ij}\}}.$$

Note that $\tilde{h}_i(d_k)$ is finite, as $\bar{h}_{ij}$ is bounded from above and $|\bar{x}|(\exp\{x\}/(1 + \exp\{x\})) \rightarrow 0$ when $x \rightarrow \infty$. We can use similar derivations as earlier when invoking Lemma A.1(iii) from Jenish and Prucha (2009). Then, we obtain the following:

$$\sum_{j=i_2}^{i_n} E[w_{ij} |h_{ij}(W_{m,-ij}; \theta)|] \leq C \sum_{d=1}^{\infty} d^{T-1} \frac{\exp\{\tilde{h}_i(d)\}}{1 + \exp\{\tilde{h}_i(d)\}} |\tilde{h}_i(d)|.$$

Again, using the Cauchy criteria for convergent series, we can see that the latter summation converges if

$$\lim_{d \rightarrow \infty} \left[(T-1) \frac{\ln d}{d} + \frac{\ln |\tilde{h}_i(d)|}{d} + \frac{\tilde{h}_i(d)}{d} - \frac{\ln(1 + \tilde{h}_i(d))}{d} \right] < 0,$$

which is true for $\lambda = 1$ in Assumption 5.1.

So, we have shown that

$$\left| \sum_{j=i_2}^{i_n} E[q_{ij}(\theta)] \right| = \left| \sum_{j=i_2}^{i_n} E[w_{ij} h_{ij}(W_{m,-ij}; \theta) - \ln[1 + \exp\{h_{ij}(W_{m,-ij}; \theta)\}]] \right| < \infty.$$

As can be seen in the proof of Claim A.1, $E[q_{ij,m}(\theta)]$ is a weighting average of

$$q_{ij,m}(w_{ij,m} = 1 | W_{m,-ij}; \theta) = \ln \frac{\exp\{h_{ij}(W_{m,-ij}; \theta)\}}{1 + \exp\{h_{ij}(W_{m,-ij}; \theta)\}}$$

and

$$q_{ij,m}(w_{ij,m} = 0 | W_{m,-ij}; \theta) = -\ln[1 + \exp\{h_{ij}(W_{m,-ij}; \theta)\}],$$

where the weights are all probabilities (non-negative). Because $q_{ij,m}(w_{ij,m} = 1 | W_{m,-ij}; \theta)$ and $q_{ij,m}(w_{ij,m} = 0 | W_{m,-ij}; \theta)$ are all non-positive uniformly over θ , $W_{m,-ij}$, m and ij , then $E[q_{ij,m}(\theta)] \leq 0$.

Therefore, A_i is a non-positive finite constant (i.e. $-\infty < A_i \leq 0$, such that $E[Z_i] \equiv \sum_{j=i_2}^{i_n} E[q_{ij,m}(\theta)] = A_i$). Thus, $\kappa(m)\mathcal{L}_m(\theta)$ can be rewritten as a summation from 1 to n of Z_i . We have

$$E[\mathcal{L}_m(\theta)] = \frac{1}{\kappa(m)} \sum_i^n A_i.$$

Therefore, $\kappa(m) = O(n)$.

CLAIM A.3. Assumption 4.1 holds. The objective is to show that, as the distance (in X) between the pairs located in U and the pairs located in V goes to infinity, we find that $|P(A^U | B^V) - P(A^U)|$ goes to 0 (at the appropriate rate). The proof is rather lengthy and proceeds in a series of steps. In order to clarify the argument, we briefly explain each step.

- STEP 1. Because ϕ -mixing is defined as the distance between sets (i.e. $d(U, V)$) of pairs goes to infinity, we describe what happens to the distance between the individuals as $d(U, V)$ grows. This allows us to separate three cases for $s : x_s \in U$, and $r : x_r \in V$.
- STEP 2. (First case): $d(s, r) = O(d(s_1, s_2))$.
- STEP 3. (Second case): $d(s, r) = O(d(r_1, r_2))$.
- STEP 4. (Third case): $d(s, r) = O(\min\{d(s_1, r_1), d(s_1, r_2), d(s_2, r_1), d(s_2, r_2)\})$.

We also make the following remarks and notations.

- (a) Technically, A^U gives realizations of $\{q_{ij}\}_{ij}$ for individuals i and j , such that the pair (i, j) is located in the subspace U of X (i.e. $x_{ij} \in U \subset X$). To lighten the text, we often use ‘ i belongs to U ’, instead of ‘ i belongs to a pair located in U ’, when there is no confusion. Also, we often use U instead of $\{s\}_{s:x_s \in U}$.
- (b) The sets A^U and B^V are realizations over the random field $\{q_s(\theta)\}$, which is itself a function of the random field $\{w_s\}$. Formally, $q_s(\theta) = w_s h_s(W_{m,-s}; \theta) - \ln(1 + \exp\{h_s(W_{m,-s}; \theta)\})$. From Assumption 2.2, the function $h_s(W_{m,-s}; \theta)$ only depends on neighbourhood links, so $q_s(\theta) = f(w_s, W_{N(s)})$.
- (c) We expand the notion of neighbourhoods to set of pairs: $N(U) = \cup_{s:x_s \in U} N(s) \setminus \{s\}_{s:x_s \in U}$ is the set of neighbored pairs of pairs located in U (excluding pairs located in U).

STEP 1. Distance between sets of pairs, and distance between individuals.

By definition, $d(U, V) = d(s^*, r^*) = \inf\{d(s, r) | s : x_s \in U, r : x_r \in V\}$. First, note that if $U \cap V \neq \emptyset$, there is a pair located in U , which is also located in V . We have $d(U, V) = 0$. We can therefore abstract for this case, and concentrate on the case where $U \cap V = \emptyset$.

Figure A.1. Pairs do not share any individual. [Colour figure can be viewed at wileyonlinelibrary.com]

Figure A.2. Pairs share an individual. [Colour figure can be viewed at wileyonlinelibrary.com]

Figure A.3. Cases 2 and 3: $V \cap N(U) \neq \emptyset$. [Colour figure can be viewed at wileyonlinelibrary.com]

Second, note that any couple of pairs (r, s) such that $x_s \in U$ and $x_r \in V$ is either such that s and r do not share any individual (as in Figure A.1), or such that they share an individual (as in Figure A.2). Note that, in Figures A.1–A.3, neighbouring pairs of pair s are represented by solid lines. Consider the case where s and r do not share any individual, i.e. $s \cap r = \emptyset$, as in Figure A.1. As $d(s, r)$ grows, there are three (exhaustive, but non-mutually exclusive) possibilities: (a) $d(s_1, s_2)$ goes to infinity, formally $d(s, r) = O(d(s_1, s_2))$; (b) $d(r_1, r_2)$ goes to infinity, formally $d(s, r) = O(d(r_1, r_2))$; (c) $\min\{d(s_1, r_1), d(s_1, r_2), d(s_2, r_1), d(s_2, r_2)\}$ goes to infinity, formally $d(s, r) = O(\min\{d(s_1, r_1), d(s_1, r_2), d(s_2, r_1), d(s_2, r_2)\})$.

Figure A.4. Case 4(a) $N(U) \cap N(V) \neq \emptyset$, where r_1 belongs to a pair in U , as well as to a pair in V .
[Colour figure can be viewed at wileyonlinelibrary.com]

Figure A.5. Case 4(b) $N(U) \cap N(V) \neq \emptyset$, where r_1 belongs to a pair in U and r_2 belongs to a pair in V .
[Colour figure can be viewed at wileyonlinelibrary.com]

To see why the order is as above, consider the first case (the other cases are symmetric): $d(s, r) = O(d(s_1, s_2))$. Consider Figure A.1, where we let $d(s_2, r_2) = \epsilon_1 > 0$ and $d(r_1, r_2) = \epsilon_2 > 0$. This forces the distance between s_1 and s_2 to go to infinity as $d(s, r)$ goes to infinity. We have $d(r_1, s_2) \leq \epsilon_1 + \epsilon_2$. Moreover, we have $d(s_1, r_1) \leq d(s_1, s_2) + d(s_2, r_1)$, so

$$d(s, r) \leq \epsilon_1/2 + d(s_1, s_2)/2 + (\epsilon_1 + \epsilon_2)/2,$$

which implies $d(s, r) = O(d(s_1, s_2))$. This argument can be repeated for all possible cases, and also for the case where s and r share an individual.

Consider the case where s and r share one individual, i.e. $s \cap r \neq \emptyset$, as in Figure A.2. Similarly to the case where $s \cap r = \emptyset$, we have that as $d(s, r)$ grows, there are two exhaustive, but non-mutually exclusive possibilities: (a) $d(s, r) = O(d(s_1, s_2))$; (b) $d(s, r) = O(d(r_1, r_2))$. Note that here, in either case, we necessarily have that $d(s, r) = O(d(s_2, r_1))$.

This first step shows that we need to look at three cases when $d(s, r)$ goes to infinity: (a) $d(s, r) = O(d(s_1, s_2))$; (b) $d(s, r) = O(d(r_1, r_2))$; (c) $d(s, r) = O(\min\{d(s_1, r_1), d(s_1, r_2), d(s_2, r_1), d(s_2, r_2)\})$. Here, $s : x_s \in U$ and $r : x_r \in V$. We study these cases in Steps 2, 3 and 4.

STEP 2. Consider the case where $d(U, V) = O(d(s_1, s_2))$ for some $s : x_s \in U$. For $A^U \in \mathcal{A}_m^U$ and $B^V \in \mathcal{B}_m^V$, we have

$$|P(A^U | B^V) - P(A^U)| \leq \sum_{t=1}^k |P(A^{s_t} | B, A^{s_1}, \dots, A^{s_{t-1}}) - P(A^{s_t} | A^{s_1}, \dots, A^{s_{t-1}})|, \quad (\text{A.1})$$

from the chain rule, where $A^U = (A^{s_1}, \dots, A^{s_k})$.

Note that, given a fixed m , for any pair s , q_s is a discrete random variable because it has a countable range. Indeed, recall that q_s is a function of a countable finite number of binary random variable, i.e.

$$q_s = \ln \frac{\exp\{w_{s,m} h_s(W_{m,-s}; \theta)\}}{1 + \exp\{h_s(W_{m,-s}; \theta)\}}.$$

Let $\{a_1^s, \dots, a_{\tilde{m}}^s\}$ be the support of q_s where $\tilde{m}$ is the cardinality of this support. Note that because of the homophily assumption (Assumption 5.1) when the distance between the pairs $s = (s_1, s_2)$ converges to infinity (i.e. $d(s_1, s_2) \rightarrow \infty$), then the support is $\{-\infty, 0\}$.

There are three cases.

CASE 1. $A^s \in (-\infty, 0)$. Then, letting $d(s_1, s_2) \rightarrow \infty$, there exists $K > 0$ such that $P(q_s = A^s) = 0$ for all $d(s_1, s_2) > K$, as A^s falls outside the range of q_s . Therefore, we have $P(A^{s_t} | B, A^{s_1}, \dots, A^{s_{t-1}}) = P(A^{s_t} | A^{s_1}, \dots, A^{s_{t-1}}) = 0$.

CASE 2. $A^s = -\infty$. However, under Assumption 4.4(a), we know that $\sum_{k=1}^{\tilde{m}} |a_k^s| P(q_s = a_k^s) < \infty$. Then, $P(q_s = a_k^s) = 0$ whenever $a_k^s = -\infty$, otherwise q_s will not have a finite moment. Therefore, we have $P(A^{s_t} | B, A^{s_1}, \dots, A^{s_{t-1}}) = P(A^{s_t} | A^{s_1}, \dots, A^{s_{t-1}}) = 0$.

CASE 3. $A^s = 0$. This only happens as $d(s_1, s_2) = \infty$, where the range only has two elements, i.e. $-\infty$ and 0 . Because $P(-\infty) = 0$, we have $P(q_s = A^s) = 1$. Therefore, we have $P(A^{s_t} | B, A^{s_1}, \dots, A^{s_{t-1}}) = P(A^{s_t} | A^{s_1}, \dots, A^{s_{t-1}}) = 1$ when $d(s_1, s_2) = \infty$ while $P(A^{s_t} | B, A^{s_1}, \dots, A^{s_{t-1}}) = P(A^{s_t} | A^{s_1}, \dots, A^{s_{t-1}}) = 0$ when $d(s_1, s_2) < \infty$.

STEP 3. Consider, now, the case where $d(U, V) = O(d(r_1, r_2))$ for some $r : x_r \in V$. For $A^U \in \mathcal{A}_m^U$ and $B^V \in \mathcal{B}_m^V$, we need to consider

$$|P(A^U | B^V) - P(A^U)|,$$

for all A^U and B^V , such that $P(B^V) > 0$. Similarly to Step 2, we find that $P(q_r = B^r)$ has a range in $(-\infty, 0)$ if $d(r_1, r_2) < \infty$ so whenever there exists r such that $B^r \in (-\infty, 0)$, there exists $K > 0$ such that $P(B) = 0$ for all $d(r_1, r_2) > K$. In those cases, ϕ -mixing does not apply. Similarly, if $B^r = -\infty$, we also have $P(B) = 0$ from Assumption 4.4(a) (which must hold for all pairs in the random field). Finally, suppose that $B^r = 0$, we have that $P(B^r = 0) = 0$ for $d(r_1, r_2) < \infty$ and $d(r_1, r_2) < \infty$ for $P(B^r = 0) = 1$ for $d(r_1, r_2) = \infty$.

In all cases, whenever there exists $r : x_r \in V$ such that $d(U, V) = O(d(r_1, r_2))$, ϕ -mixing holds trivially, either because $P(B) = 0$, or because $P(B)$ is a Dirac distribution.

STEP 4. It remains to check the case where $d(U, V)$ is not $O(d(s_1, s_2))$, or $O(d(r_1, r_2))$, for any pairs s or r such that $s : x_s \in U$ and $r : x_r \in V$. That is, in (A.1), the elements of the sum for which $d(U, V)$ is not $O(d(s_1, s_2))$ (Step 2), and assuming that no pair $r : x_r \in V$ are such that $d(U, V) = O(d(r_1, r_2))$.

We define the following set:

$$C_{U,V} = \{s \in S_m | s_1 \in r : x_r \in U \text{ and } s_2 \in t : x_t \in V\}.$$

That is, the set of pairs (i, j) such that i is in a pair located in U , and such that j is in a pair located in V . We also let $D_{U,V} = 0$ if $W_{C_{U,V}} = 0$, and $D_{U,V} = 1$ otherwise. Then, $D_{U,V} = 0$ implies that there is no link between an individual in U and an individual in V .

From Step 1, and because $d(U, V)$ is not $O(d(s_1, s_2))$, or $O(d(r_1, r_2))$, for any pairs s or r such that $s : x_s \in U$ and $r : x_r \in V$, we necessarily have that $d(U, V) = O(d(i, j))$ for all $(i, j) \in C_{U,V}$.

We use the following result.

STEP 4.1. $P(A^U | B^V, D_{U,V} = 0) = P(A^U | D_{U,V} = 0)$. Equivalently, $P(A^U | B^V, W_{C_{U,V}} = 0) = P(A^U | W_{C_{U,V}} = 0)$. Indeed, typically, $W_{C_{U,V}}$ does contain information also contained in B^V . We

claim that it contains all the information contained in B^V . Note that $W_{C_{U,V}} = 0$ means that all pairs between an individual in a pair in U and an individual in a pair in V are not linked. Now, by construction of the random field $\{q_s\}$ from $\{w_s\}$ (i.e. $q_s = f(w_s, W_{N(s)})$), B^V depends only on $W_{V \cup N(V)}$, while A^U depends only on $W_{U \cup N(U)}$. Note that this heavily relies on Proposition 3.1. There are four cases.

Case 1. $U \cap V \neq \emptyset$. Because this implies that $d(U, V) = 0$, we can disregard this case (see Step 1). We can therefore interpret the three following cases as being conditional on $U \cap V = \emptyset$.

Case 2. $U \cap N(V) \neq \emptyset$. In words, this is the set of pairs in U that have an individual in a pair in V . In Figure A.3 for instance, this is pair s . Clearly, any pair in $U \cap N(V)$ has an individual in a pair in U and an individual in a pair in V . Therefore, $U \cap N(V) \subseteq C_{U,V}$.

Case 3. $V \cap N(U) \neq \emptyset$. Similarly, in Figure A.3, this is a pair in V that has an individual in U (i.e. (r_1, s_1)). We also have $V \cap N(U) \subseteq C_{U,V}$.

Case 4. $N(U) \cap N(V) \neq \emptyset$. We need to describe the set $N(U) \cap N(V)$. Any pair $r \in N(U)$ has one (and only one) individual in a pair in U (i.e. $r_1 \in (r_1, s_1) \in U$). Similarly, $r \in N(V)$ has one individual in a pair in V . There are two cases.

Case 4(a). $r_1 \in (r_1, s_2) \in V$. Then $(r_1, s_2) \in V \cap N(U)$ (i.e. U and V share an individual, here r_1). However, this implies that $(r_1, s_1) \in C_{U,V}$ and $(r_1, s_2) \in C_{U,V}$ so conditional on $W_{C_{U,V}} = 0$, the realization of w_r (which is in the intersection of $N(U)$ and $N(V)$) brings no additional information on either (r_1, s_1) or (r_1, s_2) . See Figure A.4.

Case 4(b). $r_2 \in (r_2, s_2) \in V$, where r is a link between an individual in a pair in U and an individual in a pair in V . Therefore, $r \in C_{U,V}$. See Figure A.5.

Therefore, in all cases, B^V provides no information on A^U , conditional on $W_{C_{U,V}} = 0$.

STEP 4. (CONTINUED) Consider $A^U \in \mathcal{A}_m^U$ and $B^V \in \mathcal{B}_m^V$. Let partition U in U_1 and $U_2 = U \setminus U_1$ such that U_1 contains the pairs s for which $d(U, V) = O(d(s_1, s_2))$ (provided U_1 is non-empty). We have

$$\begin{aligned} |P(A^U | B^V) - P(A^U)| &= |P(A^{U_1} | A^{U_2}, B^V)P(A^{U_2} | B^V) - P(A^{U_1} | A^{U_2})P(A^{U_2})| \\ &\leq |P(A^{U_1} | A^{U_2}, B^V)P(A^{U_2} | B^V) - P(A^{U_1} | A^{U_2})P(A^{U_2} | B^V)| \\ &\quad + |P(A^{U_1} | A^{U_2})P(A^{U_2}) - P(A^{U_1} | A^{U_2})P(A^{U_2} | B^V)| \\ &\leq |P(A^{U_1} | A^{U_2}, B^V) - P(A^{U_1} | A^{U_2})|P(A^{U_2} | B^V) \\ &\quad + |P(A^{U_2}) - P(A^{U_2} | B^V)|P(A^{U_1} | A^{U_2}). \end{aligned}$$

So the mixing coefficient $|P(A^U | B^V) - P(A^U)|$ is bounded by $|P(A^{U_1} | A^{U_2})P(A^{U_2}) - P(A^{U_1})|$ and $|P(A^{U_2}) - P(A^{U_2} | B^V)|$. The first quantity was treated in Step 1.

We now proceed to show that $|P(A^{U_2} | B^V) - P(A^{U_2})| \leq 2[P(D_{U,V} = 1 | B^V) + P(D_{U,V} = 1)]$,

$$\begin{aligned} P(A^{U_2} | B^V) - P(A^{U_2}) &= P(A^{U_2} | B^V, D_{U,V} = 1)P(D_{U,V} = 1 | B^V) \\ &\quad + P(A^{U_2} | B^V, D_{U,V} = 0)P(D_{U,V} = 0 | B^V) \\ &\quad - P(A^{U_2} | D_{U,V} = 1)P(D_{U,V} = 1) \\ &\quad - P(A^{U_2} | D_{U,V} = 0)P(D_{U,V} = 0). \end{aligned}$$

Therefore, using Step 4.1,

$$\begin{aligned} |P(A^{U_2} | B^V) - P(A^{U_2})| &\leq |P(A^{U_2} | B^V, D_{U,V} = 1)P(D_{U,V} = 1 | B^V) \\ &\quad - P(A^{U_2} | D_{U,V} = 1)P(D_{U,V} = 1)| \\ &\quad + |P(A^{U_2} | D_{U,V} = 0)P(D_{U,V} = 0 | B^V) - P(D_{U,V} = 0)| \end{aligned}$$

$$\begin{aligned}
&\leq |P(D_{U,V} = 1|B^V) + P(D_{U,V} = 1)| \\
&\quad + P(A^{U_2}|D_{U,V} = 0)|P(D_{U,V} = 0|B^V) - P(D_{U,V} = 0)| \\
&\leq 2[P(D_{U,V} = 1|B^V) + P(D_{U,V} = 1)].
\end{aligned}$$

This shows that the ϕ -mixing coefficient is bounded by $2[P(D_{U,V} = 1|B^V) + P(D_{U,V} = 1)]$. In other words, the mixing coefficient is bounded by the probability that there exists a link between an individual in U and an individual in V . Recall that we consider the case where $d(U, V) = O(d(i, j))$ for all $(i, j) \in C_{U,V}$.

Because there are at most $4kl$ such pairs, where $|U| = k$ and $|V| = l$, the second bound is bounded by

$$(16kl) \sup_{W_m \in \mathbb{W}_m} \sup_{\{i \in S_m^U, j \in S_m^V\}} P(w_{ij} = 1 | W_{m,-ij}),$$

where $d(U, V) = O(d(i, j))$.

It is then sufficient to show that, as $d(i, j) \rightarrow \infty$, we have

$$\begin{aligned}
\sup_{A^U, B^V} |P(A^U | B^V) - P(A^U)| &\leq kl \sup_{W_m \in \mathbb{W}_m} \sup_{\{i \in S_m^U, j \in S_m^V\}} P(w_{ij} = 1 | W_{m,-ij}) \\
&= kl \sup_{W_m \in \mathbb{W}_m} \sup_{\{i \in S_m^U, j \in S_m^V\}} \frac{\exp\{w_{ij} h_{ij}(W_{m,-ij}; \theta)\}}{1 + \exp\{h_{ij}(W_{m,-ij}; \theta)\}} \\
\bar{\phi}_{k,l}(d) &\leq kl \sup_m \sup_{W_m \in \mathbb{W}_m} \sup_{\{i \in S_m^U, j \in S_m^V : d(i, j) > d\}} \frac{\exp\{w_{ij} h_{ij}(W_{m,-ij}; \theta)\}}{1 + \exp\{h_{ij}(W_{m,-ij}; \theta)\}}. \tag{A.2}
\end{aligned}$$

Under Assumption 5.1, we have

$$\bar{\phi}_{k,l}(d) = O(d^{-2T+1-\epsilon}),$$

for any $k, l < \infty$, which is sufficient for Assumptions 4.1(a) and 4.1(b).

The case of Assumption 4.1(c) is different. Consider the bound (A.2) for $k = 1$, and denote by s the only pair in U ,

$$\begin{aligned}
\bar{\phi}_{1,l}(d) &\leq l \sup_m \sup_{W_m \in \mathbb{W}_m} \sup_{\{i \in s, j \in S_m^V : d(i, j) > d\}} \frac{\exp\{w_{ij} h_{ij}(W_{m,-ij}; \theta)\}}{1 + \exp\{h_{ij}(W_{m,-ij}; \theta)\}} \\
&= lF[\tilde{h}(d)],
\end{aligned}$$

where

$$F[\tilde{h}(d)] \equiv \sup_m \sup_{W_m \in \mathbb{W}_m} \sup_{\{i \in s, j \in S_m^V : d(i, j) > d\}} \frac{\exp\{w_{ij} h_{ij}(W_{m,-ij}; \theta)\}}{1 + \exp\{h_{ij}(W_{m,-ij}; \theta)\}}.$$

Here, we implicitly assume that it is possible to have l pairs located at a distance of exactly d from s . For $d_0 > 0$, as l is increasing, this is impossible. From Lemma A.1(iii) of Jenish and Prucha (2009), there can be at most Cd^{T-1} pairs located at a distance in $[d, d+1]$, where C is a finite constant. The other pairs must be located at a distance of at least $d+1$. In fact, there can be at most $C(d+\iota)^{T-1}$ pairs located at a distance in $[d+\iota, d+\iota+1]$..., and so on. We should therefore rewrite our bound as

$$\bar{\phi}_{1,\infty}(d) \leq C \sum_{\iota=0}^{\infty} C_{d+\iota} (d+\iota)^{T-1} F[\tilde{h}(d+\iota)].$$

From Assumption 5.1, this can be rewritten as

$$\bar{\phi}_{1,\infty}(d) \leq C \sum_{\iota=0}^{\hat{\iota}(d)} (d+\iota)^{T-1} F[\tilde{h}(d+\iota)] + CK \sum_{\iota=\hat{\iota}(d)+1}^{\infty} (d+\iota)^{-T-\epsilon},$$

where K is finite and positive. The function $\hat{\iota}(d)$ maps to a finite and positive integer and represents the first value of ι such that $F[\tilde{h}(d+\iota)] \leq K(d+\iota)^{-2T+1-2\epsilon}$ holds. Note that $\hat{\iota}(d)$ is necessarily non-increasing in d , and hence $\sup_d \hat{\iota}(d) = \hat{\iota}(0)$. Notice that $\hat{\iota}(0)$ is necessarily finite under Assumption 5.1. Thus, we have

$$\bar{\phi}_{1,\infty}(d) \leq C \sum_{\iota=0}^{\hat{\iota}(0)} (d+\iota)^{T-1} F[\tilde{h}(d+\iota)] + CK \sum_{\iota=0}^{\infty} (d+\iota)^{-2T-2\epsilon}$$

The first term is a finite sum of sequences of d , each of which is (at least) $O(d^{-T-\epsilon})$, so the first term is $O(d^{-T-\epsilon})$. For a fixed d , $\sum_{\iota=0}^{\infty} (d+\iota)^{-T-\epsilon}$ converges. Then, the rest of the series converges to zero. This means that for every $\epsilon > 0$, there is a positive integer R such that for all $r \geq r' \geq R$ we have $\sum_{\iota=r'}^{r'} (d+\iota)^{-T-\epsilon} < \epsilon$. Thus, $\sum_{\iota=0}^{\infty} (d+\iota)^{-T-\epsilon} \approx \sum_{\iota=0}^R (d+\iota)^{-T-\epsilon}$. Because every term of the last finite summation is $O(d^{-T-\epsilon})$, we have finally $\bar{\phi}_{1,\infty}(d) = O(d^{-T-\epsilon})$.

CLAIM A.4. If $\mathbb{E}[\Gamma_s \Gamma'_s]$ is non-singular, then $q_s(w_s; \theta) \neq q_s(w_s; \theta_0)$ for $\theta \neq \theta_0$ with positive probability under θ_0 .

Proof: Indeed, if $E[\Gamma_s \Gamma'_s]$ is non-singular for $\theta \neq \theta_0$, $E[\{\Gamma'_s(\theta - \theta_0)\}^2] = (\theta - \theta_0)' E[\Gamma_s \Gamma'_s] (\theta - \theta_0) > 0$. This implies that $\Gamma'_s(\theta - \theta_0) \neq 0$ and then $\Gamma'_s \theta \neq \Gamma'_s \theta_0$ with positive probability under θ_0 . Because the functions $e^x/(1+e^x)$, $1/(1+e^x)$ and $\ln(x)$ are strictly monotonic, the result holds given that $q_s(w_s; \theta)$ can be written also as

$$q_s(w_s; \theta) = w_s \ln \mathbb{P}(w_s = 1 | W_{m,-s}; \theta) + (1 - w_s) \ln \mathbb{P}(w_s = 0 | W_{m,-s}; \theta). \quad \square$$

CLAIM A.5. If $\mathbb{E}[|\Gamma_s \Gamma'_s|] < \infty$, then $\mathbb{E}[|q_s(w_s; \theta)| | W_{m,-s}] < \infty$.

Proof: Indeed, because we can show that $\max(|x - \ln[1 + e^x]|, |\ln[1 + e^x]|) < 1 + x^2$, we have then $E[|q_s(w_s; \theta)| | W_{m,-s}] \leq 1 + \theta' E[\Gamma_s \Gamma'_s] \theta$. $\square$

CLAIM A.6. If

$$\begin{aligned} & |\ln \mathbb{P}(w_s = 1 | W_{m,-s}; \theta)| \mathbb{P}(w_s = 1 | W_{m,-s}; \theta_0) \\ & + |\ln \mathbb{P}(w_s = 0 | W_{m,-s}; \theta)| \mathbb{P}(w_s = 0 | W_{m,-s}; \theta_0) < \infty, \end{aligned}$$

for each $\theta \in \Theta$, then if $q_s(w_s; \theta) \neq q_s(w_s; \theta_0)$ for $\theta \neq \theta_0$, we have $\mathbb{E}[q_s(w_s; \theta)] < \mathbb{E}[q_s(w_s; \theta_0)]$.

Proof: We have

$$\begin{aligned} E[|q_s(w_s; \theta)| | W_{m,-s}, \theta_0] &= \ln \mathbb{P}(w_s = 1 | W_{m,-s}; \theta) \mathbb{P}(w_s = 1 | W_{m,-s}; \theta_0) \\ &+ \ln \mathbb{P}(w_s = 0 | W_{m,-s}; \theta) \mathbb{P}(w_s = 0 | W_{m,-s}; \theta_0). \end{aligned}$$

Thus,

$$\begin{aligned} & \mathbb{E}[|q_s(w_s; \theta)| | W_{m,-s}, \theta_0] - \mathbb{E}[|q_s(w_s; \theta_0)| | W_{m,-s}, \theta_0] \\ &= \ln \left\{ \frac{\mathbb{P}(w_s = 1 | W_{m,-s}; \theta)}{\mathbb{P}(w_s = 1 | W_{m,-s}; \theta_0)} \right\} \mathbb{P}(w_s = 1 | W_{m,-s}; \theta_0) \end{aligned}$$

$$\begin{aligned}
& + \ln \left\{ \frac{\mathbb{P}(w_s = 0 | W_{m,-s}; \theta)}{\mathbb{P}(w_s = 0 | W_{m,-s}; \theta_0)} \right\} \mathbb{P}(w_s = 0 | W_{m,-s}; \theta_0) \\
& < \ln \{ \mathbb{P}(w_s = 1 | W_{m,-s}; \theta) + \mathbb{P}(w_s = 0 | W_{m,-s}; \theta) \} = 0.
\end{aligned}$$

The inequality holds using Jensen's Inequality for strictly concave function. Then, we have $\mathbb{E}[|q_s(w_s; \theta)|] - \mathbb{E}[|q_s(w_s; \theta_0)|] < 0$ by iterated expectation. $\square$

CLAIM A.7. *If $\mathbb{E}[\Gamma_{s,m} \Gamma'_{s,m}]$ is finite and non-singular and $\kappa(m) = O(n)$, then $\lim_{m \rightarrow \infty} E[\mathcal{L}_m(\theta)] < \lim_{m \rightarrow \infty} E[\mathcal{L}_m(\theta_0)]$ for $\theta \neq \theta_0$.*

Proof: Indeed, a direct implication of Claims A.4, A.5, and A.6 is that $\mathbb{E}[q_s(w_s; \theta)] < \mathbb{E}[q_s(w_s; \theta_0)]$. Then, we have $(1/\kappa(m)) \sum_s^m \mathbb{E}[q_s(w_s; \theta)] < (1/\kappa(m)) \sum_s^m \mathbb{E}[q_s(w_s; \theta_0)]$. This also remains valid at the limit if we choose the right order of $\kappa(m)$. $\square$

Proof of Theorem 4.1: Under Assumption 4.2, it is sufficient to show that²⁶

$$\sup_{\theta \in \Theta} |\mathcal{L}_m(\theta) - E[\mathcal{L}_m(\theta)]| \rightarrow_{a.s.} 0, \text{ as } m \rightarrow \infty.$$

In order to show that this condition holds, it is sufficient to show that the conditions of Theorems 2 and 3 from Jenish and Prucha (2009) hold. Specifically,

- (1) $d(r, s) > d_0 > 0$ for any $r, s \in S_m$;
- (2) $(\Theta, \|\cdot\|)$ is a totally bounded metric space;
- (3) domination,

$$\lim_{m \rightarrow \infty} \sup \frac{1}{|S_m|} \sum_{s=1}^m E[\bar{q}_{s,m}^p \mathbf{1}_{\{\bar{q}_{s,m}^p > k\}}] \rightarrow 0 \text{ as } k \rightarrow \infty,$$

for some $p \geq 1$, and where $\bar{q}_{s,m} = \sup_{\theta} |q_{s,m}(w_{s,m}|x, W_m, \theta)|$;

- (4) stochastic equicontinuity, where for every $\epsilon > 0$,

$$\limsup_m \frac{1}{|S_m|} \sum_{s=1}^m P(\sup_{\theta' \in \Theta} \sup_{\theta \in B(\theta', \delta)} |q_{s,m}(\theta) - q_{s,m}(\theta')| > \epsilon) \rightarrow 0 \text{ as } \delta \rightarrow 0,$$

where $B(\theta', \delta)$ is the open ball $\{\theta \in \Theta : \|\theta' - \theta\| < \delta\}$;

- (5) $\sup_m \sup_{s \in S_m} E[\sup_{\theta \in \Theta} |q_{s,m}(\theta)|^{(1+\eta)}] < \infty$ for some $\eta > 0$;
- (6) $\sum_{d=1}^{\infty} d^{T-1} \bar{\phi}_{1,1}(d) < \infty$.

Condition 1 is implied by Assumption 4.3. Condition 2 is verified by construction. Condition 5 holds from Assumption 4.4. Condition 6 holds under Assumption 4.1, as shown in Lemma A.1. Conditions 3 and 4 hold from the following. Under Condition 5, $\sup_{\theta} |q_{s,m}(w_{s,m}|W_m, \theta)|$ is $L^{(1+\eta)}$ -integrable, which implies the uniform $L^{(1+\eta)}$ integrability of $|q_{s,m}(w_{s,m}|W_m, \theta)|$, and then Condition 3 holds.

Also, Lemma A.2 shows that Assumption 4.4 implies Condition 4. Hence, Theorems 2 and 3 from Jenish and Prucha (2009) apply. This completes the proof. $\square$

Proof of Theorem 4.2: We want to show that $\sqrt{\kappa(m)}(\hat{\theta}_m - \theta_0) \Rightarrow N(0, D_0(\theta_0)^{-1} B_0(\theta_0) D_0(\theta_0)^{-1})$. From the mean value theorem, we have that

$$\begin{aligned}
\frac{\partial \mathcal{L}_m(\hat{\theta}_m)}{\partial \theta} &= \frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} + \frac{\partial^2 \mathcal{L}_m(\bar{\theta}_m)}{\partial \theta \partial \theta'} \\
0 &= \frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} + \frac{\partial^2 \mathcal{L}_m(\bar{\theta}_m)}{\partial \theta \partial \theta'} (\hat{\theta}_m - \theta_0).
\end{aligned}$$

²⁶ See, for example, Gallant and White (1988, p. 18).

and

$$\begin{aligned}\sqrt{\kappa(m)}(\hat{\theta}_m - \theta_0) &= -\sqrt{\kappa(m)}\left[\frac{\partial^2 \mathcal{L}_m(\bar{\theta}_m)}{\partial \theta \partial \theta'}\right]^{-1} \frac{\partial \mathcal{L}_m(\theta_0)}{\partial \theta} \\ &= -\left[\frac{\partial^2 \mathcal{L}_m(\bar{\theta}_m)}{\partial \theta \partial \theta'}\right]^{-1} \left[\frac{\sigma_m}{\sqrt{\kappa(m)}}\right] [\sigma_m^{-1} Q_m],\end{aligned}$$

where $\sigma_m^2 = \text{Var}(Q_m)$ and $Q_m = \sum_{s=1}^m \frac{\partial q_{s,m}(\theta_0)}{\partial \theta}$.

Then, it is sufficient to show the following:

$$(a) \frac{\sigma_m^2}{\kappa(m)} \rightarrow B_0(\theta_0); \quad (b) \sigma_m^{-1} Q_m \Rightarrow N(0, I); \quad (c) \left[\frac{\partial^2 \mathcal{L}_m(\bar{\theta}_m)}{\partial \theta \partial \theta'}\right] \xrightarrow{p} D_0(\theta_0).$$

Putting together Lemmata A.3–A.5 completes the proof. $\square$

Proof of Proposition 5.1: The proof proceeds by using Claims A.1–A.3. $\square$

Proof of Proposition 6.1: For simplicity, we omit the dependence on m . Let $h_s(W_{m,-s}; \theta) = \Gamma'_s \theta$ where Γ_s is a vector of explanatory variables. We know from Proposition 3.1 that the conditional distribution is well specified, i.e.

$$\mathbb{P}(w_s | W_{m,-s}; \theta) = \frac{\exp\{w_s \Gamma'_s \theta\}}{1 + \exp\{\Gamma'_s \theta\}}.$$

Therefore, if θ_0 is the true value of interest. The true conditional distribution takes the form

$$\mathbb{P}(w_s | W_{m,-s}; \theta_0) = \frac{\exp\{w_s \Gamma'_s \theta_0\}}{1 + \exp\{\Gamma'_s \theta_0\}}.$$

The proof proceeds by using Claims A.4–A.7. Recall that we have $q_s(w_s; \theta) = \ln \mathbb{P}(w_s | W_{m,-s}; \theta) = w_s \Gamma'_s \theta - \ln[1 + \exp\{\Gamma'_s \theta\}]$.

APPENDIX B: EMPIRICAL CONSIDERATIONS

Bester et al. (2012)

Let $\mathcal{X}$ be partitioned into groups, or clusters: $c = 1, \dots, C$. Bester et al. (2012) propose using the following cluster-variance estimator:

$$\hat{B}_m(\theta) = \frac{1}{\kappa(m)} \sum_{s \in S} \sum_{r \in S} \mathbb{I}(c_s = c_r) \frac{\partial q_{s,m}(\theta)}{\partial \theta} \left(\frac{\partial q_{s,m}(\theta)}{\partial \theta} \right)'$$

where c_s is the group in which $s \in S$ is located. They show in their Proposition 1(ii) that one can use the t -statistic of the cluster-variance estimator, provided that it is rescaled by a factor of $\sqrt{C/C-1}$.

This approach has the advantage of being fast and easy to implement. In practice, the construction of those groups is not necessarily straightforward. Bester et al. (2012) recommend the use of a relatively small number of large groups. However, an important requirement is a boundary condition that states that most of the pairs in groups are located in the interior (i.e. not on the boundary) of those groups in $\mathcal{X}$. Specifically, let $\partial(c_s)$ be the boundary of the group c_s , and let $\bar{c}_m$ be the average number of pairs in a group. Then, one should have $\partial(c_s) < \bar{c}_m^{(T-1)/T}$, where $T \geq 1$ is the dimension of $\mathcal{X}$.

Application: distance and convergence rate

The Add Health database provides information on the (normalized) location of individuals' homes. Following the survey's definition, individuals can only be friends if they belong to the same community. This is similar to Example 6.2. Therefore, we dropped the links between individuals from different neighbourhoods. Also, because the GPS coordinates are normalized to 100 for confidentiality issues, it does not go well with Assumption 4.3. Indeed, under Assumption 4.3 (and, more generally, in reality), more populated communities should occupy more physical space. Therefore, we rescale the distances between pairs using $\sqrt{n_c}$, where n_c is the number of individuals in the community c . Indeed, from Lemma A.1(ii) in Jenish and Prucha (2009), there are $O(d)^T$ individuals located in a ball of radius d . Here, $T = 2$, so a coherent scaling is such that $n_c = d^2$, or $d = \sqrt{n_c}$.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article at the publisher's web site:

Replication files