Research

Cite this article: Hoffmann T, Jones NS. 2020 Inference of a universal social scale and segregation measures using social connectivity kernels. *J. R. Soc. Interface* **17**: 20200638. <http://dx.doi.org/10.1098/rsif.2020.0638>

Received: 7 August 2020

Accepted: 6 October 2020

Subject Category:

Life Sciences—Mathematics interface

Subject Areas:

biomathematics, biocomplexity

Keywords:

social networks, segregation, ego networks, inference

Author for correspondence:

Till Hoffmann

e-mail: t.hoffmann@imperial.ac.uk

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.5178747>.

Inference of a universal social scale and segregation measures using social connectivity kernels

Till Hoffmann and Nick S. Jones

Department of Mathematics, Imperial College London, London, UK

TH, 0000-0003-4403-0722; NSJ, 0000-0002-4083-972X

How people connect with one another is a fundamental question in the social sciences, and the resulting social networks can have a profound impact on our daily lives. Blau offered a powerful explanation: people connect with one another based on their positions in a social space. Yet a principled measure of social distance, allowing comparison within and between societies, remains elusive. We use the connectivity kernel of conditionally independent edge models to develop a family of segregation statistics with desirable properties: they offer an intuitive and universal characteristic scale on social space (facilitating comparison across datasets and societies), are applicable to multivariate and mixed node attributes, and capture segregation at the level of individuals, pairs of individuals and society as a whole. We show that the segregation statistics can induce a metric on Blau space (a space spanned by the attributes of the members of society) and provide maps of two societies. Under a Bayesian paradigm, we infer the parameters of the connectivity kernel from 11 ego-network datasets collected in four surveys in the UK and USA. The importance of different dimensions of Blau space is similar across time and location, suggesting a macroscopically stable social fabric. Physical separation and age differences have the most significant impact on segregation within friendship networks with implications for intergenerational mixing and isolation in later stages of life.

1. Introduction

Peter Blau proposed that individuals connect with one another based on their positions in a high-dimensional space [1], e.g. a space spanned by demographic attributes. With an accrual of large-scale survey data, we now have access to demographic and relational information internationally, for whole societies, and over time [2–4]. Despite this wealth of data, we lack two important quantities: (a) a natural notion of distance in this space, allowing us to determine how far apart individuals are in society, and (b) a universal characteristic scale, allowing the distance between pairs of individuals in one society to be compared to another society with different social dimensions. We will argue that the probability of forming a friendship, intimately related to work on homophily by Blau and his successors [2], offers both a universal characteristic scale and a notion of distance.

Homophily, the tendency for people to connect with others who are alike, is one of the most robust observations of the social sciences and shapes how our society is connected [2]. Quantifying homophily is not only important for understanding why social ties form between some people yet not between others, but the manifestation of homophily as poorly connected social networks can have a significant impact on dynamics unfolding upon them [5]. For example, users of online social networks, such as Facebook and Twitter, tend to connect with others who hold similar political views [6]. They are more likely to be exposed to information that confirms rather than challenges their beliefs [7]. An ‘echo chamber’ effect ensues, leading to polarized opinions [8].

Homophily can also have a detrimental impact on public health: clusters of individuals who mutually reinforce their belief that vaccinations are harmful can raise the likelihood of significant disease outbreaks [9]—even if the vaccination rate is above herd immunity levels on average.

Homophily can be observed in friendships [10,11], networks of discussion partners [3], communication networks [12,13], marital ties [14] and online social networks [15]. Relationships are homogeneous with respect to a wide range of attributes, including age [16,17], sex [17], ethnicity [10,15,18], education [3,17,19], occupation [20], income [12,13,19], religion [21], parental [19] and marital status [22], political ideology [6,7] and geographical location [23–27].

Segregation statistics, also referred to as segregation measures, are often used to quantify homophily [28,29]. Many approaches are based on co-presence in organizational units such as schools [30], voluntary associations [31], occupations [32] or census tracts [33], and we refer to them as *organizational* statistics. Typically, they compare how the distribution of demographic attributes within organizational units differs from the distribution of attributes in the general population. While organizational statistics are applicable whenever data can be stratified according to a variable of interest, they cannot capture segregation at smaller scales than the strata [18]. For example, the ethnic composition of a set of schools may be representative of the general population, indicating that there is no (organizational) segregation. But the social networks *within* schools often exhibit strong ethnic homophily [10,34]. Organizational statistics cannot capture such social segregation.

Social statistics of segregation, such as the assortativity coefficient [35], overcome these limitations by explicitly considering the interactions among individuals [18], but have their own difficulties: first, they usually rely on the existence of discrete groups, such as sex, ethnicity or religion [29], and they are not applicable to continuous attributes, such as age or income. Attributes are often discretized [22,36,37], but the boundaries between categories always suffer from some degree of arbitrariness [33]. Second, segregation for multiple attributes can be quantified independently, but Pelechris & Wei [38], who consider a multivariate generalization of the assortativity coefficient, note that ‘a formal metric that is generally applicable’ remains elusive. Furthermore, social statistics are typically defined as summary statistics of a fully observed social network. Consequently, we cannot easily quantify uncertainties.

In practice, the study of homophily is complicated by the scarcity of high-quality data [18,39]: we need social network data together with demographic information for each person. Online social networks and the widespread use of mobile phones provide us with detailed information about connections between individuals [40], and seemingly private traits such as socio-economic status [41,42], sexual orientation [43], age, gender and political ideology can be inferred [44]. Unfortunately, network features are often used to predict demographic attributes [12,41,42,44], which would confound any study of homophily. Furthermore, ‘data are [...] too revealing in terms of privacy’ but, at the same time, do not provide enough information for researchers [40]. Individuals can be identified in anonymized social networks [45,46], and augmenting the network data with demographic information would make re-identification even easier.

However, censuses and large-scale surveys collect comprehensive demographic information from respondents but usually lack data about their associates. Fortunately, some surveys have included questions about respondents’ friends [19,47], discussion partners [3,48] or support networks [22,49]. The questions used to elicit social ties provide an imperfect observation of the immediate neighbourhood of respondents [50–52].

Building on the successes of conditionally-independent edge models [53] and, in particular, latent space models for social networks [54,55], we consider a generative model for social networks whose members occupy a multidimensional Blau space in §2.1. We discuss desirable properties for social segregation statistics, and, using the generative network model, we develop a suite of statistics applicable to arbitrary attributes in §2.2. The statistics capture segregation at different scales: single individuals, pairs of individuals and society as a whole. Because of both their probabilistic foundations and their construction from the universal notion of a social tie, the segregation statistics have a universal scale, i.e. one unit of segregation has the same implications across different societies and at different times. We show that the segregation statistic for pairs of individuals can be a metric and can thus be used to quantify distance in Blau space. We illustrate the statistics with a simple example, and we show that it reduces to a well-known segregation statistic if the attributes are univariate and categorical: the natural logarithm of Moody’s α index [34].

In §2.4, we derive the posterior for parameters of the conditionally independent edge model given partial observations of social networks obtained from surveys. We apply our approach to nine existing datasets from the UK and two from the USA in §3. Our analysis reveals that the effects of homophily on society are remarkably stable in both countries regardless of time and the specific nature of relationships. Using the suite of segregation statistics, we find that physical separation and age are the most important factors contributing to the segregation of society. In §4, we provide recommendations for conducting surveys to infer homophily in social networks and discuss future work.

2. Methods

2.1. Generative network model

We consider a generative model for social networks for a population of n individuals N who occupy a Blau space $\mathbb{B}$ spanned by their demographic attributes, such as age, income or sex. In contrast to common latent space models [54,55], the attributes are observed, although Hoff *et al.* [54] also consider an extension including covariates. The q -dimensional attribute vector $x_i \in \mathbb{B}$ for each individual $i \in N$ is drawn independently from a distribution $P(x_i)$ of demographic attributes. Elements of the attribute vector can take continuous, ordinal or categorical values. Connections between individuals are encoded by the binary adjacency matrix A such that $A_{ij} = 1$ if j considers i to be a friend and $A_{ij} = 0$ otherwise. We assume that people do not interact with themselves such that $A_{ii} = 0$ for all i , and that connections are undirected, although social ties need not be reciprocated in general [56].

Given the positions of two individuals i and j in Blau space, we assume that connections form independently with probability $p(x_i, x_j)$, i.e. edges are conditionally independent given the attributes of nodes [57]. The assumption of conditionally independent edges can be problematic. For example, it is not

possible to reproduce heavy-tailed degree distributions if the node density is homogeneous and the kernel is translationally invariant [58]. Furthermore, the average degree scales linearly with the number of nodes unless the connectivity kernel ρ is adjusted to compensate [59]. Nevertheless, we use conditionally independent edge models because the connectivity kernel is intuitive, and they can capture salient features of social networks. For example, nodes in high-density regions have larger degrees on average [58]. Similarly, members of the ethnic majority have more social ties in social networks in US high schools [10].

2.2. Developing a model-based segregation statistic

We have so far emphasized the desirability of metrics on social spaces with a universal characteristic scale (in the sense of being comparable between societies). After first developing universal social segregation statistics, including a notion of social separation, we will formulate both a metric and a notion of scale in §2.3.

In addition to addressing the challenges mentioned in §1, a social segregation statistic should satisfy the following properties: first, the statistic should be insensitive to the overall edge density to facilitate comparison of segregation across different networks. Otherwise, the segregation statistic would depend on the size of the population because the edge density scales as n^{-1} if the average degree is approximately constant. Second, following Freeman [60], we would like the statistic to capture the notion that segregation places ‘restrictions on the access of people to one another’. Third, the statistic should be easily interpretable, and it should have a natural notion of the absence of segregation when individuals form connections without regard to their positions in Blau space. For example, the difference of within- and between-group ties considered by Krackhardt & Stern [61] depends on the sizes of the groups even if there is no homophily: there is no natural reference point.

A single statistic cannot capture the complexities of social networks, and we develop a family of statistics applicable at different scales: (a) the *social separation* between any two individuals, (b) the *isolation* experienced by any one individual and (c) the *social strain* experienced by society as a whole. Starting at the microscopic level, we define the *social separation* between two individuals i and j with attributes x and y as the relative log odds for j to connect with i compared to someone who is alike: the log odds intuitively capture the probabilistic nature of the conditionally independent edge model. In particular,

$$\begin{aligned} \varphi(x, y) &= \text{logit } \rho(y, y) - \text{logit } \rho(x, y), \\ \text{where } \text{logit } \rho &= \log\left(\frac{\rho}{1 - \rho}\right) \end{aligned} \quad (2.1)$$

are the log odds for a connection to form with probability ρ [62]. The probability $\rho(y, y)$ for j to connect with someone who is alike serves as a reference point, and the statistic does not depend on the overall edge density. The social separation φ may be understood as the isolation experienced by i with attributes x as a result of the behaviour of j with attributes y . The statistic is zero if two individuals have the same demographic attributes or if they do not discriminate with respect to the attributes on which they differ. For a homophilous connectivity kernel, the statistic is positive and is a *semimetric* for Blau space; we will consider a family of connectivity kernels for which φ is a true metric in §2.3.

Proposition 2.1. *If the connectivity kernel is homophilous, symmetric and the probability $\rho(x, x)$ to connect with others who are alike is independent of x , the social separation φ is a semimetric [63]: it satisfies the properties of a metric, including non-negativity, symmetry and the identity of indiscernibles—except the triangle inequality.*

Proof. First, $\varphi(x, y) \geq 0$ because homophily implies that $\rho(x, y) < \rho(y, y)$ and *logit* is a monotonically increasing function. Second, $\varphi(x, y) = \varphi(y, x)$ because the first term of equation (2.1) is constant by assumption and the second is symmetric because the kernel is symmetric. Third, the statistic is zero for any two individuals with the same attributes by substitution into equation (2.1). Similarly, if $\varphi(x, y) = 0$, then $x = y$ because $\rho(x, y) < \rho(y, y)$ due to homophily. ■

The less likely two people are to connect, the larger the social separation between them. The assumptions required for proposition 2.1 to hold may seem restrictive, but they are satisfied by most studies of spatial networks [23,24,39,58].

Defining social separation in terms of a generative model, i.e. using the connectivity kernel rather than a summary statistic of a particular dataset, provides us with two advantages: first, any uncertainty associated with inferred connectivity kernels naturally propagates to the segregation statistic, as discussed in §3. Second, we can easily consider the properties of the segregation statistic under a variety of generative models without having to resort to computationally expensive Monte Carlo simulations.

For example, consider a stochastic block model (SBM) [53] with K blocks, intra-group connection probability ρ_{same} , and inter-group connection probability $\rho_{\text{different}} < \rho_{\text{same}}$. Substituting into equation (2.1), the social separation between two nodes with block membership x and y is

$$\varphi(x, y) = (1 - \delta_{xy})(\text{logit } \rho_{\text{same}} - \text{logit } \rho_{\text{different}}), \quad (2.2)$$

where δ_{xy} is the Kronecker delta. The social separation only depends on block membership, and it is not affected by the size of each block. For members of different blocks, φ is the difference of log odds ratios for the existence of intra-group ties as opposed to inter-group ties. The social separation is equal to the natural logarithm of the α -index proposed by Moody [34] for categorical attributes x , but φ is applicable to arbitrary attributes and connectivity kernels.

The social separation $\varphi(x, y)$ is not sufficient to quantify segregation at the level of an individual: we also need to consider the distribution $P(y)$ of attributes y of all other members of society, such as their age, sex or other demographics. For an individual with attributes x , we define the *social isolation*

$$\phi(x) = \int dy P(y) \varphi(x, y), \quad (2.3)$$

which quantifies the average social separation between an individual with attribute x and other members of society. For the SBM, we substitute equation (2.2) into equation (2.3) and obtain

$$\phi(x) = (1 - P(x))(\text{logit } \rho_{\text{same}} - \text{logit } \rho_{\text{different}}), \quad (2.4)$$

where $P(x)$ is the probability to belong to block x , and we have used the identity $\sum_{y=1}^K P(y)(1 - \delta_{xy}) = 1 - P(x)$. Members of all blocks experience the same degree of isolation if the blocks are of the same size. If the sizes are unequal, minorities experience more isolation and majority groups experience less isolation. Indeed, ethnic minorities in schools tend to be more isolated and have fewer social ties [10]. The off-diagonal terms of the Hessian of the social isolation ϕ quantify interactions, such as the joint effect of age and ethnic differences.

To understand how segregated society is as a whole, we would like to aggregate the social isolation ϕ , but the appropriate statistic depends on the question at hand. For example, if we wanted to study the most isolated subpopulation of society, we should consider $\max_{x \in \mathbb{B}} \phi(x)$. Here, we take a utilitarian approach and, in line with equation (2.3), define the *social strain* as

$$\Phi = \int dx P(x) \phi(x), \quad (2.5)$$

which quantifies the average social separation among members of the society. It is zero when individuals do not discriminate based on attributes, and it can reach arbitrarily large values in a society comprising multiple groups that are completely disconnected. For the SBM, we substitute equation (2.4) into equation (2.5) and obtain

$$\Phi = \gamma(\logit \rho_{\text{same}} - \logit \rho_{\text{different}}),$$

$$\text{where } \gamma = 1 - \sum_{x=1}^K P^2(x) \quad (2.6)$$

is an index of dispersion [34] and accounts for the relative sizes of the K blocks. The social strain is maximal when the groups are of equal size. If one of the blocks is larger, the social strain and index of dispersion approach zero as the sizes of the minority blocks decrease: members of the majority group experience little social isolation. It is unsurprising that there is no social strain if the society is homogeneous, but the utilitarian approach has a serious limitation: it has little concern for minorities that are not well integrated in society. For equal group sizes, the social strain increases with the number of groups, asymptotically reaching a maximum value of $\logit \rho_{\text{same}} - \logit \rho_{\text{different}}$.

2.3. Distance and scale in Blau space

The social separation takes a simple form if the probability for two individuals to connect is a logistic kernel [54], i.e.

$$\logit \rho(x, y, \theta) = \sum_{l=1}^p \theta_l f_l(x, y), \quad (2.7)$$

where the p -dimensional vector θ_l parametrizes the kernel, and $f_l(x, y)$ is a set of p -dimensional features that are predictive of the connection probability, such as the age difference $f_{\text{age}} = |x_{\text{age}} - y_{\text{age}}|$. Intersectionality can be accounted for by including interaction terms in the feature set. The social separation between x and y comprises contributions from the features of the logistic kernel

$$\varphi(x, y) = \sum_{l=1}^p \varphi_l(x, y), \quad (2.8)$$

$$\text{where } \varphi_l(x, y) = \theta_l(f_l(y, y) - f_l(x, y)) \quad (2.9)$$

is the contribution due to a single feature l . In fact, φ is a true metric for many logistic connectivity kernels.

Proposition 2.2. *The social separation $\varphi(x, y)$ is a metric if the kernel is homophilous, i.e. $\theta_l < 0$, and each feature $f_l(x, y)$ is a constant or a positive affine transform of a metric $d_l(x, y)$, i.e.*

$$f_l(x, y) = a_l d_l(x, y) + b_l, \quad (2.10)$$

where $a_l > 0$ and b_l are the parameters of the affine transform.

Proof. According to proposition 2.1, the social separation $\varphi(x, y)$ is a semimetric, and it comprises contributions from individual features, as illustrated by equation (2.8). Showing that each contribution $\varphi_l(x, y)$ satisfies the triangle inequality is sufficient for $\varphi(x, y)$ to satisfy it, i.e. we require

$$\varphi_l(x, z) \leq \varphi_l(x, y) + \varphi_l(y, z) \quad (2.11)$$

for all l . Substituting equation (2.10) into equation (2.11) yields

$$- \theta_l a_l d_l(x, z) \leq - \theta_l a_l [d_l(x, y) + d_l(y, z)], \quad (2.12)$$

where we have used the metric property $d_l(x, x) = 0$ for all x , and the constant b_l in equation (2.10) vanishes by equation (2.9). The inequality in equation (2.12) holds because $\theta_l < 0$ for homophilous kernels, $a_l > 0$ by assumption, and $d_l(x, y)$ is a metric.

Equation (2.11) is trivially satisfied for a constant feature, such as a bias term controlling the overall edge density. ■

In other words, the social separation statistic is a true measure of *distance* in the social space with a probabilistic interpretation if features are themselves measures of distance, including all the features we consider subsequently. This observation puts Peter Blau's [1] hypothesis that 'the macrostructure of societies can be defined as a multidimensional space of social positions among which people are distributed and which affect their social relations' on a sound statistical footing: *fitting conditionally independent edge models allows us to learn the metric of Blau space*. The metric has a universal scale: *one unit of social separation has the same probabilistic meaning independent of the society under consideration*, facilitating comparison across disparate datasets. Even if two societies have different Blau space dimensions, e.g. a society might exist that strongly discriminates based on characteristics that are not found in other societies, the social separation between a pair of individuals has a common meaning.

2.4. Parameter inference given ego network data

A representative sample of dyads between individuals together with their demographic attributes is not generally available. However, a number of surveys have collected information about the social ties of respondents using name-generator questions which elicit social ties by asking respondents to nominate their friends [22], individuals they feel close to [11] or discussion partners [3,48]. To generate examples of disconnected dyads, we consider a random sample of pairs of individuals. To account for this non-ignorable data collection process, we introduce a variable $I_{ij} \in \{0, 1\}$ indicating whether a particular dyad A_{ij} was observed [64, ch. 8]. The available data thus comprise the demographic attributes x of individuals included in the sample and the dyad state A_{ij} (1 if i and j are connected and 0 otherwise) if it was observed, i.e. $I_{ij} = 1$. Adapting the argument presented by King & Zeng [65] to a Bayesian paradigm, we consider the posterior distribution over kernel parameters θ given the available data

$$P(\theta|A, f, I = 1) \propto P(A|\theta, f, I = 1)P(\theta), \quad (2.13)$$

where $P(\theta)$ is the kernel parameter prior, and $f = f(x, y)$ are features sufficient to evaluate the connectivity kernel given demographic attributes x and y . The observed-data likelihood is

$$P(A|\theta, f, I = 1) = \frac{P(f|A, \theta, I = 1)P(A|\theta, I = 1)}{P(f|\theta, I = 1)}. \quad (2.14)$$

Considering the first term in the numerator of equation (2.14), we note that the distribution over kernel features given the state A of the dyad does not depend on whether it was included in the sample or not. More formally,

$$\begin{aligned} P(f|A, \theta, I = 1) &= P(f|A, \theta) \\ &= \frac{P(A|f, \theta)P(f|\theta)}{P(A|\theta)}. \end{aligned} \quad (2.15)$$

Turning to the denominator in equation (2.14), we find

$$\begin{aligned} P(f|\theta, I = 1) &= \sum_{\alpha=0}^1 P(f|A = \alpha, \theta, I = 1)P(A = \alpha|\theta, I = 1) \\ &= P(f|\theta) \sum_{\alpha=0}^1 P(A = \alpha|f, \theta) \frac{P(A = \alpha|\theta, I = 1)}{P(A = \alpha|\theta)}, \end{aligned} \quad (2.16)$$

where we used the identity in equation (2.15) to arrive at the second line. Substituting equations (2.15) and (2.16) into

equation (2.14), the observed-data likelihood is

$$P(A|\theta, f, I = 1) = \frac{P(A|f, \theta)r(\alpha)}{\sum_{\alpha=0}^1 P(A = \alpha|f, \theta)r(\alpha)},$$

$$\text{where } r(\alpha) = \frac{P(A = \alpha|\theta, I = 1)}{P(A = \alpha|\theta)} \quad (2.17)$$

is the ratio of prevalences of dyad state α in the sample and the general population. In practice, we approximate the prevalence ratio r using the empirical sample prevalence and prior knowledge about the prevalence in the population. The posterior can be evaluated by substituting equation (2.17) into equation (2.13), and we can thus infer the parameters θ from ego network data. See electronic supplementary material, appendix A.2 for details on how to evaluate the observed-data log-likelihood in a numerically stable fashion and electronic supplementary material, appendix B for a validation of the inference methodology using synthetic data. For logistic connectivity kernels, the observed-data likelihood in equation (2.17) resembles a conventional case-control likelihood, e.g. as used by Smith *et al.* [17].

3. Application

3.1. Ego network data collected in surveys

The social ties identified through name-generator questions depend on the nature of the relationship, the mode of administration of the questionnaire (e.g. face-to-face, telephone interview or online survey), and the interviewer [50,51]. Consequently, we do not expect the kernel parameters inferred from different datasets to be completely consistent. In the following investigation of ego networks, we restrict the nature of relationships to friends who are not relatives as much as the available data permit: we are interested in *voluntary* association among members of the population rather than the social structures they were born into [22].

Demographic information about nominees can be collected either by asking seeds about their friends' demographic background [3,48] or by conducting follow-up surveys with nominated friends [19]. The latter seems preferable because respondents may not have complete information about their social contacts. For example, the age of nominees in the British Household Panel Survey (BHPS), a dataset we consider in §3.4, is 60% more likely to be an integer multiple of 10 than it is for seeds—presumably because seeds round the age of their friends to the nearest decade. In anticipation of such challenges, the coding for the nominees is often coarser than for seeds. To compare the demographic attributes of seeds and nominees, we need to unify the coding (see electronic supplementary material, appendix C for details for each dataset). Unfortunately, follow-up surveys require additional resources to interview the nominees and may suffer from low response rates.

3.2. General social survey

The general social survey (GSS) is a nationally representative face-to-face survey of non-institutionalized adults living in the USA. Demographic attributes of seeds are collected regularly and include age, sex, ethnicity, religion and education [16,48]. In 2004, respondents were asked about the demographic background of people 'with whom they discuss important matters', which tends to elicit close ties [50]. We omit all nominees who are not considered to be friends or who are family. Some of the demographic attributes of

seeds and nominees are missing because respondents did not know or refused to provide the information, and we drop dyads associated with individuals with one or more missing attributes, as shown in electronic supplementary material, table C.2. Such a complete-case analysis can introduce biases if the data are not missing completely at random, but handling the missing data in a principled fashion would require us to develop a model for demographic attributes [66].

The coding of age and sex is consistent among seeds and nominees. We aggregate the detailed coding of ethnic and religious attributes of seeds to match the coding of nominees, as shown in electronic supplementary material, table C.1. Kernel features include the absolute age and ordinal education level difference as well as binary indicators for differences along the sex, ethnicity and religion dimensions. For each demographic attribute, we define a feature for the logistic kernel in equation (2.7), as shown in electronic supplementary material, table C.1. To standardize the features $f(x_i, x_j)$, we subtract their mean and divide non-binary features by twice their standard deviation [67]; binary features are not rescaled. The statistics are calculated with respect to a random sample of pairs of seeds. Feature standardization allows us to compare kernel parameters more easily [67] and simplifies the formulation of priors: we use independent, weakly informative Cauchy priors for the kernel parameters such that

$$P(\theta_l) \propto \left[1 + \left(\frac{\theta_l}{\alpha_l} \right)^2 \right]^{-1}.$$

Following Gelman *et al.* [68], we chose the scale parameters $\alpha_l = 2.5$ for $l > 1$ to represent our weak prior belief that changing a feature by one standard deviation is unlikely to change the log odds by more than five: the independent Cauchy distributions regularize the kernel parameters by placing significant prior probability near zero, but their heavy tails allow for significant departures from zero should the data be in support of large parameters. We set $\alpha_1 = 10$ because the parameter θ_1 associated with the constant bias term could change significantly depending on the population size [64, ch. 16].

The inference is performed in two steps: first, we maximize the posterior with respect to the parameters θ using a gradient ascent algorithm. Second, we run a Metropolis–Hastings algorithm to draw samples from the posterior [69]. Summary statistics of the posterior are shown in figure 1a. The connection probabilities decrease quickly with increasing age differences: the odds of connection are reduced by a multiplicative factor of about 0.3 per decade. Ethnic, sex and religious differences all seem to have a similar effect and decrease the odds by a factor of about 0.3 each; a difference of one educational level reduces the odds by a factor of 0.7.

Hipp & Perrin [11] used the logarithm of physical separation as a benchmark to translate the effect of other attributes into distance-equivalents. We instead use age as a benchmark because age is available for most datasets and is typically coded uniformly in years. By contrast, physical separation is often not available or coded heterogeneously across different datasets. For example, the American Life Panel (ALP) only provides location information at the state level (see §3.3), while the BHPS recorded distance between seeds and

Figure 1. Age and physical separation have a strong impact on connection probabilities, and converting parameters into age equivalents makes feature comparison more intuitive. (a) Kernel parameters inferred from ego network data for each dataset. (b) Age equivalents. For binary features (sex, occupation, religion, ethnicity and distance for the American Life Panel), the equivalent number of years corresponds to a change from having the same attribute to having a different attribute. Age equivalents for the American Life Panel are overestimated (see §3.3 for details). Markers represent the posterior median, thick error bars the interquartile range, and thin error bars the 95% credible interval.

nominees as ordinal data (see §3.4). For the GSS, being of a different ethnicity is equivalent to a 9-year age difference, and having a different sex or religion translates to 8 and 7 years, respectively. One educational level, as defined in electronic supplementary material, table C.1, corresponds to 3 years, as shown in figure 1b.

3.3. American Life Panel

The ALP is a nationally representative panel of adults resident in the USA [70]. Panel members are interviewed either using their own internet connection or are provided with a web television to access surveys. Data are collected regularly and each survey has a different focus. In 2009, information about social networks and financial literacy was collected. Demographic attributes included sex, age, ethnicity, education, their state of residence, and whether respondents identified as Hispanic. Respondents were also asked to nominate others with whom they ‘discuss financial matters’ [71]. We only include nominees who are friends of seeds and exclude kinship ties; see electronic supplementary material, table C.3 for details of harmonization of attributes across seeds and nominees.

Homophily with respect to sex and ethnicity is slightly stronger than in the GSS, and educational homophily is weaker, but the inferred parameters are broadly consistent with the GSS. Age differences appear to play less of a role in the discussion of financial matters at first sight, but the inference is severely biased for age. We cannot resolve strong age homophily because data are only recorded in 15-year bins: the small age parameter is likely a result of regression dilution caused by measuring ages imprecisely [72]. Consequently, the age equivalents in figure 1b are inflated. Being resident in a different state has by far the most significant impact on friendship formation.

3.4. British Household Panel Survey and understanding society

The BHPS was a nationally representative face-to-face survey in the UK. It was conducted from 1991 to 2008 and has since been replaced by the Understanding Society Survey (USS). Respondents were asked questions about ‘their closest friends’ every other year as part of the BHPS and every 3 years in the USS. Data include sex, age, occupational status, ethnicity (only in the USS), and how far away friends live [73,74] (see electronic supplementary material, table C.4 for details).

The inferred kernel parameters are largely consistent with the inference for the ALP and GSS in the US suggesting that friendship formation proceeds similarly in the two countries. We have omitted data from the BHPS in 2008 because we identified errors in the coding which have since been confirmed by the Institute of Social and Economic Research [75]. Similarly, we omitted data from the BHPS in 1996 because physical separation between friends was not recorded. As shown in figure 1, homophily seems to have increased in recent years, but the changes are likely the result of a change in methodology rather than a change in behaviour: the BHPS collected friendship information as part of the main survey, whereas the USS used a self-completion questionnaire [76].

3.5. Inferred segregation

To get a better understanding of Blau space and the metric induced by the connectivity kernel, we consider a sample S of 1000 respondents from the GSS and USS. For each sample, we compute the social separation between pairs of respondents to obtain a distance matrix

$$\hat{\phi}_{ij} = \hat{\theta}^T(f(x_j, x_i) - f(x_i, x_j)),$$

Figure 2. A lower-dimensional embedding of the inter-node distances reveals an interpretable social space in the UK and USA. (a,b) The mean and standard deviation of ages as a function of the first embedding dimension as a solid line and a shaded region, respectively. (c,d) A scatter plot of respondents in a two-dimensional embedding space whose coordinates were obtained from the social separation φ using multidimensional scaling. The colour of a marker indicates the respondent's ethnicity. The heat map represents a smoothed estimate of the social isolation ϕ . The 'bands' of individuals in (c) correspond to different occupational statuses, such as employed or retired.

where $\hat{\theta}$ is the posterior median of the kernel parameters discussed in §3.2 and 3.4. We use multidimensional scaling to embed the respondents in a two-dimensional space [77], as shown in figure 2. Figure 2c,d shows the two-dimensional embedding that best approximates the distance matrix in the high-dimensional social space (we omit contribution due to physical space for the USS to make the embeddings comparable).

The first dimension captures the age of respondents, as illustrated in figure 2a,b the mean age increases monotonically as a function of the first embedding dimension, and the standard deviation is small. We evaluated both statistics using Gaussian kernel smoothing [62, ch. 6]. The second dimension captures sex and ethnicity as well as occupational status (for the USS) and education and religion (for the GSS). As expected from equation (2.4), ethnic minorities are more isolated and live on the outskirts of society while the ethnic majority occupies the centre. The embedding suggests that age has the strongest impact on how people form friendships.

Figure 2c,d shows the social isolation ϕ experienced by individuals as a greyscale heat map which we obtained in two steps: first, we evaluated an estimate of the social isolation

$$\hat{\phi}_i = \frac{1}{|S| - 1} \sum_{j \in S: j \neq i} \hat{\varphi}_{ij}.$$

Second, we applied Gaussian kernel smoothing to the social isolation in the embedding space. Respondents occupying the centre of society experience little isolation whereas individuals in the periphery are more isolated. For example, members of the ethnic majority experience an average social isolation of 4.53 (4.47–4.59 95% credible interval) in

the USS and 4.46 (4.10–4.84 95% credible interval) in the GSS. By contrast, the average social isolation among ethnic minorities is 4.99 (4.91–5.07 95% credible interval) and 5.16 (4.73–5.58 95% credible interval): significantly higher than for the ethnic majority.

Similar to the social separation in equation (2.9), the social strain can be broken down into components

$$\Phi_l = \theta_l \int dx dy (f_l(y, y) - f_l(x, y))P(x)P(y) \quad (3.1)$$

because it is a linear functional of φ : each component contributes to the social strain in society. For each of the datasets, we evaluate an estimate of the contributions to the social strain

$$\hat{\Phi}_l = \frac{2\theta_l}{|U|(|U| - 1)} \sum_{i < j \in U: i \neq j} [f_l(x_i, x_j) - f_l(x_i, x_i)],$$

where the sum is over all distinct pairs of seeds U . The contribution Φ_l quantifies the average social separation due to feature l , and it captures the effect of both the connectivity kernel $\rho(x, y)$ and the attribute distribution $P(x)$ in Blau space: neither is sufficient on its own to quantify segregation. Furthermore, Φ and its contributions in equation (3.1) can facilitate comparison across different datasets, as illustrated in figure 3: they have an intuitive interpretation (the probability of edges decreases with increasing segregation) and universal scale (one unit of segregation has the same effect on edge probability).

As might be expected based on previous studies [23–27], physical space has by far the most significant impact on how people connect with one another. Age homophily places the second strongest restriction social connections, and it is more

Figure 3. Physical separation and age differences are the most important factors preventing integration of society for all datasets. Markers represent the posterior median of the contributions to social strain for each feature, thick error bars the interquartile range, and thin error bars the 95% credible interval.

than three times as restrictive as any other feature except physical space. The ALP survey is an exception because of the regression dilution [72] discussed in §3.3. Age homophily is known to be particularly strong for friendship networks [2]. Homophily with respect to sex, ethnicity, education, religion and occupation make similar, smaller contributions to the segregation of friendships. Importantly, the social strain captures the average contribution to social isolation: it can be small either because there is little homophily or because there is a large majority group, as evident from equation (2.6). For example, almost 80% of respondents in the GSS identify as ‘white’ and experience little social isolation due to their ethnicity, whereas minority groups experience more social isolation. On average, social isolation due to ethnicity is small.

4. Discussion

We considered a generative model for social networks embedded in Blau space, a space spanned by the demographic attributes of members of society. We developed a family of segregation statistics with a universal scale (since they are based on the common notion of the probability of a social tie), facilitating comparison between datasets collected at different times or in different cultural contexts. Furthermore, the segregation statistics are applicable to mixed attribute types, have a natural reference point, and an intuitive interpretation: the probability to form connections decreases with increasing segregation. They are applicable at different resolutions: connections, individuals and society as a whole. For certain logistic connectivity kernels, the social separation is a metric for Blau space and allows us to quantify social distance in a principled fashion. The model-based approach facilitates the study of segregation in synthetic social networks, the effect of interventions, and principled quantification of uncertainties in an applied setting.

Based on 11 ego network datasets collected in the UK and USA, we inferred the connectivity kernel $\rho(x, y)$, i.e. the probability for an individual with demographic attributes x to

connect with another with attributes y . Using the kernel, we compared segregation across different datasets along different demographic dimensions and found that physical distance and age have the most significant impact on how well society is connected. We used the Blau space metric to evaluate the social distance among respondents of the GSS and USS. Using a lower-dimensional embedding of the respondents, we explored Blau space, corroborating our findings that age has a profound impact on restricting friendship formation.

The importance of physical distance highlights that our suite of segregation statistics does not distinguish between *choice homophily* and *opportunistic homophily* [78]. The former is a result of individuals having an active preference to connect with others who are alike, whereas the latter is a result of individuals being exposed to others who are similar to them. Opportunistic homophily is likely to be a large contributing factor to spatial homophily because individuals are less likely to encounter people who live far from them. Similarly, the statistics do not discriminate between choice homophily and *social influence*, i.e. the tendency for people to become more alike given a connection [79].

Other features, including sex, ethnicity, religion, education and occupation, have smaller effects on the presence of connections. Notably, the social strain due to ethnicity in the GSS and ALP is larger than in the USS: first, the effect of ethnicity is more pronounced in the US, as shown in figure 1. Second, society in the USA is more ethnically diverse than in the UK (79% white in GSS ‘04 compared with 89% white in USS ‘14), increasing strain on average, as exemplified with a SBM in §2.2.

Even though we did not expect the kernel parameters to be consistent across countries, time, or even different surveys, people connected with one another in a surprisingly similar fashion across the different datasets (the BHPS and USS are longitudinal studies such that consistent parameter estimates are less surprising). Our observations, together with a study by Mossong *et al.* [4] finding that ‘mixing patterns [...] were remarkably similar across different European countries’, suggest that connectivity kernels for friendships vary little

across societies and time. To test this hypothesis, further surveys should be conducted in a unified fashion to minimize the effects of question wording and how the survey is administered [51]. In particular, such surveys should explore options to explicitly incentivise nominees to provide data about themselves [80]: seeds may not recall certain attributes, or nominees may deliberately portray themselves inaccurately [81]. Questions regarding ethnicity should allow respondents to provide multiple answers such that people with mixed ethnic backgrounds can express their identity. Rather than asking respondents about potentially sensitive information, such as income, proxy information that is more readily available—and potentially more informative of how individuals interact with society—could be collected [82]. Whenever possible, aggregation of attributes such as age into bins should be avoided because it limits the ability to infer kernel parameters [72], as we saw in §3.3. Connectivity kernels should be inferred jointly for all dimensions of Blau space to control for social preferences on correlated attributes.

The connectivity kernel is an intuitive model of how people connect with one another, and it is able to reproduce some of the statistics of real social networks. For example, people in high-density regions of Blau space have been observed to have more connections [10]. However, exponential random graph models may be able to better capture the nature of social networks [83]. Furthermore, we have used a connectivity kernel that (a) is symmetric and cannot identify whether there is a status order in society [20,56] and (b) only depends on differences between individuals. For example, young men tend to have more social contacts than

young women, and older women have more social contacts than older men [84]—an observation that cannot be captured by a kernel of the form we have considered. The connectivity kernel could be refined by adding the demographic attributes of the seeds and nominees as features, capturing sociability and popularity, respectively. Furthermore, it should be determined whether the number of ‘intervening opportunities’ [85], absolute distance in Blau space, or a hybrid thereof are most predictive of tie probability. Ultimately, learning a connectivity kernel without a pre-specified parametric form should be considered [86] because they can better capture complex patterns, such as interactions between different demographic attributes. We note that, irrespective of the choice of connectivity kernel, the interpretable segregation statistics considered here remain valid and useful.

Data accessibility. All data are available from the UK data service (<http://doi.org/10.5255/UKDA-SN-6614-13>), RAND American Life Panel (<https://alpdata.rand.org/index.php?page=data&p=showsurvey&syid=86>) and the US General Social Survey (<https://gss.norc.umd.edu/>). Code to reproduce results is available at <https://github.com/tillahoffmann/kernels>.

Authors’ contributions. T.H. data collection, performing data analysis, interpretation of data, statistical and theoretical design, manuscript writing. N.J. interpretation of data, statistical and theoretical design, manuscript writing.

Competing interests. We declare we have no competing interests.

Funding. This work was supported by the Engineering and Physical Sciences Research Council grant no. EP/N014529/1.

Acknowledgements. We thank Sahil Loomba and two anonymous referees for useful feedback on the manuscript and acknowledge support from the grant no. EP/N014529/1.

References

- Blau PM. 1977 A macrosociological theory of social structure. *Am. J. Sociol.* **83**, 26–54. (doi:10.1086/226505)
- McPherson M, Smith-Lovin L, Cook JM. 2001 Birds of a feather: homophily in social networks. *Annu. Rev. Sociol.* **27**, 415–444. (doi:10.1146/annurev.soc.27.1.415)
- McPherson M, Smith-Lovin L, Brashears ME. 2006 Social isolation in America: changes in core discussion networks over two decades. *Am. Sociol. Rev.* **71**, 353–375. (doi:10.1177/000312240607100301)
- Mossong J *et al.* 2008 Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS Med.* **5**, e74. (doi:10.1371/journal.pmed.0050074)
- Golub B, Jackson MO. 2012 How homophily affects the speed of learning and best-response dynamics. *Q. J. Econ.* **127**, 1287–1338. (doi:10.1093/qje/qjs021)
- Boutyline A, Willer R. 2017 The social structure of political echo chambers: variation in ideological homophily in online networks. *Polit. Psychol.* **38**, 551–569. (doi:10.1111/pops.12337)
- Bakshy E, Messing S, Adamic LA. 2015 Exposure to ideologically diverse news and opinion on facebook. *Science* **348**, 1130–1132. (doi:10.1126/science.aaa1160)
- DeMarzo PM, Vayanos D, Zwiebel J. 2003 Persuasion bias, social influence, and unidimensional opinions. *Q. J. Econ.* **118**, 909–968. (doi:10.1162/00335530360698469)
- Salathé M, Bonhoeffer S. 2008 The effect of opinion clustering on disease outbreaks. *J. R. Soc. Interface* **5**, 1505–1508. (doi:10.1098/rsif.2008.0271)
- Currarini S, Jackson MO, Pin P. 2009 An economic model of friendship: homophily, minorities, and segregation. *Econometrica* **77**, 1003–1045. (doi:10.3982/ECTA7528)
- Hipp JR, Perrin AJ. 2009 The simultaneous effect of social distance and physical distance on the formation of neighborhood ties. *City Community* **8**, 5–25. (doi:10.1111/j.1540-6040.2009.01267.x)
- Wang Y, Zang H, Faloutsos M. 2013 Inferring cellular user demographic information using homophily on call graphs. In *2013 IEEE Conf. on Computer Communications Workshops (INFOCOM WKSHPS)*, Turin, Italy, pp. 211–216. See <http://ieeexplore.ieee.org/document/6562897/>.
- Leo Y, Fleury E, Alvarez-Hamelin JL, Sarraute C, Karsai M. 2016 Socioeconomic correlations and stratification in social-communication networks. *J. R. Soc. Interface* **13**, 20160598. (doi:10.1098/rsif.2016.0598)
- Blau P, Schwartz J. 1984 *Crosscutting social circles: testing a macrostructural theory of intergroup relations*. New York, NY: Routledge.
- Chang J, Rosen I, Backstrom L, Marlow C. 2010 ePluribus: ethnicity on social networks. In *ICWSM*.
- Marsden PV. 1988 Homogeneity in confiding relations. *Soc. Netw.* **10**, 57–76. (doi:10.1016/0378-8733(88)90010-X)
- Smith JA, McPherson M, Smith-Lovin L. 2014 Social distance in the United States: sex, race, religion, age, and education homophily among confidants, 1985 to 2004. *Am. Sociol. Rev.* **79**, 432–456. (doi:10.1177/0003122414531776)
- Blumenstock J, Fratamico L. 2013 Social and spatial ethnic segregation: a framework for analyzing segregation with large-scale spatial network data. In *ACM DEV-4 ‘13: Proceedings of the 4th Annual Symposium on Computing for Development, Cape Town, South Africa, December*, pp. 1–10, art. no.: 11. New York, NY: ACM. See <https://doi.org/10.1145/2537052.2537061>.
- Johnson MA. 1989 Variables associated with friendship in an adult population. *J. Soc. Psychol.* **129**, 379–390. (doi:10.1080/00224545.1989.9712054)
- Chan TW, Goldthorpe JH. 2004 Is there a status order in contemporary british society? Evidence from the occupational structure of friendship. *Eur. Sociol. Rev.* **20**, 383–401. (doi:10.1093/esr/jch033)
- Platt L. 2012 Exploring social spaces of Muslims. In *Muslims in Britain: making social and political space*, pp. 53–83. London, UK: Routledge.

22. Kalmijn M, Vermunt JK. 2007 Homogeneity of social networks by age and marital status: a multilevel analysis of ego-centered networks. *Soc. Netw.* **29**, 25–43. (doi:10.1016/j.socnet.2005.11.008)
23. Lambiotte R, Blondel VD, De Kerchove C, Huens E, Prieur C, Smoreda Z, Van Dooren P. 2008 Geographical dispersal of mobile communication networks. *Physica A* **387**, 5317–5325. (doi:10.1016/j.physa.2008.05.014)
24. Expert P, Evans TS, Blondel VD, Lambiotte R. 2011 Uncovering space-independent communities in spatial networks. *Proc. Natl Acad. Sci. USA* **108**, 7663–7668. (doi:10.1073/pnas.1018962108)
25. Backstrom L, Sun E, Marlow C. 2010 Find me if you can: improving geographical prediction with social and spatial proximity. In *WWW '10: Proceedings of the 19th international conference on World Wide Web, April*, pp. 61–70. (doi:10.1145/1772690.1772698)
26. Scellato S, Noulas A, Lambiotte R, Mascolo C. 2011 Socio-spatial properties of online location-based social networks. In *5th Int. AAAI Conf. on Weblogs and Social Media (ICWSM-11), Barcelona, Spain, 17–21 July*. See <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2751>.
27. Illenberger J, Nagel K, Flötteröd G. 2013 The role of spatial interaction in social networks. *Netw. Spat. Econ.* **13**, 255–282. (doi:10.1007/s11067-012-9180-4)
28. Rodriguez-Moral A, Vorsatz M. 2016 An overview of the measurement of segregation: classical approaches and social network analysis. In *Complex networks and dynamics*, pp. 93–119. New York, NY: Springer.
29. Bojanowski M, Corten R. 2014 Measuring segregation in social networks. *Soc. Netw.* **39**, 14–32. (doi:10.1016/j.socnet.2014.04.001)
30. Orfield G, Frankenberg E. 2014 Brown at 60: great progress, a long retreat and an uncertain future. Tech. rep. Civil Rights Project. See <https://civilrightsproject.ucla.edu/research/k-12-education/integration-and-diversity/brown-at-60-great-progress-a-long-retreat-and-an-uncertain-future/Brown-at-60-051814.pdf>.
31. Popielarz PA. 1999 (In)voluntary association: a multilevel analysis of gender segregation in voluntary organizations. *Gender Soc.* **13**, 234–250. (doi:10.1177/089124399013002005)
32. Charles M, Grusky DB. 1995 Models for describing the underlying structure of sex segregation. *Am. J. Sociol.* **100**, 931–971. (doi:10.1086/230605)
33. Reardon SF, O'Sullivan D. 2004 Measures of spatial segregation. *Sociol. Methodol.* **34**, 121–162. (doi:10.1111/j.0081-1750.2004.00150.x)
34. Moody J. 2001 Race, school integration, and friendship segregation in America. *Am. J. Sociol.* **107**, 679–716. (doi:10.1086/338954)
35. Newman MEJ. 2003 Mixing patterns in networks. *Phys. Rev. E* **67**, 026126. (doi:10.1103/PhysRevE.67.026126)
36. Lam Morgan D. 2012 A spatial econometric approach to the study of social influence. PhD thesis. University of Texas Austin.
37. Kim M, Leskovec J. 2012 Multiplicative attribute graph model of real-world networks. *Internet Math.* **8**, 113–160. (doi:10.1080/15427951.2012.625257)
38. Pelechris K, Wei D. 2016 VA-index: quantifying assortativity patterns in networks with multidimensional nodal attributes. *PLoS ONE* **11**, e0146188.
39. Butts CT, Acton RM, Hipp JR, Nagle NN. 2012 Geographical variability and network structure. *Soc. Netw.* **34**, 82–100. (doi:10.1016/j.socnet.2011.08.003)
40. Golder SA, Macy MW. 2014 Digital footprints: opportunities and challenges for online social research. *Annu. Rev. Sociol.* **40**, 129–152. (doi:10.1146/annurev-soc-071913-043145)
41. Blumenstock J, Cadamuro G, On R. 2015 Predicting poverty and wealth from mobile phone metadata. *Science* **350**, 1073–1076. (doi:10.1126/science.aac4420)
42. Luo S, Morone F, Sarraute C, Travizano M, Makse HA. 2017 Inferring personal economic status from social network location. *Nat. Commun.* **8**, 15227. (doi:10.1038/ncomms15227)
43. Wang Y, Kosinski M. 2017 Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Open Sci. Framework* **114**, zn79k.
44. Kosinski M, Stillwell D, Graepel T. 2013 Private traits and attributes are predictable from digital records of human behavior. *Proc. Natl Acad. Sci. USA* **110**, 5802–5805. (doi:10.1073/pnas.1218772110)
45. Backstrom L, Dwork C, Kleinberg J. 2011 Wherefore art thou R3579X?: anonymized social networks, hidden patterns, and structural steganography. *Commun. ACM* **54**, 133–141. (doi:10.1145/2043174.2043199)
46. Narayanan A, Shmatikov V. 2008 Robust de-anonymization of large sparse datasets. In *Symp. on Security and Privacy*, pp. 111–125. (doi:10.1109/SP.2008.33)
47. Huckfeldt RR. 1983 Social contexts, social networks, and urban neighborhoods: environmental constraints on friendship choice. *Am. J. Sociol.* **89**, 651–669. (doi:10.1086/227908)
48. Marsden PV. 1987 Core discussion networks of Americans. *Am. Sociol. Rev.* **52**, 122–131. (doi:10.2307/2095397)
49. Banerjee A, Chandrasekhar AG, Dufflo E, Jackson MO. 2013 The diffusion of microfinance. *Science* **341**, 1236498. (doi:10.1126/science.1236498)
50. Marin A. 2004 Are respondents more likely to list alters with certain characteristics? Implications for name generator data. *Soc. Netw.* **26**, 289–307. (doi:10.1016/j.socnet.2004.06.001)
51. Eagle DE, Proeschold-Bell RJ. 2015 Methodological considerations in the use of name generators and interpreters. *Soc. Netw.* **40**, 75–83. (doi:10.1016/j.socnet.2014.07.005)
52. Eveland Jr WP, Appiah O, Beck PA. 2017 Americans are more exposed to difference than we think: capturing hidden exposure to political and racial difference. *Soc. Netw.* **52**, 192–200. (doi:10.1016/j.socnet.2017.08.002)
53. Snijders TAB. 2011 Statistical models for social networks. *Annu. Rev. Sociol.* **37**, 131–153. (doi:10.1146/annurev.soc.012809.102709)
54. Hoff PD, Raftery AE, Handcock MS. 2002 Latent space approaches to social network analysis. *J. Am. Stat. Assoc.* **97**, 1090–1098. (doi:10.1198/016214502388618906)
55. Hoff PD. 2008 Multiplicative latent factor models for description and prediction of social networks. *Comput. Math. Organ. Theory* **15**, 261–272. (doi:10.1007/s10588-008-9040-4)
56. Ball B, Newman MEJ. 2013 Friendship networks and social status. *Netw. Sci.* **1**, 16–30. (doi:10.1017/nws.2012.4)
57. Fienberg SE. 2012 A brief history of statistical models for network analysis and open challenges. *J. Comput. Graph. Stat.* **21**, 825–839. (doi:10.1080/10618600.2012.738106)
58. Barnett L, Di Paolo E, Bullock S. 2007 Spatially embedded random networks. *Phys. Rev. E* **76**, 056115. (doi:10.1103/PhysRevE.76.056115)
59. Caron F, Fox EB. 2017 Sparse graphs using exchangeable random measures. *J. R. Stat. Soc. B* **79**, 1295–1366. (doi:10.1111/rssb.12233)
60. Freeman LC. 1978 Segregation in social networks. *Sociol. Methods Res.* **6**, 411–429. (doi:10.1177/004912417800600401)
61. Krackhardt D, Stern RN. 1988 Informal networks and organizational crises: an experimental simulation. *Soc. Psychol. Q.* **51**, 123–140. (doi:10.2307/2786835)
62. Hastie T, Tibshirani R, Friedman J. 2009 *The elements of statistical learning: data mining, inference and prediction*. New York, NY: Springer.
63. Wilson WA. 1931 On semi-metric spaces. *Am. J. Math.* **53**, 361–373. (doi:10.2307/2370790)
64. Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB. 2013 *Bayesian data analysis*. Boca Raton, FL: Chapman and Hall/CRC.
65. King G, Zeng L. 2001 Logistic regression in rare events data. *Pol. Anal.* **9**, 137–163. (doi:10.1093/oxfordjournals.pan.a004868)
66. Pigott TD. 2001 A review of methods for missing data. *Educ. Res. Eval.* **7**, 353–383. (doi:10.1076/edre.7.4.353.8937)
67. Gelman A. 2008 Scaling regression inputs by dividing by two standard deviations. *Stat. Med.* **27**, 2865–2873. (doi:10.1002/sim.3107)
68. Gelman A, Jakulin A, Pittau MG, Su Y-S. 2008 A weakly informative default prior distribution for logistic and other regression models. *Ann. Appl. Stat.* **2**, 1360–1383. (doi:10.1214/08-AOAS191)
69. Hastings WK. 1970 Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**, 97–109. (doi:10.1093/biomet/57.1.97)
70. Pollard M, Baird MD. 2017 *The RAND American Life Panel*. Research Report. Santa Monica, CA: RAND. See <https://doi.org/10.7249/RR1651>.
71. Mihaly K. 2009 *American Life Panel: well-being 86 questionnaire*. Santa Monica, CA: RAND corporation.

72. Hutcheon JA, Chiolerio A, Hanley JA. 2010 Random measurement error and regression dilution bias. *BMJ* **340**, 1402–1406. (doi:10.1136/bmj.c2289)
73. Institute for Social and Economic Research. 2000 British Household Panel Survey Questionnaire Wave 10. See https://www.iser.essex.ac.uk/bhps/documentation/pdf_versions/questionnaires/bhpsw10q.pdf.
74. Institute for Social and Economic Research. 2017 UK Household Longitudinal Study Mainstage Questionnaire Wave 3. See https://www.understandingsociety.ac.uk/documentation/mainstage/questionnaire/questionnaire-documents/mainstage/wave-3/Understanding_Society_Wave_3_Questionnaire_v03.pdf.
75. Understanding Society User Support. 2016 Unusual age distribution after conditioning on WJBSTAT in wave R. See <https://www.understandingsociety.ac.uk/support/issues/687>.
76. Understanding Society User Support. 2017 Unexpectedly strong gender homophily in Understanding Society compared with the BHPS. See <https://www.understandingsociety.ac.uk/support/issues/869>.
77. Borg I, Groenen P. 1996 *Modern multidimensional scaling: theory and applications*. New York, NY: Springer.
78. Franz S, Marsili M, Pin P. 2010 Observed choices and underlying opportunities. *Sci. Cult.* **76**, 471–476.
79. Shalizi CR, Thomas AC. 2011 Homophily and contagion are generically confounded in observational social network studies. *Sociol. Methods Res.* **40**, 211–239. (doi:10.1177/0049124111404820)
80. Biernacki P, Waldorf D. 1981 Snowball sampling: problems and techniques of chain referral sampling. *Sociol. Methods Res.* **10**, 141–163. (doi:10.1177/004912418101000205)
81. Bruch E, Feinberg F, Lee KY. 2016 Extracting multistage screening rules from online dating activity data. *Proc. Natl Acad. Sci. USA* **113**, 10 530–10 535. (doi:10.1073/pnas.1522494113)
82. Po JYT, Finlay JE, Brewster MB, Canning D. 2012 Estimating household permanent income from ownership of physical assets. PGDA Working Papers 9712, Program on the Global Demography of Aging. See https://cdn1.sph.harvard.edu/wp-content/uploads/sites/1288/2013/10/PGDA_WP_97.pdf.
83. Wimmer A, Lewis K. 2010 Beyond and below racial homophily: ERG models of a friendship network documented on facebook. *Am. J. Sociol.* **116**, 583–642. (doi:10.1086/653658)
84. Bhattacharya K, Ghosh A, Monsivais D, Dunbar RIM, Kaski K. 2016 Sex differences in social focus across the life cycle in humans. *Open Sci.* **3**, 160097.
85. Stouffer SA. 1940 Intervening opportunities: a theory relating mobility and distance. *Am. Sociol. Rev.* **5**, 845–867. (doi:10.2307/2084520)
86. Frölich M. 2006 Non-parametric regression for binary dependent variables. *Econ. J.* **9**, 511–540. (doi:10.1111/j.1368-423X.2006.00196.x)