

Kent Academic Repository

Full text document (pdf)

Citation for published version

Advani, Arun and Malde, Bansi (2016) Methods to Identify Linear Network Models: A Review. Swiss Journal of Economics and Statistics . ISSN 0303-9692. (In press)

DOI

Link to record in KAR

<http://kar.kent.ac.uk/59894/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Methods to Identify Linear Network Models: A Review

ARUN ADVANI and BANSI MALDE*

30th September 2016

Abstract

In many contexts we may be interested in understanding whether direct connections between agents, such as declared friendships in a classroom or family links in a rural village, affect their outcomes. In this paper we review the literature studying econometric methods for the analysis of linear models of *social effects*, a class that includes the ‘linear-in-means’ local average model, the local aggregate model, and models where network statistics affect outcomes. We provide an overview of the underlying theoretical models, before discussing conditions for identification using observational and experimental/quasi-experimental data.

Key Words: Networks, Social Effects, Peer Effects, Econometrics

JEL Classification: C31, C81, Z13

*Affiliations: Advani - University College London and Institute for Fiscal Studies (E-mail: arun_a@ifs.org.uk); University of Kent and Institute for Fiscal Studies (E-mail: bansi_m@ifs.org.uk). We are grateful to Imran Rasul for his support and guidance on this project. We also thank Richard Blundell, Andreas Dzemski, Toru Kitagawa, Aureo de Paula, and Yves Zenou for their useful comments and suggestions. Financial support from the ESRC-NCRM Node ‘Programme Evaluation for Policy Analysis’, Grant reference RES-576-25-0042 is gratefully acknowledged. This paper draws on material from a longer working paper entitled “Empirical Methods for Networks Data: Social Effects, Network Formation and Measurement Error”.

1 Introduction

Researchers and policymakers are often interested in identifying whether and the extent to which direct connections between agents affect their outcomes. For example, does the schooling performance of an individual depend on that of her friends? Does the health seeking behaviour of one's relatives influence one's own health seeking behaviour? Are firms' investment and pay decisions influenced by the behaviour of firms in the same or closely-related industry? The identification and estimation of such social or network effects – direct spillovers from the characteristics or outcomes of one agent to the outcome of others – is of central interest in empirical research on networks in economics.

This paper reviews recent developments in methods to identify linear peer effect models using networks data – data with detailed information on the exact interactions between agents –, when a single cross-section of data is available. Linear models are the most widely used in empirical work, with many econometric methods developed to work with these, making them a natural choice to consider in this review.¹ Moreover, panel data on the network is rarely available, so that a large swathe of methods focus on the case with a single cross-section of data.²

We provide an overview of a number of commonly used empirical specifications, the underlying theoretical models that generate them; and the conditions for the causal identification of parameters with cross-sectional data. We first consider three 'local' models, where only an agent's direct connections (or *neighbours*) affect his outcome. The three specifications allow this effect to depend on the average outcome, total outcome, or both, of his neighbours. In the absence of information on interactions within a network (or group), identification of social effects parameters is greatly complicated by the so-called reflection problem, a form of simultaneity where it is not possible to identify who influences whom within the network or reference group (Manski, 1993). Information on the exact interactions within a network can break this simultaneity for a wide range of network structures, allowing for identification of social effect parameters.³

Recent theoretical analyses have shown that the structure of networks, as well as positions of agents within them influences agents' overall outcomes on dimensions such as information diffusion, risk sharing, among others (Bloch et al. (2008), Jackson et al. (2012), Banerjee et al. (2013) among others). The availability of detailed networks data has motivated the testing of implications of these

¹Brock and Durlauf (2001) and Brock and Durlauf (2007) study identification of social effects in models with discrete outcomes.

²Panel data on networks is becoming increasingly available. Such longitudinal information can be useful in making identification more credible by, for example, allowing for models that account for fixed unobserved variables that affect both the network and the outcomes of interest. Studies making use of such data to estimate peer effects include Patnam (2013), Goldsmith-Pinkham and Imbens (2013), Comola and Prina (2014), among others. The former studies use the panel data to account for network endogeneity. More commonly, studies often have available data with a panel dimension of outcomes, but only a single measurement of the network. The methods outlined in this paper can be applied in these cases, though stronger restrictions could also be imposed improving the credibility of identification.

³It is possible to overcome the reflection problem without detailed networks data. Possible methods include using conditional variance restrictions as in Graham (2008), or using variation in network/group size as in Lee (2007). Blume et al. (2010) provide a more complete overview of the different methods.

models in recent work. We thus next discuss models where the entire structure of the network might matter for an individual’s outcome. Finally, we discuss how experimental and quasi-experimental variation could be used to provide additional variation to uncover social effects.

Our objective is to provide an overview of methods developed under the assumption that the network is (conditionally) exogenously formed: that is, there are no unobserved individual variables within the network that determine who links with whom. This is an important issue that is the subject of much recent research. A companion paper (Advani and Malde, 2016), as well as other recent reviews (e.g. Graham, 2015; Chandrasekhar, 2015; and de Paula, 2016) provide overviews of this issue and of possible methods to deal with this.

An important issue in the practical estimation of social effects relates to the definition and measurement of the network. Our review proceeds assuming that the researcher perfectly observes and measures the network (or network neighbours) relevant for the outcome(s) of interest. Clearly, the researcher’s choice of network to use will influence the estimated parameters, and potentially lead to different policy implications. However, few existing datasets collect information on more than one type of network, so that in practice, researchers are often restricted in their definition of the network by the available data. Nonetheless, existing studies indicate that the definition of the network matters (often in line with theory): For example, Sacerdote (2001) finds that only students’ individual room-mates affect their college performance, while a broader set of peers matters for decisions related to social group participation; Renna et al. (2008) document that adolescents’ weights are more responsive to those of their friends of the same gender; while Patacchini et al. (2016) find that adolescent friendships lasting longer than one year have persistent effects on adolescents’ education outcomes, while shorter-lived friendships do not. Measurement error on the network will also bias parameter estimates. A more detailed overview of this issue, as well as of methods to deal with it are provided in Advani and Malde (2014) and another companion paper, Advani and Malde (2016).

To illustrate the practical restrictions imposed by each of the different models, empirical specifications and conditions for causal identification, we will use a simple and widely studied question in the education and labour economics literatures: How is a teenager’s schooling performance influenced by his friends? This is also a question of great policy interest.⁴ More specifically, this paper will provide an overview of methods that can yield answers to questions such as “Is a teenager’s schooling performance influenced by the average schooling performance of her friends?”; “Do teenagers gain more utility from studying if their friends also study?”; “What is the relative importance of the average schooling performance of a teenager’s friends, and complementarities arising from one’s friends’ studying decisions in shaping a teenager’s overall schooling performance?”; and “How does a student’s popularity influence his/her schooling performance?” The same methods can also be applied to answer analogous questions about interactions between many other types of agents.

The literature on methods for networks data is broad and developing rapidly. We therefore focus

⁴See Sacerdote (2011) for an overview of this literature.

our review on the issues outlined above, leaving aside a number of other interesting areas, including methods to deal with endogenous network formation and measurement error in the network, that have been surveyed elsewhere. In addition, though many of the methods reviewed here either build on or apply methods developed in spatial econometrics, it is not our objective to provide an overview of spatial econometric methods (see instead Anselin, 1988). Boucher and Fortin (2015) provide a complementary review, though they do not cover methods using experimental and quasi-experimental variation.

The rest of the paper is organised as follows. Section 2 outlines the notation used in this paper. Section 3 provides an overview of the local average model, where the average of neighbours' outcomes are allowed to affect an individual's outcome. In Section 4 we cover the local aggregate model, where instead it is the total of neighbours' outcomes that matter. In Section 5 we consider hybrid local models, which allow both local average and local aggregate type effects. We consider models including network statistics that are not purely local, in Section 6. Finally, we show how experimental and quasi-experimental variation can be used to identify social effects in Section 7, before concluding.

2 Notation

We begin by outlining the notation we use throughout the paper. We define a *network* or *graph* $g = (\mathcal{N}_g, \mathcal{E}_g)$ as a set of nodes, \mathcal{N}_g , and edges or links, \mathcal{E}_g .⁵ The nodes represent individual agents, and the edges represent the links between pairs of nodes. In economic applications, nodes are usually individuals, households, firms or countries. Edges could be social ties such as friendship, kinship, or co-working, or economic ties such as purchases, loans, or trade. The number of nodes present in g is $N_g = |\mathcal{N}_g|$, and the number of edges is $E_g = |\mathcal{E}_g|$. We define $\mathcal{G}_N = \{g : |\mathcal{N}_g| = N\}$ as the set of all possible networks on N nodes.

In the simplest case – the *binary network* – any (ordered) pair of nodes $i, j \in \mathcal{N}_g$ is either linked, $ij \in \mathcal{E}_g$, or not linked, $ij \notin \mathcal{E}_g$. If $ij \in \mathcal{E}_g$ then j is described as being a *neighbour* of i . We denote by $nei_{i,g} = \{j : ij \in \mathcal{E}_g\}$ the *neighbourhood* of node i , which contains all nodes with whom i is linked. Nodes that are neighbours of neighbours will be referred to as ‘*second degree neighbour*’. Typically it is convenient to assume that $ii \notin \mathcal{E}_g \forall i \in \mathcal{N}_g$. Edges may be directed, so that a link from node i to node j is not the same as a link from node j to node i ; in this case the network is a *directed graph* (or *digraph*).

Network graphs, whether directed or not, can also be represented by an *adjacency matrix*, \mathbf{G}_g , with typical element $G_{ij,g}$. This is an $N_g \times N_g$ matrix with the leading diagonal normalised to 0. When the network is binary, $G_{ij,g} = 1$ if $ij \in \mathcal{E}_g$, and 0 otherwise, while for weighted graphs, $G_{ij,g} = wei(i, j)$. We will use the notation $\mathbf{G}_{i,g}$ to denote the i^{th} row of the adjacency matrix \mathbf{G}_g , and $\mathbf{G}'_{i,g}$ to denote

⁵In a slight abuse of notation, we will also use g to index individual networks when data from multiple networks is available.

its i^{th} column.⁶ Many models defined for binary networks make use of the row-stochastic adjacency matrix or *influence matrix*, $\tilde{\mathbf{G}}_g$, whose elements are defined as $\tilde{G}_{ij,g} = G_{ij,g}/\sum_j G_{ij,g}$.⁷

In what follows, we will frequently work with data from a number of network graphs. Graphs for different networks will be indexed, in a slight abuse of notation, by $g = 1, \dots, M$, where M is the total number of networks in the data. Node-level variables will be indexed with $i = 1, \dots, N_g$, where N_g is the number of nodes in graph g . Node-level outcomes will be denoted by $y_{i,g}$, while exogenous covariates will be denoted by the $1 \times K$ vector $\mathbf{x}_{i,g}$ and common network-level variables will be collected in the $1 \times Q$ vector, \mathbf{z}_g .

The node-level outcomes, covariates and network-level variables can be stacked for each node in a network. In this case, we will denote the stacked $N_g \times 1$ outcome vector as \mathbf{y}_g and the $N_g \times K$ matrix stacking node-level vectors of covariates for graph g as \mathbf{X}_g . Common network-level variables for graph g will be gathered in the matrix $\mathbf{Z}_g = \boldsymbol{\iota}_g \mathbf{z}_g$ where $\boldsymbol{\iota}_g$ denotes an $N_g \times 1$ vector of ones. The adjacency and influence matrices for network g will be denoted by \mathbf{G}_g and $\tilde{\mathbf{G}}_g$. At times we will also make use of the $N_g \times N_g$ identity matrix, \mathbf{I}_g , consisting of ones on the leading diagonal, and zeros elsewhere.

Finally, we introduce notation for vectors and matrices stacking together the network-level outcome vectors, covariate matrices and adjacency matrices for all networks in the data. $\mathbf{Y} = (\mathbf{y}'_1, \dots, \mathbf{y}'_M)'$ is an $\sum_{g=1}^M N_g \times 1$ vector that stacks together the outcome vectors; $\mathbf{G} = \text{diag}\{\mathbf{G}_g\}_{g=1}^{g=M}$ denotes the $\sum_{g=1}^M N_g \times \sum_{g=1}^M N_g$ block-diagonal matrix with network-level adjacency matrices along the leading diagonal and zeros off the diagonal, and analogously $\tilde{\mathbf{G}} = \text{diag}\{\tilde{\mathbf{G}}_g\}_{g=1}^{g=M}$ (with similar dimensions as \mathbf{G}) for the influence matrices; and $\mathbf{X} = (\mathbf{X}'_1, \dots, \mathbf{X}'_M)'$ and $\mathbf{Z} = (\mathbf{Z}'_1, \dots, \mathbf{Z}'_M)'$ are respectively, $\sum_{g=1}^M N_g \times K$ and $\sum_{g=1}^M N_g \times Q$ matrices, that stack together the covariate matrices across networks. Finally, we define the vector $\boldsymbol{\iota}$ as a $\sum_{g=1}^M N_g \times 1$ vector of ones and the matrix $\mathbf{L} = \text{diag}\{\boldsymbol{\iota}_g\}_{g=1}^{g=M}$, as an $\sum_{g=1}^M N_g \times M$ matrix with each column being an indicator for being in a particular network.

3 Local Average Models

3.1 Set Up

In local average models, an agent's outcome (or choice) is influenced by the average outcome of its neighbours.⁸ Thus, an individual's schooling effort or performance is influenced by the average schooling effort or performance of his friends. We characterise the individual as a node, i , in network g , with outcome $y_{i,g}$. This outcome is modelled as being influenced by the individual's own observed characteristics, $\mathbf{x}_{i,g}$, scalar unobserved heterogeneity $\varepsilon_{i,g}$, observed network characteristics \mathbf{z}_g , an unobserved network characteristic ν_g , and the average outcomes and characteristics of neighbours,

⁶ $\mathbf{G}'_{i,g}$ is the i^{th} row of \mathbf{G}'_g , which is the i^{th} column of \mathbf{G}_g .

⁷A row stochastic (also called 'right stochastic' matrix) is one whose rows are normalised so they each sum to one.

⁸Identification of parameters in such a model was first systematically considered by Manski (1993).

$\sum_{j=1}^{N_g} \tilde{G}_{ij,g} y_{j,g}$ and $\sum_{k=1}^K \sum_{j=1}^{N_g} \tilde{G}_{ij,g} x_{j,k,g}$. Below, we consider identification conditions when data are available from multiple networks, though some results apply to data from a single network.⁹

Stacking together data from multiple networks yields the following empirical specification, expressed in matrix terms:

$$\mathbf{Y} = \alpha \boldsymbol{\iota} + \beta \tilde{\mathbf{G}} \mathbf{Y} + \mathbf{X} \boldsymbol{\gamma} + \tilde{\mathbf{G}} \mathbf{X} \boldsymbol{\delta} + \mathbf{Z} \boldsymbol{\eta} + \mathbf{L} \boldsymbol{\nu} + \boldsymbol{\varepsilon} \quad (3.1)$$

where \mathbf{Y} , $\boldsymbol{\iota}$, $\tilde{\mathbf{G}}$, \mathbf{X} , \mathbf{Z} , \mathbf{L} and $\boldsymbol{\nu}$ are as defined previously. Pre multiplying the vector \mathbf{Y} by $\tilde{\mathbf{G}}$ gives a vector containing, for each individual, the average outcome of his neighbours, and similarly $\tilde{\mathbf{G}} \mathbf{X}$ is a vector of the average characteristics of his friends. The social effect of interest is β , the effect of an increase in the mean of neighbours' outcome on the individual's outcome. This is often described as the 'endogenous social effect', in contrast to the (vector of) 'contextual effect(s)' (or 'exogenous social effect(s)'), $\boldsymbol{\delta}$, which represent the effect of an increase in neighbours' *characteristics*. $\boldsymbol{\nu}$ is described as the 'correlated effect', capturing the correlation in individuals' outcomes simply due to common (unobserved) shocks.

Given the simple empirical form of this model, it has been widely applied in the economics literature. Examples include:

- Understanding how the average schooling performance of an individual's peers influences the individual's own performance in a setting where students share a number of different classes (*e.g.* De Giorgi et al., 2010), or where students have some (but not all) common friends (*e.g.* Bramoullé et al., 2009).
- Understanding how non-market links between firms arising from company directors being members of multiple company boards influence firm choices on investment and executive pay (*e.g.* Patnam, 2013).

Although this specification is widely used in the empirical literature, few studies consider or acknowledge the form of its underlying economic model, even though parameter estimates are subsequently used to evaluate alternative policies and to make policy recommendations. Indeed, parameters are typically interpreted as in the econometric model of Manski (1993), whose parameters do not map back to 'deep' structural (*i.e.* policy invariant) parameters without an economic model.

3.2 Theoretical Foundations

An economic model that leads to this specification is one where nodes have a desire to conform to the average behaviour and characteristics of their neighbours (Akerlof, 1980; Jones, 1984; Bernheim,

⁹When data on only a single network are available, the empirical specification is as follows: $\mathbf{y}_g = \tilde{\nu}_g \boldsymbol{\iota} + \beta \tilde{\mathbf{G}}_g \mathbf{y}_g + \mathbf{X}_g \boldsymbol{\gamma} + \tilde{\mathbf{G}}_g \mathbf{X}_g \boldsymbol{\delta} + \boldsymbol{\varepsilon}_g$, where $\tilde{\nu}_g = \alpha + \mathbf{z}_g \boldsymbol{\eta} + \nu_g$ in our earlier notation, capturing all of the network-level characteristics.

1994; Patacchini and Zenou, 2012). In our schooling example, conformism implies that individuals would want to exert a similar amount of effort in their school work as their friends so as to ‘fit in’. Conditional on the individual’s own characteristics which affect his school effort, if his friends exert relatively low effort in their school work, he will reduce his effort. Below we write out this model more formally and demonstrate how it leads to Equation 3.1.

Conformism is commonly modelled by individual payoffs that are decreasing in the distance between own outcome and network neighbours’ average outcomes.¹⁰ Payoffs are also allowed to vary with an individual heterogeneity parameter, $\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g})$, which captures the individual’s ability or productivity associated with the outcome:

$$U_i(y_{i,g}; \mathbf{y}_{-i,g}, \mathbf{X}_g, \tilde{\mathbf{G}}_{i,g}) = \left(\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g})_{i,g} - \frac{1}{2} \left(y_{i,g} - 2\beta \sum_{j=1}^{N_g} \tilde{G}_{ij,g} y_{j,g} \right) \right) y_{i,g} \quad (3.2)$$

β in Equation 3.2 can be thought of as a taste for conformism. Although we write this model as though individuals are perfectly able to observe each others’ actions, this assumption can be relaxed. In particular, an econometric specification similar to Equation 3.1 can be obtained from a static model with imperfect information (see Blume et al., 2015).

The best response function derived from the first order condition with respect to $y_{i,g}$ is:

$$y_{i,g} = \pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g}) + \beta \sum_{j=1}^{N_g} \tilde{G}_{ij,g} y_{j,g} \quad (3.3)$$

Patacchini and Zenou (2012) derive the conditions under which a Nash equilibrium exists, and characterise properties of this equilibrium.

Note that this is not the only economic model that leads to an empirical specification of this form: a similar specification arises from, for example, models of perfect risk sharing.¹¹ Here, when preferences are homogeneous and risk is perfectly shared, the consumption of risk-averse households will co-move with the average household consumption in the risk sharing group or network (Townsend, 1994).

The individual heterogeneity parameter, $\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g})$, can be modelled as a linear function of observed and unobserved individual and network characteristics:

$$\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g}) = \mathbf{x}_{i,g} \boldsymbol{\gamma} + \sum_{j=1}^{N_g} \tilde{G}_{ij,g} \mathbf{x}_{j,g} \boldsymbol{\delta} + \mathbf{z}_g \boldsymbol{\eta} + \nu_g + \varepsilon_{i,g} \quad (3.4)$$

¹⁰Notice that in Equation 3.2, $\sum_{j=1}^{N_g} \tilde{G}_{ij,g} y_{j,g}$ is identical to the i^{th} row of $\tilde{\mathbf{G}}_g \mathbf{y}_g$, which appears in Equation 3.1.

¹¹Consumer demand models where, given prices, an individual’s demand increases with average demand of some reference group also generate a similar specification (see Gaertner, 1974; Pollak, 1976; Alessie and Kapteyn, 1991; Case, 1991).

Substituting for this in Equation 3.3, we obtain the following best response function for individual outcomes:

$$y_{i,g} = \beta \sum_{j=1}^{N_g} \tilde{G}_{ij,g} y_{j,g} + \mathbf{x}_{i,g} \boldsymbol{\gamma} + \sum_{j=1}^{N_g} \tilde{G}_{ij,g} \mathbf{x}_{j,g} \boldsymbol{\delta} + \mathbf{z}_g \boldsymbol{\eta} + \nu_g + \varepsilon_{i,g} \quad (3.5)$$

Then, stacking observations for all nodes in multiple networks, we obtain Equation 3.1, which can be taken to the data.

3.3 Identification

3.3.1 Without Network Fixed Effects

Bramoullé et al. (2009) study the identification and estimation of Equation 3.1 in observational data with detailed network information or data from partially overlapping peer groups.¹² To proceed further, one needs to make some assumptions on the relationship between the unobserved variables – $\boldsymbol{\nu}$ and $\boldsymbol{\varepsilon}$ – and the other right hand side variables in Equation 3.1.

We first consider identification under the assumptions that $\mathbb{E}[\boldsymbol{\varepsilon} | \mathbf{X}, \mathbf{Z}, \tilde{\mathbf{G}}] = 0$ and $\mathbb{E}[\boldsymbol{\nu} | \mathbf{X}, \mathbf{Z}, \tilde{\mathbf{G}}] = 0$ *i.e.* both the individual level error term, $\boldsymbol{\varepsilon}$ and the network level unobservable are assumed to be mean independent of the observed individual and network-level characteristics and of the network. We will later relax the assumption on $\boldsymbol{\nu}$.

Under these assumptions, the parameters $\{\alpha, \beta, \boldsymbol{\gamma}, \boldsymbol{\delta}, \boldsymbol{\eta}\}$ are identified if $\{\mathbf{I}, \tilde{\mathbf{G}}, \tilde{\mathbf{G}}^2\}$ are linearly independent. Identification thus relies on the network structure. In particular, the condition would not hold in networks composed only of cliques – subnetworks comprising of completely connected components – of the same size, and where the diagonal terms in the influence matrix, $\tilde{\mathbf{G}}$ are not set to 0. In this case, $\tilde{\mathbf{G}}^2$ can be expressed as a linear function of \mathbf{I} and $\tilde{\mathbf{G}}$. Moreover, the model is then similar to the single peer group case of Manski (1993), and the methods outlined in Blume et al. (2010) apply.

In an undirected network (such as the in the left panel in Figure 1 below), this identification condition holds when there exists a triple of nodes (i, j, k) such that i is connected to j but not k , and j is connected to k . The exogenous characteristics of k , $\mathbf{x}_{k,g}$, directly affect j 's outcome, but not (directly) that of i , hence forming valid instruments for the outcome of i 's neighbours (*i.e.* j 's outcome) in the equation for node i . Intuitively this method uses the characteristics of second-degree neighbours who are not direct neighbours as instruments for outcomes of direct neighbours.

¹²Similar identification results have been independently described by De Giorgi et al. (2010), who have data with overlapping peer groups of students who share a number of classes.



Figure 1: Intransitive triad in a undirected network (left panel) and a directed network (right panel)

It is thus immediately apparent why identification fails in networks composed only of cliques: in such networks, there is no triple of nodes (i, j, k) such that i is connected to j , and j is connected to k , but i is not connected to k .

In the directed network case, the condition is somewhat weaker, requiring only the presence of an intransitive triad: that is, a triple such that $ij \in \mathcal{E}$, $jk \in \mathcal{E}$ and $ik \notin \mathcal{E}$ (as in the right panel of Figure 1 above).¹³ This is weaker than in undirected networks, which would also require that $ki \notin \mathcal{E}$.

These conditions impose strong restrictions on behaviour. The identification condition for undirected networks, for example, relies on i and k not influencing one another directly, which might be too strong an assumption in contexts such as within-classroom networks, where it is not unreasonable to assume that all students are likely to interact with one another, and hence influence one another. Similarly, the identification assumption in a directed network is likely to plausibly hold in specific contexts only: for example, the effort of check-out workers who face the same direction (or are arranged in a circle) would be influenced by that of the colleagues that they can see, satisfying the identification condition. Conversely, such an assumption might be too strong to justify in other teams where interactions are not as rigidly fixed.

Let us consider how this method could be applied to identify the influence of the average schooling performance of an individual's friends on the individual. Potential variables that one might include as controls include the individual's age, gender, and parental income; the average age, gender, and parental income of his friends; and some observed school characteristics such as expenditure per pupil. Assume first that the underlying friendship network in this school is undirected as in the left panel of Figure 1, so that if i considers j to be his friend, j also considers i to be his friend. j also has a friend k who is not friends with i . We could then use the age, gender, and parental income of k as instruments for the schooling performance of j in the equation for i . If instead, the network were directed as in the right panel of Figure 1, which might be the case for outcomes such as risky behaviours, where the arrows indicate who is affected by whom (*i.e.* i is affected by j in the Figure, and so on), we can still use the age, gender, and parental income of k as instruments for the school performance of j in the equation for i even though k is connected with i . This is

¹³Equivalently, a triple such $ji \in \mathcal{E}$, $kj \in \mathcal{E}$ and $ki \notin \mathcal{E}$ forms an intransitive triad.

possible since the direction of the relationship is such that k 's school performance is affected by i 's performance, but the converse is not true.

3.3.2 Including Network Fixed Effects

The identification result above requires that the network-level unobservable term be mean independent of the observed covariates, \mathbf{X} and \mathbf{Z} , and of the network, $\tilde{\mathbf{G}}$. However, in many circumstances one might be concerned that unobservable characteristics of the network might be correlated with \mathbf{X} , so that $\mathbb{E}[\boldsymbol{\nu}|\mathbf{X}, \mathbf{Z}, \tilde{\mathbf{G}}] \neq 0$. In the schooling context, where the network of interest is often constrained to be within the same school, a large literature (e.g. Black, 1999, Gibbons and Machin, 2003, Bayer et al., 2007 among others) indicates that wealthier parents choose to live in areas with good schools, making it likely that children with higher parental income will be in schools with teachers who have better unobserved teaching abilities. A natural solution is to include network fixed effects, $\mathbf{L}\tilde{\boldsymbol{\nu}}$, which will control for the network-level unobservable, $\mathbf{L}\boldsymbol{\nu}$, though at the cost of meaning we cannot identify $\boldsymbol{\nu}$.

Since the fixed effects themselves are generally not of interest, to ease estimation they are removed using a *within transformation*. This is done by pre multiplying Equation 3.1 by \mathbf{J} , a block diagonal matrix that stacks the network-level transformation matrices $\mathbf{J}_g = \mathbf{I}_g - \frac{1}{N_g}(\boldsymbol{\iota}_g\boldsymbol{\iota}'_g)$ along the leading diagonal, and off-diagonal terms are set to zero. This subtracts the network-level mean of the outcome, giving a vector $\mathbf{J}\mathbf{Y}$ of deviations from the mean. The resulting model is of the following form:

$$\mathbf{J}\mathbf{Y} = \beta\mathbf{J}\tilde{\mathbf{G}}\mathbf{Y} + \mathbf{J}\mathbf{X}\boldsymbol{\gamma} + \mathbf{J}\tilde{\mathbf{G}}\mathbf{X}\boldsymbol{\delta} + \mathbf{J}\boldsymbol{\varepsilon} \quad (3.6)$$

The identification condition here imposes a stronger requirement on network structure: now, the matrices $\{\mathbf{I}, \tilde{\mathbf{G}}, \tilde{\mathbf{G}}^2, \tilde{\mathbf{G}}^3\}$ need to be linearly independent. This requires that there exists a pair of agents (i, l) such that the shortest path between them is of length 3. That is, i would need to go through at least two other nodes to get to l (as in Figure 2 below). The presence of at least two intermediate agents allows the researcher to use the characteristics of third-degree neighbours (neighbours-of-neighbours-of-neighbours who are not direct neighbours or neighbours-of-neighbours) as an additional instrument to account for the network fixed effect.




Figure 2: Identification with network fixed effects. The picture on the left panel shows an undirected network with an agent l who is at least 3 steps away from i , while the picture on the right panel shows the same for a directed network.

A concern that arises when applying this method is that of instrument strength. Bramoullé et al. (2009) find that this varies with graph *density*, *i.e.*, the proportion of node pairs that are linked; and the level of *clustering*, *i.e.* the proportion of node triples such that precisely two of the possible three edges are connected.¹⁴ Instrument strength is declining in density, since the number of intransitive triads tends to zero, reducing the variation in the instrument. The results for clustering are non-monotone, and depend on density.

The discussion thus far has assumed that the network through which the endogenous social effect operates is the same as the network through which the contextual effect operates. It is possible to allow for these two networks to be distinct. This could be useful in the school setting, for instance, where contextual effects could be driven by the average characteristics of all students in the school, while endogenous effects by the outcomes of a subset of students who are friends. Such a formulation allows for a more flexible representation of the environment: the contextual effect could operate through, for example, the resources available to a school, which might depend on the parental income of all students (if schools are financed through local taxation); while peer influences on effort or performance might come only from friends.

Let $\mathbf{G}_{\mathbf{X},g}$ and $\mathbf{G}_{\mathbf{y},g}$ denote the network-level adjacency matrices through which, respectively, the contextual and endogenous effects operate. As before we define the block diagonal matrices $\mathbf{G}_{\mathbf{X}} = \text{diag}\{\mathbf{G}_{\mathbf{X},g}\}_{g=1}^{g=M}$ and $\mathbf{G}_{\mathbf{y}} = \text{diag}\{\mathbf{G}_{\mathbf{y},g}\}_{g=1}^{g=M}$. Blume et al. (2015) study identification of this model assuming that both networks are (conditionally) exogenous and show that when the matrices $\mathbf{G}_{\mathbf{y}}$ and $\mathbf{G}_{\mathbf{X}}$ are observed by the econometrician, and at least one of δ and γ is non-zero, then the necessary and sufficient conditions for the parameters of Equation 3.1 to be identified are that the matrices \mathbf{I} , $\mathbf{G}_{\mathbf{y}}$, $\mathbf{G}_{\mathbf{X}}$ and $\mathbf{G}_{\mathbf{y}}\mathbf{G}_{\mathbf{X}}$ are linearly independent.

¹⁴This definition is also referred to as the clustering coefficient.

4 Local Aggregate Model

The local aggregate class of models, studied theoretically by Ballester et al. (2006), Calvó-Armengol et al. (2009) and Bramoullé et al. (2014); and empirically by Calvó-Armengol et al. (2009), Lee and Liu (2010) and Liu et al. (2014b), considers settings where agents' utilities are a function of the aggregate outcomes (or choices) of their neighbours. This is in contrast with local average models, where utilities are a function of the difference between own outcome and the average outcome of his neighbours. The local aggregate model thus encompasses a different assumption on behaviour – it applies to situations where there are strategic complementarities or strategic substitutabilities – and, allows dyad-level social effects to aggregate at network or group-level. An agent's utility is thus influenced by complementarities/substitutabilities with all his neighbours, rather than by the distance between his choices and those of his neighbours, as is the case with the local average model. Consequently, as will be seen below, the social effect will vary across agents depending on their network position.

Examples of cases where strategic complementarities and substitutabilities are likely to be at play, and thus where a local aggregate, rather than local average model would be more appropriate include:

- An individual's costs of engaging in crime may be lower when his neighbours also engage in crime (*e.g.* Bramoullé et al., 2014).¹⁵
- An agent is more likely to learn about a new product and how it works if more of his neighbours know about it and have used it.
- A student gains more utility from undertaking effort in studying if his friends also undertake more effort (*e.g.* Calvó-Armengol et al. 2009)

4.1 Set Up and Theoretical Foundations

The local aggregate model corresponds to the following empirical specification:

$$\mathbf{Y} = \alpha\mathbf{1} + \beta\mathbf{G}\mathbf{Y} + \mathbf{X}\boldsymbol{\gamma} + \mathbf{G}\mathbf{X}\boldsymbol{\delta} + \mathbf{Z}\boldsymbol{\eta} + \mathbf{L}\boldsymbol{\nu} + \boldsymbol{\varepsilon} \quad (4.1)$$

where \mathbf{Y} , \mathbf{G} , \mathbf{X} and \mathbf{Z} are as defined earlier. In contrast to the local average model, $\tilde{\mathbf{G}}$ has been replaced by \mathbf{G} . As before, the social effect of interest is β , which now represents the effect of the aggregate outcome of my neighbours on my own outcome.

This specification can be motivated by the best responses of a game in which nodes have linear-quadratic utility and there are strategic complementarities or substitutabilities between the actions

¹⁵The games considered in both Bramoullé and Kranton (2007) and Bramoullé et al. (2014) are not strictly linear models, since there are corner solutions at zero.

of a node and those of its neighbours. A model of this type has been studied by among others, Ballester et al. (2006) and Bramoullé et al. (2014). The former focus on the case of strategic complementarities; while the latter study the case with strategic substitutabilities and characterise all equilibria of this game. In this model, the utility function for agent i in network g takes the following form:

$$U_i(y_{i,g}; \mathbf{y}_{-i,g}, \mathbf{X}_g, \mathbf{G}_g) = \left(\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g}) - \frac{1}{2} \left(y_{i,g} - 2\beta \sum_{j=1}^{N_g} G_{ij,g} y_{j,g} \right) \right) y_{i,g} \quad (4.2)$$

where $y_{i,g}$ is i 's action or choice, and $\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g})$ is, as before, an individual heterogeneity parameter.¹⁶ This has the same form as Equation 3.2, but with \mathbf{G} replacing $\tilde{\mathbf{G}}$. $\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g})$ is again typically parameterised as:

$$\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g}) = \mathbf{x}_{i,g} \boldsymbol{\delta} + \sum_{j=1}^n \tilde{G}_{ij,g} \mathbf{x}_{j,g} \boldsymbol{\gamma} + \mathbf{z}_g \boldsymbol{\eta} + \nu_g + \varepsilon_{i,g}$$

so that individual heterogeneity is a function of a node's own characteristics, the *average* characteristics of its neighbours, network-level observed characteristics, and some unobserved network- and individual-level terms. This model shares some features with the model of Blume et al. (2015), as different network matrices are used to capture the effects of neighbours' outcomes and characteristics, which helps to ease identification.

The quadratic cost of own actions means that in the absence of any network, there would be a unique optimal amount of effort the node would exert, as in the local average model. $\beta > 0$ implies that neighbours' actions are complementary to a node's own actions, so that the node increases his actions in response to those of his neighbours. If $\beta < 0$, then nodes' actions are substitutes, and the reverse is true. Nodes choose $y_{i,g}$ so as to maximise their utility.

The best response function is:

$$y_{i,g}^*(\mathbf{G}_g) = \beta \sum_{j=1}^n G_{ij,g} y_{j,g} + \mathbf{x}_{i,g} \boldsymbol{\delta} + \sum_{j=1}^n \tilde{G}_{ij,g} \mathbf{x}_{j,g} \boldsymbol{\gamma} + \mathbf{z}_g \boldsymbol{\eta} + \nu_g + \varepsilon_{i,g} \quad (4.3)$$

Ballester et al. (2006) solve for the Nash equilibrium of this game when $\beta > 0$ and show that when $|\beta \omega_{max}(\mathbf{G}_g)| < 1$, where $\omega_{max}(\mathbf{G}_g)$ is the largest eigenvalue of the matrix \mathbf{G}_g , the equilibrium is unique and the equilibrium outcome relates to a node's Katz-Bonacich centrality, which is defined as $\mathbf{b}(\mathbf{G}_g, \beta) = (\mathbf{I}_g - \beta \mathbf{G}_g)^{-1} \boldsymbol{\iota}_g$.^{17,18}

¹⁶Notice that $\sum_{j=1}^{N_g} G_{ij,g} y_{j,g} = \mathbf{G}_{i,g} \mathbf{y}_g$.

¹⁷A more general definition for Katz-Bonacich centrality is $\mathbf{b}(\mathbf{G}_g, \beta, a) = (\mathbf{I}_g - \beta \mathbf{G}_g)^{-1} (a \mathbf{G}_g \boldsymbol{\iota}_g)$, where $a > 0$ is a constant (Jackson, 2008).

¹⁸The condition that $|\beta \omega_{max}(\mathbf{G}_g)| < 1$ ensures that positive feedback loops set off by these complementarities don't diverge without bound. The extent to which it binds will depend on the underlying network topology: the largest

Bramoullé et al. (2014) study the game with strategic substitutabilities between the action of a node and those of its neighbours. In this case, when one agent chooses one action, his neighbours would choose the opposite action, inducing their neighbours to choose a similar action to the first agent, and so on. In equilibrium, the agent’s choice is influenced by the direct choice of his neighbours, as well as by the aggregated sum of these opposing choices across different network paths. Bramoullé et al. (2014) characterise the set of Nash equilibria of the game and show that they depend on the lowest eigenvalue of the network adjacency matrix, $\omega_{min}(\mathbf{G}_g)$. This eigenvalue is negative and relates to the aggregated effect of agents’ choices on others. When it is large in magnitude, the opposing forces outlined above could go in many different directions, leading to multiple equilibria. When it is sufficiently small, *i.e.* $\beta|\omega_{min}(\mathbf{G}_g)| < 1$, a unique equilibrium exists since the opposing forces move in the same direction to converge to a single point. When multiple equilibria are possible, they must be accounted for in any empirical analysis. Methods developed in the literature on the econometrics of games may be applied here (Bisin et al., 2011). See de Paula (2013) for an overview. Below, we focus on conditions under which the social effect parameter is identified when a unique equilibrium exists. This is done to keep the exposition simple.

When a unique equilibrium exists, this theoretical set-up implies the following empirical model (stacking data from multiple networks):

$$\mathbf{Y} = \alpha\mathbf{I} + \beta\mathbf{G}\mathbf{Y} + \mathbf{X}\boldsymbol{\gamma} + \tilde{\mathbf{G}}\mathbf{X}\boldsymbol{\delta} + \mathbf{Z}\boldsymbol{\eta} + \mathbf{L}\boldsymbol{\nu} + \boldsymbol{\varepsilon} \quad (4.4)$$

where variables and parameters are as defined above .

4.2 Identification

Identification of Equation 4.4 using observational data has been studied by Calvó-Armengol et al. (2009), Lee and Liu (2010) and Liu et al. (2014b). They proceed under the assumption that $\mathbb{E}[\boldsymbol{\varepsilon}|\mathbf{X}, \mathbf{Z}, \mathbf{G}, \tilde{\mathbf{G}}] = 0$ and $\mathbb{E}[\boldsymbol{\nu}|\mathbf{X}, \mathbf{Z}, \mathbf{G}, \tilde{\mathbf{G}}] \neq 0$. That is, the node-varying error component is conditionally independent of node- and network-level observables and of the network, while the network-level unobservable could be correlated with node- and network-level characteristics and/or the network itself.

These assumptions imply a two-stage network formation process. First agents select into a network based on a set of observed individual- and network-level characteristics and some common network-level unobservables. Then in a second stage they form links with other nodes. There are no network-level unobservable factors that determine link formation once the network has been selected by the node. Moreover, there are no node-level unobservable factors that determine the choice of network or link formation within the chosen network.

eigenvalue increases with additional links, so that in larger and denser networks, the largest possible value of β for which the condition holds will be smaller.

To proceed, we assume that data are available for multiple networks. Then, as in Section 3, we replace the network-level observables, \mathbf{Z} , and the network-level unobservable, $\mathbf{L}\boldsymbol{\nu}$ in Equation 4.4 with network fixed effects, $\mathbf{L}\tilde{\boldsymbol{\nu}}$, where $\tilde{\boldsymbol{\nu}}$ is a $M \times 1$ vector that captures the network fixed effects.

To account for the fixed effect, a within-transformation is applied, as in Section 3. This transformation is represented by the block diagonal matrix \mathbf{J} that stacks the following network-level transformation matrices – $\mathbf{J}_g = \mathbf{I}_g - \frac{1}{N_g}(\boldsymbol{\nu}_g\boldsymbol{\nu}_g')$ – along the leading diagonal, with off-diagonal terms set to 0. The resulting model, analogous to Equation 3.6, is:

$$\mathbf{JY} = \beta\mathbf{JGY} + \mathbf{JX}\boldsymbol{\gamma} + \mathbf{J}\tilde{\mathbf{G}}\mathbf{X}\boldsymbol{\delta} + \mathbf{J}\boldsymbol{\varepsilon} \quad (4.5)$$

The model above suffers from the reflection problem, since \mathbf{Y} appears on both sides of the equation. However, the parameters of Equation 4.5 can be identified using linear instrumental variables (IV) if the deterministic part of the right hand side, $[\mathbb{E}(\mathbf{JGY}), \mathbf{JX}, \mathbf{J}\tilde{\mathbf{G}}\mathbf{X}]$, has full column rank. To see the conditions under which this is satisfied, we examine the term with the endogenous variable, $\mathbb{E}(\mathbf{JGY})$. Under the assumption that $|\beta\omega_{max}(\mathbf{G}_g)| < 1$, we obtain the following from the reduced form equation of Equation 4.4:

$$\begin{aligned} \mathbb{E}(\mathbf{JGY}) = & \mathbf{J}(\mathbf{GX} + \beta\mathbf{G}^2\mathbf{X} + \dots)\boldsymbol{\gamma} + \mathbf{J}(\mathbf{G}\tilde{\mathbf{G}}\mathbf{X} + \beta\mathbf{G}^2\tilde{\mathbf{G}}\mathbf{X} + \dots)\boldsymbol{\delta} \\ & + \mathbf{J}(\mathbf{GL} + \beta\mathbf{G}^2\mathbf{L} + \dots)\tilde{\boldsymbol{\nu}} \end{aligned} \quad (4.6)$$

We can thus see that if the matrices $\{\mathbf{I}, \mathbf{G}, \tilde{\mathbf{G}}, \mathbf{G}\tilde{\mathbf{G}}\}$ are linearly independent, and $\boldsymbol{\gamma}$, $\boldsymbol{\delta}$, and $\tilde{\boldsymbol{\nu}}$ each have some non-zero terms, the parameters of Equation 4.4 are identified.¹⁹ Node degree (\mathbf{GL}), along with the sum of the exogenous characteristics of the node's direct neighbours (\mathbf{GX}), and sum of the average exogenous characteristics of its second-degree neighbours ($\mathbf{G}\tilde{\mathbf{G}}\mathbf{X}$) can be used as instruments for the total outcome of the node's neighbours (\mathbf{GY}). That node degree can be used as an instrument follows intuitively from the theoretical model: when there are dyad-level strategic complementarities, an individual's own outcome will be a weakly increasing function of the number of his direct neighbours. Moreover, the availability of node degree as an instrument can allow one to identify parameters without using the exogenous characteristics, \mathbf{X} , of second- or higher-degree network neighbours, which can be particularly advantageous when only sampled data is available, as shown by Liu (2013).

In terms of practical application, consider using this method to identify whether there are complementarities between the schooling performance of an individual and that of his friends, conditional


¹⁹See Liu et al. (2014b) for a different identification condition that allows for some linear dependence among these matrices under additional restrictions.

on his own characteristics (age, gender, and parental income), the composition of his friends (average age, gender, and parental income), and some school characteristics. Then, if there are individuals in the same network with different numbers of friends, and the matrices $\{\mathbf{I}, \mathbf{G}, \tilde{\mathbf{G}}, \mathbf{G}\tilde{\mathbf{G}}\}$ are linearly independent, the individual’s degree, along with the total characteristics of his friends (*i.e.* total age, gender, and parental income) and the sum of the average age, gender, and parental income of the individual’s friends of friends can be used as instruments for the sum of the individual’s friends’ schooling performance. Similar to the local average model, this strategy relies on a researcher being able to define reasonably well direct, as well as indirect neighbours, which might be less clear in some contexts.

Parameters can still be identified if there is no variation in node degree within a network for all networks in the data, but there is variation in degree across networks. In this case, $\mathbf{G}_g = \bar{d}_g \tilde{\mathbf{G}}_g$ and $[E(\mathbf{JGY}), \mathbf{JX}, \mathbf{J}\tilde{\mathbf{G}}\mathbf{X}]$ has full column rank if the matrices $\{\mathbf{I}, \mathbf{G}, \tilde{\mathbf{G}}, \mathbf{G}\tilde{\mathbf{G}}, \tilde{\mathbf{G}}^2, \mathbf{G}\tilde{\mathbf{G}}^2\}$ are linearly independent and γ and δ each have non-zero terms.²⁰ This requires the presence of some pair of nodes i and k , who are only indirectly connected. Finally, when there is no variation in node degree within and across all networks in the data, parameters can be identified using a similar condition as encountered in Section 4 above: the matrices $\{\mathbf{I}, \tilde{\mathbf{G}}, \tilde{\mathbf{G}}^2, \tilde{\mathbf{G}}^3\}$ should be linearly independent.

It is possible to identify model parameters in the local aggregate model in networks where the local average model parameters cannot be identified. For example, in a star network (see Figure 3) there is no pair of agents that has a geodesic distance (*i.e.* shortest path) of 3 or more, so this fails the identification condition for the local average model with fixed effects. However, there is variation in node degree within the network and the matrices $\mathbf{I}_g, \mathbf{G}_g, \tilde{\mathbf{G}}_g, \mathbf{G}_g\tilde{\mathbf{G}}_g$ can be shown to be linearly independent, thus satisfying the identification conditions for the local aggregate model.

Figure 3: Star Network



²⁰With additional restrictions, identification can still be achieved when there is some linear dependence in these matrices. See Liu et al. (2014b) for details.

5 Hybrid Local Models

The local average and local aggregate models embody distinct mechanisms through which social effects arise. One may be interested in jointly testing these mechanisms, and empirically identifying the most relevant one for a particular context. Liu et al. (2014a) present a framework nesting both the local aggregate and local average models, allowing for this.

5.1 Set Up and Theoretical Foundations

The utility function for node i in network g that nests both the (linear) local aggregate and local average models has the following form:

$$U_i(y_{i,g}; \mathbf{y}_{-i,g}, \mathbf{X}_g, \tilde{\mathbf{G}}_{i,g}, \mathbf{G}_{i,g}) = \left(\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g}) + \beta_1 \sum_{j=1}^{N_g} G_{ij,g} y_{j,g} - \frac{1}{2} \left(y_{i,g} - 2\beta_2 \sum_{j=1}^{N_g} \tilde{G}_{ij,g} y_{j,g} \right) \right) y_{i,g} \quad (5.1)$$

where $\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g})$ is node-specific observed heterogeneity, which affects the node's marginal return from the chosen outcome level $y_{i,g}$. A node's utility is thus affected by the choices of its neighbours through changing the marginal returns of its own choice (*e.g.* in a schooling context, an individual's studying effort is more productive if his friends also study), as in the local aggregate model, and by a cost of deviating from the average choice of its neighbours (*i.e.* individuals face a utility cost if they study when their friends don't study), as in the local average model.

The best reply function for a node i nests both the local average and local aggregate terms. Liu et al. (2014a) prove that under the condition that $\beta_1 \geq 0$, $\beta_2 \geq 0$ and $d_g^{max} \beta_1 + \beta_2 < 1$, where d_g^{max} is the largest degree in network g , the simultaneous move game has a unique interior Nash equilibrium in pure strategies. Note that this rules out the possibility of individuals' actions being strategic substitutes (*i.e.* $\beta < 0$), as for example if one student in a group needs to supply effort on homework so that the others can copy him.

The econometric model, assuming that the node-specific observed heterogeneity parameter takes the form $\pi_{i,g}(\mathbf{X}_g, \tilde{\mathbf{G}}_{i,g}) = \mathbf{x}_{i,g} \boldsymbol{\gamma} + \sum_{j=1}^{N_g} \tilde{G}_{ij,g} \mathbf{x}_{j,g} \boldsymbol{\delta} + \mathbf{z}_g \boldsymbol{\eta}_g + \nu_g + \varepsilon_{i,g}$, is as follows:

$$\mathbf{Y} = \alpha \mathbf{1} + \beta_1 \mathbf{G} \mathbf{Y} + \beta_2 \tilde{\mathbf{G}} \mathbf{Y} + \mathbf{X} \boldsymbol{\gamma} + \tilde{\mathbf{G}} \mathbf{X} \boldsymbol{\delta} + \mathbf{Z} \boldsymbol{\eta} + \mathbf{L} \boldsymbol{\nu} + \boldsymbol{\varepsilon} \quad (5.2)$$

using the same notation as before.

With data from only a single network it is not possible to separately identify β_1 and β_2 and hence test between the local aggregate and local average models (or indeed find that the truth is a hybrid of the two effects). Identification of parameters is considered when data from multiple networks are available under the assumption that $\mathbb{E}[\varepsilon_{i,g} | \mathbf{X}_g, \mathbf{Z}_g, \mathbf{G}_g, \tilde{\mathbf{G}}_g] = 0$ and $\mathbb{E}[\nu_g | \mathbf{X}_g, \mathbf{Z}_g, \mathbf{G}_g, \tilde{\mathbf{G}}_g] \neq 0$.

Thus, as in Sections 3 and 4 above, the individual error term, $\varepsilon_{i,g}$ is assumed to be mean independent of node- and network-level observable characteristics and the network. The network-level unobservable, ν_g , by contrast is allowed to be correlated with node- and network-level characteristics and/or the network.

5.2 Identification

To proceed, as in the local average and local aggregate model, $\mathbf{Z}\eta$ and $\mathbf{L}\nu$ are replaced by a network-level fixed effect, $\mathbf{L}\tilde{\nu}$, which is then removed using the within-transformation, \mathbf{J} . The resulting transformed network model is:

$$\mathbf{J}\mathbf{Y} = \beta_1\mathbf{J}\mathbf{G}\mathbf{Y} + \beta_2\mathbf{J}\tilde{\mathbf{G}}\mathbf{Y} + \mathbf{J}\mathbf{X}\gamma + \mathbf{J}\tilde{\mathbf{G}}\mathbf{X}\delta + \mathbf{J}\varepsilon \quad (5.3)$$

When there is variation in the degree within a network g , then the reduced form equation of Equation 5.3 implies that $\mathbf{J}\mathbf{G}(\mathbf{I} - \beta_1\mathbf{G} - \beta_2\tilde{\mathbf{G}})^{-1}\mathbf{L}$ can be used as an instrument for the local aggregate term $\mathbf{J}\mathbf{G}\mathbf{Y}$ and $\mathbf{J}\tilde{\mathbf{G}}(\mathbf{I} - \beta_1\mathbf{G} - \beta_2\tilde{\mathbf{G}})^{-1}\mathbf{L}$ can be used as an instrument for the local average term $\mathbf{J}\tilde{\mathbf{G}}\mathbf{Y}$. The model parameters may in principle thus be identified even if there are no node-level exogenous characteristics, \mathbf{X} , in the model, as long as $\beta_1 \neq 0$. However, if $\beta_1 = 0$, the model excluding exogenous characteristics, \mathbf{X} , is tautological – in this case one is simply regressing individuals' outcomes on the mean of the outcomes (see Angrist, 2014 for further discussion). The availability of such characteristics offers more possible IVs: in particular, the total and average exogenous characteristics of direct and indirect neighbours can be used as instruments. These are necessary for identification when all nodes within a network have the same degree, though average degree may vary across networks. In this case, parameters can be identified if the matrices $\{\mathbf{I}, \mathbf{G}, \tilde{\mathbf{G}}, \mathbf{G}\tilde{\mathbf{G}}, \tilde{\mathbf{G}}^2, \mathbf{G}\tilde{\mathbf{G}}^2, \tilde{\mathbf{G}}^3\}$ are linearly independent.²¹ If, however, all nodes in all networks have the same degree, it is not possible to identify separately the parameters β_1 and β_2 .

This specification nests both the local average and local aggregate models, so a J-test for non-nested regression models can be applied to uncover the relevance of each mechanism. The intuition underlying the J-test is as follows: if a model is correctly specified (in terms of the set of regressors), then the fitted value of an alternative model should have no additional explanatory power in the original model, *i.e.* its coefficient should not be significantly different from zero. Thus, to identify which of the local average or local aggregate mechanisms is more relevant for a specific outcome, one could first estimate one of the models (*e.g.* the local average model), and obtain the predicted outcome value under this mechanism. In a second step, estimate the other model (in our example, the local aggregate model), and include as a regressor the predicted value from the other (*i.e.* local average) model. If the mechanism underlying the local average model is also relevant for the outcome, the coefficient on the predicted value will be statistically different from 0. The converse

²¹As with the local average and local aggregate model, this identification assumption imposes that there are pairs of individuals who influence the other (or each other) indirectly only, which might be very strong in some contexts.

can also be done to test the relevance of the second model (the local aggregate model in our case). See Liu et al. (2014a) for more details.

6 Models with Network Characteristics

The models considered thus far allow for a node’s outcomes to be influenced only by outcomes of its neighbours, so-called ‘local’ models. However, the broader network structure may affect node- and aggregate network- outcomes through more general functionals or features of the network. Depending on the theoretical model used, there are different predictions on which network features relate to different outcomes of interest. For example, the DeGroot (1974) model of social learning implies that a node’s eigenvector centrality, which measures its ‘importance’ in the network by how important its neighbours are, determines how influential it is in affecting the beliefs of other nodes.

Empirical testing and verification of the predictions of these theoretical models has greatly lagged the theoretical literature due to a lack of datasets with both information on network structure and socio-economic outcomes of interest. The recent availability of detailed network data from many contexts has begun to relax this constraint.

The following types of specification are typically estimated when assessing how outcomes vary with network structure, for node-level outcomes:

$$\mathbf{Y} = \mathbf{f}_y(\mathbf{w}_y(\mathbf{G}, \mathbf{Y}), \mathbf{X}, \mathbf{w}_x(\mathbf{G}, \mathbf{X}), \mathbf{Z}) + \varepsilon \quad (6.1)$$

and network-level outcomes:

$$\bar{\mathbf{Y}} = \mathbf{f}_{\bar{y}}(\bar{\mathbf{w}}_{\bar{y}}(\mathbf{G}), \bar{\mathbf{X}}, \bar{\mathbf{w}}_{\bar{x}}(\mathbf{G}, \bar{\mathbf{X}})) + \mathbf{u} \quad (6.2)$$

$\mathbf{f}_y(\cdot)$ and $\mathbf{f}_{\bar{y}}(\cdot)$ are functions that specify the shape of the relationship between the network statistics and the node- and network-level outcomes. Though, in principle, the shape of $\mathbf{f}_y(\cdot)$ should be guided by theory (where possible), through the rest of this Section, we take $\mathbf{f}_y(\cdot)$ to be a linear index in its argument, as is common in the literature. $\mathbf{w}_y(\mathbf{G}, \mathbf{Y})$ is an $(\sum_{g=1}^M N_g \times R)$ matrix stacking the $(1 \times R)$ node-level vector of network statistics that vary at the node- or network-level and that may be interacted with \mathbf{Y} .²² $\bar{\mathbf{w}}_{\bar{y}}(\mathbf{G})$ is an $(M \times \bar{R})$ matrix containing the \bar{R} network statistics in the network-level regression. \mathbf{X} is a matrix of observable characteristics of nodes, $\mathbf{w}_x(\mathbf{G}, \mathbf{X})$ interacts network statistics with exogenous characteristics of nodes, and \mathbf{Z} and $\bar{\mathbf{X}}$ are network-level observable characteristics. $\bar{\mathbf{w}}_{\bar{x}}(\mathbf{G}, \bar{\mathbf{X}})$ interacts network statistics with network-level observable characteristics.

²²The term $\mathbf{w}_y(\mathbf{G}, \mathbf{Y})$ will be endogenous when network statistics are interacted with \mathbf{Y} .

The complexity of networks poses an important challenge in understanding how outcomes vary with network structure. In particular, there are no sufficient statistics that fully describe the structure of a network. For example, networks with the same average degree may vary greatly on dimensions such as density, clustering and average path length among others. Moreover, the adjacency matrix, \mathbf{G} , which describes fully the structure of a network, is too high-dimensional an object to include directly in tests of the influence of broader features of network structure. Theory can provide guidance on which statistics are likely to be relevant, and also on the shape of the relationship between the network statistic and the outcome of interest. A limitation though is that theoretical results may not be available (given currently known techniques) for outcomes one is interested in studying. This is a challenge faced by, for instance Alatas et al. (2016) who study how network structure affects information aggregation.

Below we outline methods that have been applied to analyse the effects of features of network structure on socio-economic outcomes. We do so separately for node-level specifications and network-level specifications. This literature is very much in its infancy and few methods have been developed to allow for identification of causal parameters.

6.1 Node-Level Specifications

Many theoretical models predict how node-level outcomes vary with the ‘position’ of a node in the network, captured by node varying network statistics such as centrality; or with features of the node’s local neighbourhood such as node clustering; or with the ‘connectivity’ of the network, represented by statistics that vary at the network-level such as network density.

A common type of empirical specification used in the literature correlates network statistics with some relevant socio-economic outcome of interest. This approach is taken by, for example, Jackson et al. (2012) who test whether informal favours take place across edges that are supported (*i.e.* that nodes exchanging a favour have a common neighbour), which is the prediction of their theoretical model.

This corresponds with $\mathbf{w}_y(\mathbf{G}, \mathbf{Y})$ in Equation 6.1 above being defined as $\mathbf{w}_y(\mathbf{G}, \mathbf{Y}) = \boldsymbol{\omega}$, where $\boldsymbol{\omega}$ is the matrix of network statistics of interest, and $\mathbf{w}_x(\mathbf{G}, \mathbf{X})$ being defined as $\boldsymbol{\nu}$. Here, $\mathbf{w}_y(\mathbf{G}, \mathbf{Y})$ is defined to be a function of the network only.

When $\mathbf{f}_y(\cdot)$ is linear, the specification is as follows:

$$\mathbf{Y} = \alpha\boldsymbol{\nu} + \boldsymbol{\omega}\boldsymbol{\beta} + \mathbf{X}\boldsymbol{\gamma} + \mathbf{Z}\boldsymbol{\eta} + \boldsymbol{\varepsilon} \tag{6.3}$$

where the variables and parameters are as defined above and the parameter of interest is $\boldsymbol{\beta}$. Defining $\mathbf{W} = (\boldsymbol{\omega}, \mathbf{X}, \mathbf{Z})$, the key identification assumption is that $E[\boldsymbol{\varepsilon}'\mathbf{W}] = 0$, that is that the right hand side terms are uncorrelated with the error term. This may not be satisfied if there are unobserved factors that affect both the network statistic (through affecting network formation decisions) and

the outcome, \mathbf{Y} or if the network statistic is mismeasured. Both of these are important concerns that we cover in detail in Advani and Malde (2014).

In some cases, one may also be interested in estimating a model where an agent's outcome is affected by the outcomes of his neighbours, weighted by a measure of their network position. For example, in the context of learning about a new product or technology, the DeGroot (1974) model of social learning implies that nodes' eigenvector centrality determines how influential they are in influencing others' behaviour. Thus, conditional on the node's eigenvector centrality, its choices may be influenced more by the choices of his neighbours with high eigenvector centrality. Thus, one may want to weight the influence of neighbours' outcomes on own outcomes by their eigenvector centrality, conditional on own eigenvector centrality. If we model this linearly, it implies a model of the following form:

$$\mathbf{Y} = \alpha\mathbf{1} + \mathbf{w}_y(\mathbf{G}, \mathbf{Y})\boldsymbol{\beta} + \tilde{\mathbf{X}}\tilde{\boldsymbol{\gamma}} + \mathbf{w}_x(\mathbf{G}, \tilde{\mathbf{X}})\tilde{\boldsymbol{\delta}} + \mathbf{Z}\boldsymbol{\eta} + \mathbf{L}\boldsymbol{\nu} + \boldsymbol{\varepsilon} \quad (6.4)$$

$\mathbf{w}_y(\mathbf{G}, \mathbf{Y})$ is an $\sum_g N_g \times R$ matrix, with the $(i, r)^{th}$ element being the weighted sum of i 's neighbours' outcomes, $\sum_{j \neq i} G_{ij,g} y_{j,g} \omega_{j,g}^r$ or $\sum_{j \neq i} \tilde{G}_{ij,g} y_{j,g} \omega_{j,g}^r$, with weights $\omega_{j,g}^r$ being the neighbour's r^{th} network statistic (*e.g.* the neighbour's eigenvector centrality in the DeGroot's model of social learning). $\tilde{\mathbf{X}} = (\tilde{\mathbf{X}}'_1, \tilde{\mathbf{X}}'_2, \dots, \tilde{\mathbf{X}}'_M)'$, where $\tilde{\mathbf{X}}_g = (\mathbf{X}_g, \boldsymbol{\omega}_g)$ is a matrix stacking together the network-level matrices of exogenous explanatory variables and (own) network statistics of interest. $\mathbf{w}_x(\mathbf{G}, \tilde{\mathbf{X}})$ could be defined as $\mathbf{G}\tilde{\mathbf{X}}$ or $\tilde{\mathbf{G}}\tilde{\mathbf{X}}$. Note that this formulation allows for the influence of neighbours' background characteristics on the outcome to be weighted by the values of their network statistics. Identification of parameters in this case is complicated by the fact that $\mathbf{w}_y(\mathbf{G}, \mathbf{Y})$ is a (possibly non-linear) function of \mathbf{Y} , and thus endogenous. It may be possible to achieve identification using network-based instrumental variables, as above, though it is not immediately obvious how such an IV could be constructed. Further research is needed to shed light on these issues.

6.2 Network-level Specifications

Aggregate network-level outcomes, such as the degree of risk sharing or the aggregate penetration of a new product, may also be affected by how 'connected' the network is, or the 'position' of nodes that experience a shock or who first hear about a new product.

Empirical tests of the relationship between aggregate network-level outcomes and network statistics involve estimating specifications such as Equation 6.2. The shape of the function $\mathbf{f}_{\bar{\mathbf{y}}}(\cdot)$ and the choice of statistics in $\bar{\mathbf{w}}_{\bar{\mathbf{y}}}(\mathbf{G}) = \bar{\boldsymbol{\omega}}$, where $\bar{\boldsymbol{\omega}}$ is an $(M \times \bar{R})$ matrix of network statistics, are again, ideally, motivated by theory. With linear $\mathbf{f}_{\bar{\mathbf{y}}}(\cdot)$, this implies the following equation:

$$\bar{\mathbf{Y}} = \phi_0 + \bar{\boldsymbol{\omega}}\boldsymbol{\phi}_1 + \tilde{\mathbf{X}}\boldsymbol{\phi}_2 + \bar{\mathbf{w}}_{\bar{\mathbf{x}}}(\mathbf{G}, \tilde{\mathbf{X}})\boldsymbol{\phi}_3 + \mathbf{u} \quad (6.5)$$

where the variables are as defined after Equation 6.2. The parameter of interest is typically ϕ_1 . Defining $\bar{\mathbf{W}} = (\boldsymbol{\omega}, \bar{\mathbf{X}}, \bar{\mathbf{w}}_{\bar{\mathbf{x}}}(\mathbf{G}, \bar{\mathbf{X}}))$, the key identification assumption is that $E[\mathbf{u}\bar{\mathbf{W}}] = 0$, which will not hold if there are unobserved variables in \mathbf{u} that affect both the formation of the network and the outcome $\bar{\mathbf{y}}$; or if the network statistics are mismeasured. Recent empirical work, such as that by Banerjee et al. (2013), has used quasi-experimental variation to try and alleviate some of the challenges posed by the former issue in identifying the parameter ϕ_1 .

Since this specification uses data at the network-level, estimation will require a large sample of networks in order to recover precise estimates of the parameters, even in the absence of endogeneity from network formation and mismeasurement of the network. This is a problem in practice, since networks data is difficult and costly to collect, meaning that in practice researchers have data for a small number of networks only.

7 Experimental Variation

Thus far, we have considered the identification of the social effect parameters using observational data. We now consider identification of these parameters using experimental data. We focus on the case where a policy is assigned randomly to a sub-set of nodes in a network. Throughout we assume that the network is pre-determined and unchanged by the exogenously assigned policy.²³

We focus the discussion on identifying parameters of the local average model specified in Section 3 above, studied by Dieye et al. (2014). The issues related to using experimental variation to uncover the parameters of the local aggregate model are similar. As outlined above, this model implies that a node’s outcome is affected by the average outcome of its network neighbours, its own and network-level exogenous characteristics (the latter may be subsumed into a network fixed effect), and the average characteristics of its network neighbours. We are typically interested in parameters β , γ and δ in the following equation:

$$\mathbf{Y} = \alpha\mathbf{1} + \beta\tilde{\mathbf{G}}\mathbf{Y} + \mathbf{X}\gamma + \tilde{\mathbf{G}}\mathbf{X}\delta + \mathbf{L}\tilde{\nu} + \boldsymbol{\varepsilon} \quad (7.1)$$

where the variables are as defined previously.

Throughout this section, we assume that the policy shifts outcomes for the nodes that directly receive the policy.²⁴ To proceed further, we first assume that a node that does not receive the policy (*i.e.* is untreated, to use the terminology from the policy evaluation literature), is only affected by the policy through its effects on the outcomes of the node’s network neighbours. This implies the following model for the outcome \mathbf{Y} :

²³This assumption is not innocuous. Comola and Prina (2014) provide an example where the policy intervention does change the network.

²⁴Below, we will consider identification conditions in the case where a node may be affected by the treatment status of his network neighbours even if their outcomes do not shift in response to the treatment.

$$Y = \alpha\iota + \beta\tilde{\mathbf{G}}\mathbf{Y} + \mathbf{X}\gamma + \tilde{\mathbf{G}}\mathbf{X}\delta + \rho\mathbf{t} + \mathbf{L}\tilde{\nu} + \varepsilon \quad (7.2)$$

where \mathbf{t} is the treatment vector, and ρ is the direct effect of treatment. We assume that $\mathbb{E}[\varepsilon|\mathbf{X}, \mathbf{Z}, \tilde{\mathbf{G}}, \mathbf{t}] = 0$. Moreover, random allocation of the treatment implies that $\mathbf{t} \perp\!\!\!\perp \mathbf{X}, \mathbf{Z}, \tilde{\mathbf{G}}, \varepsilon$.

Applying the same within-transformation as in Section 3 above to account for the network-level fixed effect leads to the following specification:

$$\mathbf{J}\mathbf{Y} = \alpha\mathbf{J}\iota + \beta\mathbf{J}\tilde{\mathbf{G}}\mathbf{Y} + \mathbf{J}\mathbf{X}\gamma + \mathbf{J}\tilde{\mathbf{G}}\mathbf{X}\delta + \rho\mathbf{J}\mathbf{t} + \mathbf{J}\varepsilon \quad (7.3)$$

We can use instrumental variables to identify β as long as the deterministic part of the right hand side of Equation 7.3, $[\mathbb{E}(\mathbf{J}\tilde{\mathbf{G}}\mathbf{Y}), \mathbf{J}\mathbf{X}, \mathbf{J}\tilde{\mathbf{G}}\mathbf{X}, \mathbf{J}\mathbf{t}]$ has full column rank. $\mathbf{J}\mathbf{X}$ and $\mathbf{J}\tilde{\mathbf{G}}\mathbf{X}$ can be used as instruments for themselves. We thus need an instrument for $\mathbb{E}[\mathbf{J}\tilde{\mathbf{G}}\mathbf{Y}]$. We use the following expression for $\mathbf{J}\tilde{\mathbf{G}}\mathbf{Y}$, derived from the reduced form of Equation 7.2 under the assumption that $|\beta| < 1$, to construct instruments:

$$\begin{aligned} \mathbb{E}[\mathbf{J}\tilde{\mathbf{G}}\mathbf{Y}] = & \mathbf{J}\tilde{\mathbf{G}} \sum_{s=0}^{\infty} \beta^s \tilde{\mathbf{G}}^s \alpha\iota + \mathbf{J}(\tilde{\mathbf{G}}\mathbf{X}\gamma + \beta\tilde{\mathbf{G}}^2\mathbf{X}\gamma + \dots) + \mathbf{J}(\tilde{\mathbf{G}}^2\mathbf{X}\delta + \beta\tilde{\mathbf{G}}^3\mathbf{X}\delta + \dots) \\ & + \mathbf{J}(\rho\tilde{\mathbf{G}}\mathbf{t} + \beta\rho\tilde{\mathbf{G}}^2\mathbf{t} + \dots) \end{aligned} \quad (7.4)$$

From this equation, we can see that $\tilde{\mathbf{G}}\mathbf{t}$, the average treatment status of a node's network neighbours, appears in the reduced form for $\mathbb{E}[\mathbf{J}\tilde{\mathbf{G}}\mathbf{Y}]$. However, it does not appear in Equation 7.2. It can thus be used as an instrument for $\tilde{\mathbf{G}}\mathbf{Y}$, either in addition to, or as an alternative to $\tilde{\mathbf{G}}^2\mathbf{X}$ and $\tilde{\mathbf{G}}^3\mathbf{X}$, the average characteristics of the node's second- and third-degree neighbours. Thus, the policy could be used to identify the model parameters, albeit under a strong assumption on the mechanism by which it has any effect (see below).²⁵

An advantage to using $\tilde{\mathbf{G}}\mathbf{t}$ as an instrument for the endogenous $\tilde{\mathbf{G}}\mathbf{Y}$ is that since the treatment is randomly assigned, and it only directly affects the treated individual's outcome, making it more plausible than $\tilde{\mathbf{G}}^2\mathbf{X}$ and $\tilde{\mathbf{G}}^3\mathbf{X}$, in which \mathbf{X} is not likely to be randomly assigned. A second advantage of this instrument relative to $\tilde{\mathbf{G}}^2\mathbf{X}$ and $\tilde{\mathbf{G}}^3\mathbf{X}$ is that it only requires knowledge of agents' direct neighbours. The identification conditions outlined previously in Sections 3 - 5 hinge on the fact that within the network of interest, some agents are only indirectly connected, which might be too strong in some contexts. Identification here requires sufficient variation in the proportion of an agents' direct neighbours that are treated for it to be a powerful instrument, which imposes fewer restrictions on network structure (e.g. there should be variation in degree within the network). As a result, the social effect parameter can be identified in a wider range of network structures.

²⁵Similar results can be shown for the local aggregate model when $|\beta\omega_{max}(\mathbf{G})| < 1$. However, as shown above, node degree can also be used as an additional instrument in this model.

In many cases, however, the assumption that the policy affects a node’s outcome only if it is directly treated may be too strong. The treatment status of a node’s neighbours could affect its outcome even when the neighbours’ outcomes do not shift in response to receiving the policy. An example of such a case, studied by Banerjee et al. (2013), is when the treatment involves providing individuals with information on a new product, and the outcome of interest is the take-up of the product. Then neighbours’ treatment status could affect the individual’s own adoption decision by (1) shifting his neighbours’ decision (endorsement effects), and also (2) through neighbours passing on information about the product and letting the individual know of its existence (diffusion effect).²⁶ In this case, which is studied by Dieye et al. (2014), a more appropriate model would be as follows:

$$\mathbf{Y} = \alpha\mathbf{1} + \beta\tilde{\mathbf{G}}\mathbf{Y} + \mathbf{X}\boldsymbol{\gamma} + \tilde{\mathbf{G}}\mathbf{X}\boldsymbol{\delta} + \rho\mathbf{t} + \tilde{\mathbf{G}}\mathbf{t}\mu + \boldsymbol{\varepsilon} \quad (7.5)$$

where ρ captures the direct treatment effect, *i.e.* the effect of a node itself being treated, and μ is the direct effect of the average treatment status of social contacts. This highlights the limits to using exogenous variation from randomised experiments to identify social effect parameters. We might want to use the exogenous variation in the average treatment allocation of a node’s neighbours, $\tilde{\mathbf{G}}\mathbf{t}$, as an instrument for neighbours’ outcomes, $\tilde{\mathbf{G}}\mathbf{Y}$. However, this will identify β only under the assumption that $\mu = 0$, *i.e.* there is no direct effect of neighbours’ treatment status. This rules out economic effects such as the diffusion effect.

We can still make use of the treatment effect for identification, by using the average treatment status of a node’s second-degree (and higher-degree) neighbours, $\tilde{\mathbf{G}}^2\mathbf{t}$, as instruments for the average outcome of his neighbours ($\tilde{\mathbf{G}}\mathbf{Y}$). This is the same identification result as discussed earlier, from Bramoullé et al. (2009), and simply treats $\tilde{\mathbf{G}}^2\mathbf{t}$ in the same way the other covariates of second-degree neighbours, $\tilde{\mathbf{G}}^2\mathbf{X}$. Such instruments rely not only on variation in treatment status, but also on the network structure, with identification not possible for certain network structures as we saw in Section 3. As before, note that instruments based on random treatment allocation and network structure (*e.g.* $\tilde{\mathbf{G}}\mathbf{t}$ and $\tilde{\mathbf{G}}^2\mathbf{t}$) will be more plausible than those based on the exogenous characteristics, \mathbf{X} , and the network structure (*e.g.* $\tilde{\mathbf{G}}^2\mathbf{X}$) since \mathbf{t} has been randomly allocated, whereas \mathbf{X} need not be.²⁷

Thus far, we have discussed how exogenous variation arising from the random assignment of a policy can be used to identify the social effect associated with a specific model – the local average model – which, as we saw, arises from an economic model where agents conform to their peers. In empirical work, though, it is common for researchers to directly include the average treatment status of network neighbours, rather than their average outcome, as a regressor in the model. In other words, the following type of specification is usually estimated:

²⁶The study of how to use these effects to maximise the number of people who adopt relates closely to study of the ‘key player’ in work by Ballester et al. (2006) and Liu et al. (2014b).

²⁷ However, it is important to remain aware that if some variables in \mathbf{X} are not exogenous (and so not suitable instruments when combined with network matrices), then their inclusion in estimation equation will already be a problem.

$$Y = b_1\iota + b_2\tilde{G}t + Xb_3 + \tilde{G}Xb_4 + b_5t + u \quad (7.6)$$

A non-zero value for b_2 is taken to indicate the presence of some social effect. However, without further modelling, it is not possible to shed light on the exact mechanism underlying this social effect, or the value of some ‘deep’ structural parameter.

When within-network experimental variation is used to identify social effects, careful attention must be paid to the important issue of inference. At one extreme, a researcher might have many (often thousands) of nodes embedded in the same network, and construct a valid comparison group using similar nodes embedded in a different part of the network to the treated node. The complication in computing standard errors comes from the fact that all nodes are embedded within the same network, and may face correlated unobserved shocks. Though these shocks may not affect identification, they will generate correlations in the outcomes of units, and must be accounted for when conducting inference. The availability of only a single network makes it very difficult to derive large sample approximations of distributions and thereby to calculate valid standard errors. Athey et al. (2015) extend the method of randomisation inference, which calculates exact p-values, to this setting. Under randomisation inference, the distribution of the test statistic is generated by considering all possible realisations of the treatment assignment, keeping the potential outcomes and characteristics of units fixed. A drawback of this procedure is that it allows for testing of sharp null hypotheses – e.g. the treatment has no effect whatsoever – only. However, we often want to test non-sharp hypotheses. Athey et al. (2015) develop methods for the computation of p-values for three specific null hypotheses.

8 Conclusion

In this paper, we provide an overview of methods to identify social effects in linear social effect models using a single cross-section of data. For a number of the most commonly used specifications of linear local models, we provide an overview of the theory models that could generate the empirical linear local specification, before outlining the conditions under which the social effect parameters of interest are identified. Thereafter, we describe what is known so far about non-local models, before considering how experimental and quasi-experimental variation can be used to identify the social effects. Our focus is on methods that take the network to be (conditionally) exogenous as well as perfectly observed by the researchers.

When data are available on only agents and the reference groups to which they belong, researchers have for some time worried about how social effects might be identified. However, when detailed data on nodes and their individual links are present, identification of social effects (taking the network as conditionally exogenous) is generic, and estimation is relatively straightforward. Three broader conceptual issues exist in this case.

First, theory is often silent on the precise form that peer effects should take when they exist. Since Manski (1993), many people have focused on the ‘local average’ framework, often without discussion of the implications for economic behaviour, but social effects might instead take a local aggregate, or indeed local maximum/minimum form where the best child in a classroom provides a good example to all others, or the worst disrupts the lesson. Until a non-parametric way of allowing for social effects is developed, researchers need to use theory to guide the empirical specification they use.

Second, researchers typically treat the observed network as the network which mediates the social effect, and where many networks are observed the union of these is taken. However, this could generate important biases if the actual network is different to that observed, either due to measurement error or because a different type of relationship is the relevant one for mediating a particular effect. Here again it is important that some justification is given for why the network used should be the appropriate one.

Finally, the observed network is typically treated as exogenous, or at least the errors are treated as being mean independent of the network. In many circumstances one might imagine that agents choose which links to form, and these choices may depend on characteristics that are neither observed by the econometrician nor independent of the outcome. This creates a problem for identification strategies that rely on the absence of links (*i.e.* using friends-of-friends), since the absence of a link may contain some information about the difference in the unobservables. For instance, more motivated pupils in a school may choose to link with other motivated pupils; or individuals may choose to become friends with other individuals who share a common interest (such as an interest in reading, or mathematics) that is unobserved in the data available to the researcher. In such examples, the absence of a link is due to the unobserved terms of the two agents being correlated in a specific way rather than the absence of correlation between these terms. Recent literature, including Goldsmith-Pinkham and Imbens (2013), Blume et al. (2015), Arduini et al. (2015) and Horace et al. (2015) among many others, has begun considering solutions to this issue. Many of these methods are reviewed in a longer working paper version of this article (see Advani and Malde, 2014).

Much work has been done to develop methods for working with networks data, both in economics and in other fields. Applied researchers can therefore take some comfort in knowing that many of the challenges they face using these data are ones that have been considered before, and for which there are typically at least partial solutions already available. Whilst the limitations of currently available techniques mean that empirical results should be interpreted with some caution, attempting to account for social effects is likely to be less restrictive than simply imposing that they cannot exist.

References

- A. Advani and B. Malde. "Empirical Methods for Networks Data: Social Effects, Network Formation and Measurement Error". *IFS Working Paper W14/34*, 2014.
- A. Advani and B. Malde. "Credibly Identifying Social Effects: Accounting for Network Formation and Measurement Error". *mimeo, Institute for Fiscal Studies*, 2016.
- George A. Akerlof. A Theory of Social Custom, of which Unemployment may be One Consequence. *The Quarterly Journal of Economics*, 94(4):749–775, 1980.
- V. Alatas, A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken. "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia". *American Economic Review*, 106(7):1663–1704, 2016.
- Rob Alessie and Arie Kapteyn. Habit formation, interdependent references and demographic effects in the almost ideal demand system. *Economic Journal*, 101(406):404–19, 1991.
- J. Angrist. "The Perils of Peer Effects". *Labour Economics*, 30:98–108, 2014.
- L. Anselin. *"Spatial Econometrics: Methods and Models"*. Springer, 1988.
- T. Arduini, E. Patacchini, and E. Rainone. "Parametric and Semi-Parametric IV Estimation of Network Models with Selectivity". *EIEF Working Paper*, 15/09, 2015.
- S. Athey, D. Eckles, and G. Imbens. "Exact P-values for Network Inference". *NBER Working Paper 21313*, 2015.
- C. Ballester, A. Calvó-Armengol, and Y. Zenou. "Who's who in networks. Wanted: the key player". *Econometrica*, 74:1403–1417, 2006.
- A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. Jackson. "The Diffusion of Microfinance". *Science*, 341:1236498, 2013.
- P. Bayer, F. Ferrira, and R. McMillan. "A Unified Framework for Measuring Preferences for Schools and Neighborhoods". *Journal of Political Economy*, 115(4):588–638, 2007.
- B Douglas Bernheim. A Theory of Conformity. *Journal of Political Economy*, 102(5):841–77, October 1994.
- A. Bisin, A. Moro, and G. Topa. "The Empirical Content of Models with Multiple Equilibria in Economies with Social Interactions". *NBER Working Paper*, WP 17196, 2011.
- S. Black. "Do Better Schools Matter? Parental Valuation of Elementary Education". *Quarterly Journal of Economics*, 114(2):577–599, 1999.

- F. Bloch, G. Genicot, and D. Ray. "Informal Insurance in Social Networks". *Journal of Economic Theory*, 143(1):36–58, 2008.
- L. E. Blume, W. A. Brock, S. N. Durlauf, and Y. M. Ioannides. "Identification of Social Interactions". In J. Benhabib, A. Bisin, and M. Jackson, editors, *Handbook of Social Economics*, volume 1B. North Holland, 2010.
- L. E. Blume, W. A. Brock, S. N. Durlauf, and R. Jayaraman. "Linear Social Interaction Models". *Journal of Political Economy*, 123 (2):444–496, 2015.
- V. Boucher and B. Fortin. "Some Challenges in the Empirics of the Effects of Networks". In Y. Bramouille, A. Galeotti, and B. Rogers, editors, *The Oxford Handbook of the Economics of Networks*, pages 277–302. Oxford University Press, 2015.
- Y. Bramoullé, H. Djebbari, and B. Fortin. "Identification of Peer Effects through Social Networks". *Journal of Econometrics*, 150:41–55, 2009.
- Y. Bramoullé, R. Kranton, and M. D'Amours. "Strategic Interaction and Networks". *American Economic Review*, 104(3):898–930, 2014.
- Y. Bramouille and R. Kranton. "Public Goods in Networks". *Journal of Economic Theory*, 135(1):478–494, 2007.
- W. A. Brock and S. N. Durlauf. "Discrete Choice with Social Interactions". *Review of Economic Studies*, 68:235–260, 2001.
- W. A. Brock and S. N. Durlauf. "Identification of Binary Choice Models with Social Interactions". *Journal of Econometrics*, 140:52–75, 2007.
- A. Calvó-Armengol, E. Patacchini, and Y. Zenou. "Peer Effects and Social Networks in Education". *Review of Economic Studies*, 76:1239–1267, 2009.
- Anne C Case. Spatial Patterns in Household Demand. *Econometrica*, 59(4):953–65, July 1991.
- A. G. Chandrasekhar. "Econometrics of Network Formation". In Y. Bramouille, A. Galeotti, and B. Rogers, editors, *The Oxford Handbook of the Economics of Networks*, pages 303–357. Oxford University Press, 2015.
- M. Comola and S. Prina. "Do Interventions Change the Network? A Dynamic Peer Effect Model Accounting for Network Changes". *SSRN Working Paper No. 2250748*, 2014.
- G. De Giorgi, M. Pellizzari, and S. Redaelli. "Identification of Social Interactions through Partially Overlapping Peer Groups". *American Economic Journal: Applied Economics*, 2(2):241–275, April 2010.
- A. de Paula. "Econometric Analysis of Games with Multiple Equilibria". *Annual Review of Economics*, 5:107–131, 2013.

- A. de Paula. "Econometrics of Network Models". *CeMMAP working paper CWP06/16*, 2016.
- M. DeGroot. Reaching a Consensus. *Journal of the American Statistical Association*, 69:118–121, 1974.
- R. Dieye, H. Djebbari, and F. Barrera-Osorio. "Accounting for Peer Effects in Treatment Response". *IZA Discussion Paper*, 8340, 2014.
- W. Gaertner. "A Dynamic Model of Interdependent Consumer Behavior". *Zeitschrift für Nationalökonomie*, 34:327–344, 1974.
- S. Gibbons and S. Machin. "Valuing English Primary Schools". *Journal of Urban Economics*, 53(2): 197–219, 2003.
- P. Goldsmith-Pinkham and G. W. Imbens. "Social Networks and the Identification of Peer Effects". *Journal of Business and Economic Statistics*, 31:253–264, 2013.
- B. S. Graham. "Identifying Social Interactions through Conditional Variance Restrictions". *Econometrica*, 76:643–660, 2008.
- B. S. Graham. "Methods of Identification in Social Networks". *Annual Review of Economics*, 7, 2015.
- W. Horace, X. Liu, and E. Patacchini. "Endogenous Network Production Functions with Selectivity". *Journal of Econometrics*, 190(2):222–232, 2015.
- M.O. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.
- M.O. Jackson, T. Rodriguez-Barraquer, and X. Tan. "Social Capital and Social Quilts: Network Patterns of Favor Exchange". *American Economic Review*, 102(5):1857–97, 2012.
- S. R. Jones. "*The Economics of Conformism*". Basil Blackwell New York/Oxford, 1984.
- L-F. Lee. "Identification and Estimation of Econometric Models with Group Interactions, Contextual Factors and Fixed Effects". *Journal of Econometrics*, 140:333–374, 2007.
- L-F. Lee and X. Liu. "Identification and GMM Estimation of Social Interactions Models with Centrality". *Journal of Econometrics*, 159:99–115, 2010.
- X. Liu. "Estimation of a local-aggregate network model with sampled networks". *Economics Letters*, 118:243–246, 2013.
- X. Liu, E. Patacchini, and Y. Zenou. "Endogenous Peer Effects: Local Aggregate or Local Average?". *Journal of Economic Behavior and Organization*, 103:39–59, 2014a.
- X. Liu, E. Patacchini, Y. Zenou, and L-F. Lee. "Criminal Networks: Who is the Key Player?". *Unpublished Manuscript*, 2014b.

- C. Manski. "Identification of Endogenous Social Effects: The Reflection Problem". *Review of Economic Studies*, 60:531–542, 1993.
- E. Patacchini and Y. Zenou. "Juvenile Delinquency and Conformism". *Journal of Law, Economics and Organization*, 1:1–31, 2012.
- E. Patacchini, E. Rainone, and Y. Zenou. "Heterogeneous Peer Effects in Education". *mimeo, Cornell University*, 2016.
- M. Patnam. "Corporate Networks And Peer Effects In Firm Policies". *mimeo, ENSAE-CREST*, 2013.
- Robert A Pollak. Interdependent Preferences. *American Economic Review*, 66(3):309–20, June 1976.
- F. Renna, I. Grafova, and N. Thakur. "The Effects of Friends on Adolescent Body Weight". *Economics and Human Biology*, 6:377–387, 2008.
- B. Sacerdote. "Peer Effects With Random Assignment: Results For Dartmouth Roommates". *Quarterly Journal of Economics*, 116:681–704, 2001.
- B. Sacerdote. "Peer Effects in Education: How Might They Work, How Big Are They and How Much Do We Know Thus Far?". In E. Hanushek, S. Machin, and L. Woessman, editors, *Handbook of the Economics of Education*, volume 3. Elsevier, 2011.
- R. Townsend. "Risk and Insurance in Village India". *Econometrica*, 62(3):539–591, 1994.