

IDENTIFYING PREFERENCES IN NETWORKS WITH BOUNDED DEGREE

ÁUREO DE PAULA

University College London, São Paulo School of Economics, CeMMAP, IFS, and CEPR

SETH RICHARDS-SHUBIK

Lehigh University and NBER

ELIE TAMER

Harvard University

This paper provides a framework for identifying preferences in a large network where links are pairwise stable. Network formation models present difficulties for identification, especially when links can be interdependent, for example, when indirect connections matter. We show how one can use the observed proportions of various local network structures to learn about the underlying preference parameters. The key assumption for our approach restricts individuals to have bounded degree in equilibrium, implying a finite number of payoff-relevant local structures. Our main result provides necessary conditions for parameters to belong to the identified set. We then develop a quadratic programming algorithm that can be used to construct this set. With further restrictions on preferences, we show that our conditions are also sufficient for pairwise stability and therefore characterize the identified set precisely. Overall, the use of both the economic model along with pairwise stability allows us to obtain effective dimension reduction.

KEYWORDS: Network formation, pairwise stability, partial identification, large network.

1. INTRODUCTION

THIS PAPER PROVIDES A FRAMEWORK for studying identification—what can be learned about parameters of interest from data—in strategic network formation models. The framework applies to complete information games with non-transferable payoffs in which individuals, given a particular utility function, form links with each other. Our objective is to learn about these payoffs from observed data on network linkages. In particular, we assume that observed networks are equilibrium networks and use pairwise stability, proposed in Jackson and Wolinsky (1996), as the solution concept.

Áureo de Paula: apaula@ucl.ac.uk

Seth Richards-Shubik: sethrs@lehigh.edu

Elie Tamer: elietamer@fas.harvard.edu

We thank the editor and five anonymous referees, along with seminar and conference participants at the Conference on Econometric Analysis of Social Interactions at Northwestern University (2010), the ASSA meetings in Chicago (2012), University of Western Ontario (2013), the Stats in Paris: Statistics and Econometrics of Networks (2013) conference, the Recent Advances in Set Identification (2013) conference, the 2nd European Meeting of Network Economics (2014), Heinz College (2015), the INET Econometrics of Networks Workshop at Cambridge University (2015), the Cowles Foundation (2015), Northwestern University (2015), the Chicago Fed (2015), the Econometric Society World Congress (2015), the USC Dornsife INET Conference on Networks (2015), the New York Fed (2016), the University of Pennsylvania (2016), Surrey (2016), Groningen (2016), the Tinbergen Institute (2016), Harvard (2016), Rutgers (2016), and Virginia (2016) for suggestions and comments. We are also grateful to A. Advani, J. Altonji, G. Chamberlain, C. Galdani, P. Haile, T. Holden, and M. Shum for detailed comments. de Paula gratefully acknowledges financial support from the NSF through award SES-1123990, the ERC through Starting Grant 338187, and the Economic and Social Research Council through the ESRC Centre for Microdata Methods and Practice Grant RES-589-28-0001.

A network formation model founded on a well-defined preference structure for all the players is helpful in determining how networks develop given a particular policy or incentive system. In a variety of applications, the architecture of networks is thought to influence final outcomes, and so it is important to understand how the networks themselves are formed. Empirically sound network formation models can also be helpful in describing why certain networks emerge and not others, for example, tracking the effects of policies or frictions on the kinds and shapes of networks that arise.

The problem in analyzing such strategic models is that multiplicity of solutions and other computational difficulties are pervasive, especially when more than a handful of agents are involved. This paper proposes a framework for studying large networks, which achieves dimension reduction by aggregating individuals with similar network positions. We impose restrictions on the number of links that a person would choose to have, as well as on the cardinality of the observable characteristics. As a consequence, each individual's position in the network can be described as one of a finite number of possible local network structures. This greatly reduces the dimensionality of the problem. The limitation on the number of links, or bounded degree, is of course an important restriction. It makes this framework appropriate for applications where individuals are observed to have a relatively small number of connections, but not for others where it is important to have some nodes with many links (e.g., network "hubs").

Our approach for characterizing set-identified parameters bypasses the selection of a particular equilibrium (when many are possible) and directly exploits the economic predictions under pairwise stability. The key idea is that each individual can be classified as the central node (the "ego") of one of the finite possible local network structures, which we refer to as *network types*. The set of relevant types is determined by the preference specification, so, for example, a specification where only direct connections matter will lead to a different set of types than a specification where indirect connections matter as well. The link between the observed frequencies of these network types and their model predicted ones allows us to learn about preference parameters. A utility structure generates a set of payoff-relevant network types, and a parameterization of the network formation model predicts the measures or proportions of these types in the network. The inference question then reduces to collecting all utility parameters that can predict proportions of network types that match the proportions estimated in the data.

Developing this correspondence in a computationally feasible way represents the main contribution of this paper. Our main result provides necessary conditions for a set of parameters to generate a pairwise stable network with a distribution of network types that matches the observed distribution. For preference structures that satisfy an additional restriction, we show that these conditions are sufficient as well, and so are able to characterize the identified set precisely. Then, to implement the approach, we show that the majority of the computation can be done through a series of quadratic programming problems.

In part because of the difficulties indicated above, the literature on the econometric analysis of strategic network formation models is small, but growing.¹ The payoff structures we analyze are related to those contemplated in, for example, [Currarini, Jackson, and Pin \(2009\)](#), [Christakis, Fowler, Imbens, and Kalianaraman \(2010\)](#), and (for a directed network) [Mele \(2017\)](#). The most similar papers to ours are those which also use pairwise stability and identify a set of parameters that could be consistent with the data in

¹See [Graham \(2015\)](#) and [de Paula \(2017\)](#) for recent surveys.

any equilibrium (Sheng (2014), Miyauchi (2016), Leung (2015)).² Relative to these papers, our proposed method does not require certain restrictions on preferences (e.g., ruling out negative externalities or requiring a homophilous attribute), and it may be more computationally tractable. A number of other papers in this literature rely on dynamic meeting protocols for the formation of the network (Christakis et al. (2010), Mele (2017), Badev (2013)).³ Chandrasekhar and Jackson (2014) proposed a different approach where the network is generated from overlapping sub-graphs. Also, some recent papers consider the estimation of dyadic link formation models (i.e., without link externalities) with a focus on disentangling homophily and node-specific heterogeneity (Charbonneau (2017), Graham (2017), Dzemski (2014)).

Because matching models essentially aim at characterizing a bipartite graph, and hence a particular type of network, those models are also related to strategic network formation. There is a growing literature on the econometrics of matching (e.g., Choo and Siow (2006), Fox (2010, 2018), Galichon and Salanie (2009), Echenique, Lee, and Shum (2010), Chiappori, Galichon, and Salanié (2016), Menzel (2015)). Our setting differs in substantive aspects, however: indirect connections are payoff-relevant, utility is non-transferable, and multiple equilibria are possible (in contrast to some models in that literature). Also the concept of pairwise stability in matching games is related but not identical to the Jackson and Wolinsky (1996) definition, where only one link at a time is considered.

2. NETWORK FORMATION MODEL

Our framework applies to complete information games that produce an undirected network. One example is a static game in which players simultaneously announce the set of other players they would like to be connected with, links form if they are mutually beneficial, and payoffs are received.⁴ We use a continuum of players, $i \in N \equiv [0, \mu]$, where $\mu > 0$ is their total measure. This is unlike the typical setup in the empirical games literature, but the modeling choice is natural in our setting where there is a large number of players (Khan and Sun (2002)). Each player has some predetermined characteristic(s) $X_i \in \mathcal{X}$ that is (are) observed by the econometrician, and player-pairs have a one-dimensional characteristic ε_{ij} that is not observed by the econometrician. Nature draws $\mathbf{X} = (X_i)_{i \in N}$ and $\boldsymbol{\varepsilon} = (\varepsilon_{ij})_{(i,j) \in N^2}$, and these vectors are common knowledge to all players.

The network that results from players' actions is characterized by the adjacency mapping $G : N \times N \rightarrow \{0, 1\}$. This is a *continuous graph* as there is a continuum of nodes.⁵ Such graphs (particularly refinements known as *graphings*, for limits of bounded degree graphs) are a recent development and are used as approximations for large graphs under a well-defined metric (for a review, see Lovasz (2012)). Hence, we view the continuous graph model as a close approximation to a model with a large but finite number of players.⁶

²Boucher and Mourifié (2017) developed a method that is similar to Leung (2015) but bypasses the issue of multiplicity.

³Dynamic data on link formation are rarely available. As pointed out in Mele (2017), the meeting protocol therefore acts as an equilibrium selection mechanism. Additionally, Badev (2013) extended the model to include actions beyond link formation.

⁴Pairwise stability can also describe the rest point of a sequential meeting process (Jackson (2009)).

⁵Formally, the graph consists of the adjacency mapping G and set of players N . However, we typically refer to G as the “network.”

⁶To be clear, we use the word “large” in its colloquial sense when referring to a network, not in the econometric sense that would imply the number of nodes is sent to infinity.

Payoffs depend on the network configuration and covariates, and are denoted by $u_i(G, \mathbf{X}) = u(G, \mathbf{X}; \varepsilon_i)$, where $\varepsilon_i = (\varepsilon_{ij})_{j \neq i}$. Our objective is to learn about the (parametric) utility functions $u(G, \mathbf{X}; \cdot)$ using the data on G and \mathbf{X} .

To make the model tractable, we rely on two main assumptions about the payoffs. We start with a restriction on network depth and total number of links.

ASSUMPTION 1: Only connections up to distance D affect utility, and preferences are such that players will never choose more than a total of L links.

The distance above refers to the length of the shortest path between two individuals, denoted $d(i, j; G)$. If $D = 1$, only direct connections are relevant (e.g., Currarini, Jackson, and Pin (2009)). When $D > 1$, indirect connections also matter, and here most specifications in the literature use $D = 2$ (e.g., “friends of friends”). The limit L denotes the maximum number of links an individual would have (i.e., utility would be infinitely negative if you have more than L links). This restricts our framework to networks with bounded degree distributions, where nodes have a relatively small number of links, rather than networks with approximately power-law degree distributions. Networks with such limited degree distributions are found in several social science contexts (e.g., close friendships) but not in others (e.g., Facebook).⁷ Together, the restrictions on depth and degree in Assumption 1 make payoffs depend on a finite number of direct and indirect connections in the network. This is crucial for dimension reduction. For example, with $D = 2$, there would be at most L direct alters and $L \times (L - 1)$ indirect alters that impact utility.

Our second assumption relates to the unobservable preference shocks, as well as the support of the observable characteristics. We assume that the preference shocks do not depend on the individual identities of the alters. Instead, there is one shock for each possible direct connection, combined with their possible predetermined characteristics. We further assume that the predetermined characteristics have finite support, and that the unobserved shocks are independent of these characteristics.

ASSUMPTION 2: Individuals are endowed with $L \times |\mathcal{X}|$ preference shocks, denoted $\varepsilon_l(x)$, $l = 1, \dots, L$, $x \in \mathcal{X}$, which correspond to the possible direct connections and their characteristics. This vector of preference shocks is independent of X with a known distribution (possibly up to some finite-dimensional parameter). In addition, the support of X is finite.

This assumption implies that if two potential alters have the same observables, then the ego in question is indifferent between them. Similar (though not identical) assumptions about homogeneity in preferences have been made in models of large games (e.g., Kalai (2004), Menzel (2016)) and in some matching models (e.g., Choo and Siow (2006), Galichon and Salanie (2009)). This helps control the dimensionality of the problem, and it can be a natural restriction in settings with many agents where individual identities are unknown and irrelevant to the researcher. In addition, having a limited number of shocks allows the model to retain a positive fraction of isolated individuals in equilibrium even

⁷For example, such a limitation is seen in the National Longitudinal Study of Adolescent to Adult Health (also known as the “Add Health” study), a commonly used data set on social networks. Individuals nominate up to five friends of each sex, and the number of reciprocated nominations is even smaller. The median number of such links is one, and less than five percent of individuals have more than three links to the same sex. Similarly, in data on networks among Indian villagers used in Jackson, Rodriguez-Barraquer, and Tan (2012), fewer than 1 per 1,000 respondents reached the caps of 5 or 8 nominations, respectively, for various kinds of social and financial relationships (footnote 37, p. 1879).

when the group under consideration is large.⁸ This assumption could be weakened somewhat, by extending the vector of shocks to include ones for indirect connections up to distance D . In that case, the number of shocks would be $|\mathcal{X}| \times L \sum_{d=1}^D (L - 1)^{d-1}$ (for $L > 1$). Our theoretical results would be unchanged, but the computational burden would of course increase. Last, we note that Assumption 2 allows for correlation within the vector of preference shocks. This has economic relevance since it captures preference correlation among unobservables. However, generally allowing for unrestricted correlation might lead to exceedingly large identified sets.

Assumptions 1 and 2 together imply that there is a finite number of possible configurations of alters and their characteristics within the payoff-relevant distance from any given individual. Then, much as in discrete choice models for differentiated products, utility is determined from a fixed and finite list of possible outcomes. Our proposed inference strategy relies on this feature of the model to reduce the dimensionality of the problem from the universe of possible network configurations to the much smaller set of possible payoff-relevant local subnetworks, which we refer to as the network types in the model.⁹

Next, we provide an example utility specification, which we will return to throughout the paper. Payoffs depend on the predetermined characteristics of the individual and her direct connections ($f(x_i, x_j) + \varepsilon_{il(j)}(x_j)$), any links among these direct connections (ω), and any further connections at distance 2 (ν). Such specifications are commonly used to model high school friendships, where the observable characteristics are race and perhaps other family attributes (e.g., Christakis et al. (2010), Mele (2017), Goldsmith-Pinkham and Imbens (2013), Sheng (2014), Miyauchi (2016)). We use this context descriptively as well.¹⁰ The utility function is expressed as follows:

$$\begin{aligned}
 u_i(G, \mathbf{X}) \equiv & \sum_{j \in N(i)} (f(x_i, x_j) + \varepsilon_{il(j)}(x_j)) && \text{(direct connections)} \\
 & + \left| \bigcup_{j \in N(i)} N(j) - N(i) - \{i\} \right| \nu && \text{(friends of friends)} \\
 & + \sum_{j \in N(i)} \sum_{\substack{k \in N(i): \\ k > j}} G(j, k) \omega && \text{(mutual friends)} \\
 & - \infty \cdot 1_{|N(i)| > L}, && \text{(bounded degree)}
 \end{aligned} \tag{1}$$

⁸Otherwise, if there were, for example, i.i.d. preference shocks for every potential connection in the network, the probability that a link is mutually beneficial with at least one other person increases as the size of the network increases (and hence the number of preference shocks increases).

⁹The use of subnetworks here is rather distinct from their use in Sheng (2014), although both are promising strategies for dimension reduction. We consider all possible subnetworks among individuals that are within some distance from a reference individual, where the distance is determined by the specification of preferences. Sheng (2014) considered subnetworks among arbitrary individuals, where the number of individuals in the subnetwork is chosen for computational tractability and is unrelated to the model. These approaches offer different advantages and should be viewed as complementary developments.

¹⁰We believe the approach we put forward can accommodate many other applications. Our presentation is based on nontransferable utility (NTU). Applications elsewhere, for example, in Industrial Organization (e.g., networks of buyers and suppliers or hospitals and insurers), may rely instead on transferable utility (TU). In that case, pairwise stability would require a positive surplus for observed links and a negative one for absent links, and preference classes, which we introduce later in the article, may need to rely on dyads rather than individuals. It may also be important to consider differences in “numerosity” whereby one side of the market (e.g., insurers) is substantially smaller than the other (e.g., hospitals). A careful examination of these variations merits much more space and attention.

where $N(i)$ denotes the neighbors of node i (i.e., $N(i) \equiv \{j : G(i, j) = 1\}$) and $|\cdot|$ gives the cardinality of a set. In the first line, the index function $l(j)$ assigns neighbor j to the l th link of node i and thereby assigns the preference shocks.¹¹ The second line takes the union of the neighbors of each friend but removes the friends themselves ($N(i)$) and the reference individual ($\{i\}$) to find the number of distinct friends of friends. The third line counts any links among the direct connections, and the fourth line ensures bounded degree in equilibrium. Like related specifications in the literature, this uses a maximum depth of $D = 2$, and has additively separable shocks for direct connections from some known distribution.

As in most of the empirical games literature, we assume that observed choices correspond to equilibrium play. Our solution concept is pairwise stability (Jackson and Wolinsky (1996)).

DEFINITION 1—Pairwise Stability: All links ij must be preferred by players i and j over not having the link, and all non-existing links must be damaging to at least one of the players:

$$\forall i, j: G(i, j) = 1, \quad u_i(G, \mathbf{X}) \geq u_i(G_{-ij}, \mathbf{X}) \quad \text{and} \quad u_j(G, \mathbf{X}) \geq u_j(G_{-ij}, \mathbf{X}); \quad (\text{i})$$

and

$$\forall i, j: G(i, j) = 0, \quad \text{if } u_i(G_{+ij}, \mathbf{X}) > u_i(G, \mathbf{X}) \quad \text{then} \quad u_j(G_{+ij}, \mathbf{X}) < u_j(G, \mathbf{X}). \quad (\text{ii})$$

In the definition, G_{-ij} denotes the mapping $(k, l) \mapsto G_{-ij}(k, l) = G(k, l)$ if $(k, l) \neq (i, j)$ and $(k, l) \mapsto G_{-ij}(k, l) = 0$ if $(k, l) = (i, j)$. Analogously, G_{+ij} denotes the mapping $(k, l) \mapsto G_{+ij}(k, l) = G(k, l)$ if $(k, l) \neq (i, j)$ and $(k, l) \mapsto G_{+ij}(k, l) = 1$ if $(k, l) = (i, j)$. Other solution concepts exist; see Bloch and Jackson (2006) or Jackson (2009). As discussed in those references, an advantage of pairwise stability is that it incorporates the intuition that, in a social setting, agents are likely to communicate to form mutually desirable connections. This is not the case with Nash equilibrium, where absent links can still be part of an equilibrium even though they would be mutually beneficial.

3. PREVIEW OF OUR RESULTS

Here, we provide a preview of the approach using a very simple case of specification (1). In this example, individuals have at most one link ($L = 1$, also $D = 1$), and there are two races: B (black) and W (white). This could describe a network of best friends, which consists only of isolates and linked pairs. There are no externalities from links in such a network, but it nevertheless illustrates the main features of our approach.

Outcomes can be expressed as ordered pairs, (x, y) , for the individual's race (x) and the best friend's race (y , where $y = 0$ if no best friend). For example, (B, W) corresponds to a black individual with a white best friend. These pairs represent the *network types* in this model. (More generally, network types will involve a local adjacency matrix as well as a vector of node characteristics such as (B, W) , but here the matrices are redundant.) Utility depends on an individual's network type, (x, y) . The utility function (1) simplifies here to $u_i(x, y) = f_{xy} + \varepsilon_i(y)$, where the f_{xy} , $x, y \in \{B, W\}$, are four parameters, and each individual is endowed with two preference shocks, $\varepsilon_i(B)$ and $\varepsilon_i(W)$, with some known

¹¹The function $l(j)$ implicitly depends on G as well, because in different graphs the neighbor j could be assigned to different links. This is omitted to alleviate the complexity of the notation.

distribution (up to a finite-dimensional parameter). Our goal is to use data on the linkages in the network to learn about the parameters f_{xy} .

First, we collapse the global graph and node characteristics (G, \mathbf{X}) into the shares of individuals of each network type, or *type shares*. For example, suppose, in a school with 500 students, there were 50 black individuals with a white best friend. The share of type (B, W) is 0.1 (the share of type (W, B) is also 0.1, as they must balance). We will search for parameter values f_{xy} that can generate the observed type shares while satisfying necessary conditions for pairwise stability.

To do this, we start by classifying individuals based on which network types they would not reject (i.e., they would be happy with all the links). For example, depending on the preference shocks drawn, a black individual may prefer having a black best friend to being alone, but may not prefer having a white best friend to being alone: that is, $f_{BB} + \varepsilon_i(B) \geq 0$ and $f_{BW} + \varepsilon_i(W) < 0$. Hence, the network type (B, W) could not be an equilibrium outcome for this individual, but (B, B) could be. We refer to these sets of network types that individuals would not unilaterally deviate from as *preference classes*, generically denoted as H . For this individual, the preference class would be $H = \{(B, 0), (B, B)\}$. (Because there are no connections to be dropped from an isolated type, e.g., $(B, 0)$, every preference class contains one such type.) There are four possible preference classes for blacks in this example: $H_1 = \{(B, 0)\}$, that is, prefers to be alone; $H_2 = \{(B, 0), (B, B)\}$, that is, prefers a black best friend; $H_3 = \{(B, 0), (B, W)\}$, that is, prefers a white best friend; $H_4 = \{(B, 0), (B, B), (B, W)\}$, that is, prefers a best friend of either race. The preference classes for whites are similar, replacing the first race in each type with W .

Each preference class corresponds to some region in the space of the shocks, ε , that determines which network types would be “acceptable” to an individual with those shocks (i.e., an individual with shocks from the region for preference class H would not unilaterally deviate from any network type in H , but would deviate from any type not in H). Hence, given a distribution for the preference shocks and proposed values for the preference parameters, one can compute the probability that individuals fit into each preference class. In this example, there are simple thresholds in $\varepsilon(B)$ and $\varepsilon(W)$ based on the parameters f_{xy} that yield these probabilities. With more elaborate models, the preference class probabilities can be computed, for example, via Monte Carlo integration (see Section 7 and Supplemental Material Appendix D.7.1 (de Paula, Richards-Shubik, and Tamer (2017))).

Given the preference class probabilities derived from a vector of structural parameters, we can then generate predicted type shares by allocating the individuals from each preference class to the possible network types. To do this, we define *allocation parameters*, denoted $\alpha_H(\cdot)$, one for each type in each class, which designate the proportion of individuals allocated from preference class H to network type “.”. For example, the predicted share of blacks with a white best friend, type (B, W) , is $\Pr(H_1|B)\alpha_{H_1}(B, W) + \Pr(H_2|B)\alpha_{H_2}(B, W) + \Pr(H_3|B)\alpha_{H_3}(B, W) + \Pr(H_4|B)\alpha_{H_4}(B, W)$ (multiplied by the proportion of blacks in the school to obtain the share among all students).

The key to our approach is to provide restrictions on the allocation parameters that need to be satisfied in order for the network to be pairwise stable. These restrictions correspond to necessary equilibrium restrictions and agreement with the data. First, individuals may only be allocated to network types contained in their preference classes. This restricts some of the allocation parameters to zero. Second, given any pair of network types that could feasibly add a link with each other (i.e., an isolated individual of either race in this example), the measure of individuals who would prefer to do so must be zero for at least one of the two types. Otherwise, additional mutually beneficial links could be

formed and the network would be unstable. Hence, for any pair of types, the product of the measures of individuals of one type who would prefer to add links to individuals of the other type must be zero. This defines a quadratic objective function which in equilibrium has to be zero. Finally, the predicted proportions of types must match the observed proportions of types in the network, which defines a set of linear constraints.

4. NETWORK TYPES AND PREFERENCE CLASSES

Now we formalize and extend the concepts discussed in the preceding example. Our proposed identification strategy is built on the notion of predefined *network types*. These describe the local network structure around a given individual, along with the predetermined characteristics of each person (or node) in this local subnetwork. The size of the local subnetworks depends on the preference specification, specifically the parameters D and L that control the relevant depth in the network and maximum number of links. The predetermined characteristics are fixed attributes, such as sex and race, or predetermined behaviors (i.e., those which precede the formation of the network), for example, the education levels of coworkers at a firm. Intuitively, then, network types can be described in words, for example, “a female connected to two females and one male,” “an unconnected low-income male,” “a female connected to another female with two other friends,” and so on.

More formally, a network type is characterized by a local adjacency matrix, A , and a vector of node characteristics, v . The matrix A describes the local subnetwork up to distance D from the reference individual, who is called the *ego* of the network type. It is symmetric and has one row and column for the ego and one for each possible alter up to distance D . The first row corresponds to the ego and indicates that individual’s links with a 1 in the appropriate columns, and 0 otherwise. The next L rows correspond to the possible direct alters, then the next $L(L - 1)$ rows to the possible alters at distance 2, and so on. This gives a total of $1 + L + L(L - 1) + L(L - 1)^2 + \dots + L(L - 1)^{D-1} = 1 + L \sum_{d=1}^D (L - 1)^{d-1}$ rows.¹² The vector v contains the predetermined characteristics of the ego and the alters, in the same order as the rows of A . The first element of v is the characteristic of the ego, denoted v_1 . The subsequent elements are the characteristics of the possible alters, belonging to $\mathcal{X} \cup \{0\}$, where 0 denotes the absence of an alter in that position.¹³ Thus we have the following:

DEFINITION 2—Network Type: Fix D , L , and \mathcal{X} . A network type t is characterized by $t = (A, v)$, where A is a square matrix of size $1 + L \sum_{d=1}^D (L - 1)^{d-1}$ and v is a vector of same length as the number of rows in A . This matrix describes the local subnetwork that is utility-relevant for an individual of type t . The vector v contains the predetermined characteristics of this person and the alters in the local subnetwork. The complete enumeration of network types generated from a preference structure u and set of characteristics \mathcal{X} is given by \mathcal{T} .

¹²This count simplifies to $1 + L[(L - 1)^D - 1]/(L - 2)$ lines when $L > 2$, and is $1 + 2D$ when $L = 2$ and simply 2 when $L = 1$ (since links are reciprocated).

¹³Typically, there will be a collection of isomorphic matrix and vector pairs (A, v) that can characterize a given type t . The first row and column of A and the first element of v are reserved for the ego, but the remaining rows and columns of A and elements of v could be permuted and still express the same local subnetwork that is rooted in the ego. For computational convenience, we adopt a convention that singles out a representative characterization for each network type (see Supplemental Material Appendix D.5.1.)

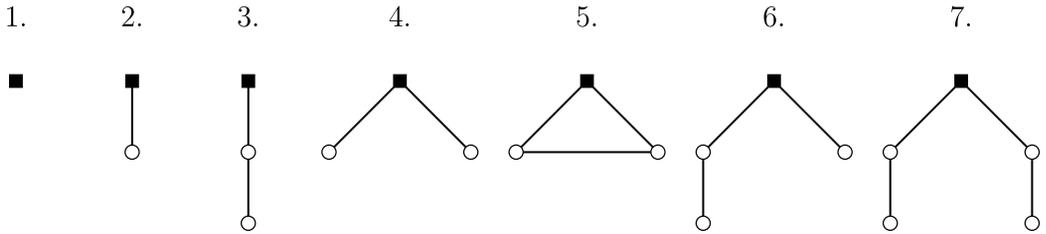


FIGURE 1.—Graphs of network types in Example 1.

As the definition says, the set of network types, \mathcal{T} , is determined completely by the preference structure and the set of predetermined characteristics. For example, if preferences are such that individuals have a taste for at most one friend (and there are no X 's), then there would be only two types: “alone” and “connected.” Each person in the network is one of these mutually exclusive and exhaustive types (i.e., each person can be identified as an ego of one particular network type), and we assume that the proportions of individuals of each network type can be consistently estimated from the observed data (see Section 6.2).

Below, we use a particular example of specification (1) to illustrate the network types that arise from a given preference structure and to show their representation as (A, v) .

EXAMPLE 1: Let preferences be given by (1), with $D = 2, L = 2$, and $\mathcal{X} = \{B, W\}$. An individual has at most two direct alters and two indirect alters at distance 2.¹⁴ The graphs of all the relevant network types under these preferences are visualized in Figure 1. The ego is represented with a square and the alters are circles. There are seven distinct graphs and a total of 62 different types, found by populating the nodes of these graphs with different combinations of the two races.

Figure 2 shows the local adjacency matrix A and vector of predetermined characteristics v for two network types in this example. The first type is a triangle (graph 5 from Figure 1) where the reference individual (the ego) is black, with one black friend and one white friend who are also friends with each other. The first row of A indicates the ego's links to the two direct alters (in columns 2 and 3). The second and third rows correspond to the direct alters and indicate that they are linked to the ego and to each other. The second type is a tree (graph 6) where the ego is black with one black friend and one white friend, and the black friend has a further connection (i.e., a friend of a friend to the ego) who is also black. Here, the direct alters are not linked to each other (elements (2, 3) and (3, 2) of A are 0), but the first direct alter (in row 2) is linked to an indirect alter represented in row and column 4.

<p>(a) First Example Type (Graph 5)</p> $A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad v = \begin{pmatrix} B \\ B \\ W \\ 0 \\ 0 \end{pmatrix}$	<p>(b) Second Example Type (Graph 6)</p> $A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad v = \begin{pmatrix} B \\ B \\ W \\ B \\ 0 \end{pmatrix}$
---	--

FIGURE 2.—The matrix A and vector v for two network types in Example 1.

¹⁴We are grateful to an anonymous referee who suggested using an example with two links.

Next, we note that under Assumptions 1 and 2, utility depends only on an individual’s network type. Assumption 1 limits payoffs to depend on a subnetwork represented by A , and Assumption 2 restricts the unobserved shocks to depend on the characteristics listed in v , not the individual identities of the alters. Accordingly, we rewrite the utility function as $u_i(G, \mathbf{X}) = u(A, v; \varepsilon_i)$. The vector ε_i contains one preference shock for each potential direct connection and each possible characteristic of those connections (i.e., each element of $L \times |\mathcal{X}|$). In Example 1, there would be four shocks: $\varepsilon_i = (\varepsilon_{i1}(B), \varepsilon_{i1}(W), \varepsilon_{i2}(B), \varepsilon_{i2}(W))$.¹⁵ For the two example types shown in Figure 2, the utilities would be $f(B, B) + \varepsilon_{i1}(B) + f(B, W) + \varepsilon_{i2}(W) + \omega$ for the triangle with mutual friends and $f(B, B) + \varepsilon_{i1}(B) + f(B, W) + \varepsilon_{i2}(W) + \nu$ for the tree with a friend of a friend (the only difference is whether ω or ν appears).

In order to make predictions from the model, we will categorize individuals based on their preferences over network types. For this, we define *preference classes*, the second important concept in our framework. These are sets of types that individuals would not unilaterally deviate from, given their own preferences. In other words, individuals would not reject any of their links if assigned to one of the types in their preference class, but would reject a link if assigned to a type outside their preference class. Naturally, each preference class corresponds to a region in the support of ε : individuals with shocks in the same region (and with the same predetermined characteristics) would be content with the same set of types. The preference classes therefore partition the support of ε into regions that can rationalize different sets of network types, from an individual perspective.

Given the predetermined characteristics ($X_i = x$) and preference shocks (ε_i) of an individual, it is fairly straightforward to find his or her preference class. We compare the utility of each network type (where $v_1 = x$) against the utility that would be obtained by removing any of the links in that type, and retain all types that are not dominated by this. To express these comparisons, we define the matrix A_{-l} to be equal to the local adjacency matrix A but with the l th link removed.¹⁶ The preference class is then the set of types (A, v) with $v_1 = x$ such that $u(A, v; \varepsilon_i) \geq u(A_{-l}, v; \varepsilon_i), \forall l = 1, \dots, L$. We indicate how this works below.

EXAMPLE 1—Continued: Suppose that the values of mutual friendships, ω , and friends of friends, ν , are positive. Consider an individual who is black ($X_i = B$), with preference shocks for black friends ($\varepsilon_{i1}(B)$ and $\varepsilon_{i2}(B)$) such that $f_{BB} + \varepsilon_{i1}(B) \geq 0, f_{BB} + \varepsilon_{i2}(B) < 0, f_{BB} + \varepsilon_{i2}(B) + \omega \geq 0$, and $f_{BB} + \varepsilon_{i2}(B) + \nu \geq 0$. The first inequality above means that this person would prefer to have one black friend over being isolated. Hence, the network type with graph 2 from Figure 1 and $v = (B, B, 0, 0, 0)'$ would be in this person’s preference class. Also, because ν is positive, the network types with one friend and one friend of a friend—that is, graph 3 and $v = (B, B, 0, B, 0)'$ or $v = (B, B, 0, W, 0)'$ —would be in the preference class as well. The second inequality ($f_{BB} + \varepsilon_{i2}(B) < 0$) means that i would not want to have a second black friend with no other connections. Hence, the preference class would not include network types with graphs 4 or 6 from Figure 1, where the second friend is black. However, the third and fourth inequalities above mean that the

¹⁵As noted earlier, this could be extended to include shocks for the indirect alters (friends of friends) as well. Then ε_i would have eight elements in total (four possible alters times two races), in this example.

¹⁶In other words, set elements $(1, l + 1)$ and $(l + 1, 1)$ in A to zero (because alter number l appears in row $l + 1$). There may remain nodes in the subnetwork represented by A_{-l} that are irrelevant to the resulting network type, because there is no path shorter than D between them and the ego once the l th link is removed. All the entries in the rows and columns of A_{-l} corresponding to these nodes could be replaced with 0, as could the corresponding elements of v . Regardless, these nodes will have no impact on $u(A_{-l}, v; \varepsilon_i)$.

shock $\varepsilon_{i2}(B)$ is large enough so that i would want to have a second black friend if there was a mutual friendship or an indirect connection (i.e., friend of friend) via this second friend. Hence, the preference class would contain the network type with graph 5 where $v = (B, B, B, 0, 0)'$ and the types with graph 7 and $v = (B, B, B, v_4, v_5)$, where v_4 and v_5 are either B or W .

The discussion above identifies seven types in individual i 's preference class (and it rules out certain other types). Similar conditions on the shocks for white friends ($\varepsilon_{i1}(W)$ and $\varepsilon_{i2}(W)$) would determine which types with white friends, or a combination of black and white friends, are also in the preference class, thereby completing its construction.

To formalize the definition of preference classes, we introduce a mapping ψ that yields the preference class for a given value of (x, ε) :

$$(x, \varepsilon) \mapsto \psi(x, \varepsilon) = H \subset \mathcal{T}, \tag{2}$$

where $H = \{(A, v) \in \mathcal{T} : v_1 = x \text{ and } u(A, v; \varepsilon) \geq u(A_{-l}, v; \varepsilon), \forall l = 1, \dots, L\}$. The definition is then as follows.

DEFINITION 3: Given a preference structure u , let ψ be as defined in (2). A *preference class* $H \subset \mathcal{T}$ is a set of network types in the range of ψ .

One important note is that every preference class includes an isolated type with the appropriate ego characteristics (e.g., in the example above, $A = 0_{5 \times 5}$ and $v = (B, 0, 0, 0, 0)'$). This is implicit from the definition of A_{-l} : there are no links to drop in the adjacency matrix for an isolated type, so $A_{-l} = A$. The presence of an isolated type in each preference class is important in order to allow for individuals who wind up isolated even if they desire connections. A second note is that preference classes must be obtainable under u from some vectors of characteristics and preference shocks. Hence, the number of preference classes is smaller than the number of subsets of \mathcal{T} , and typically much smaller. That is because the preference structures in these models imply dependencies among the network types that can appear together in a preference class. This is seen in Example 1, where, for example, the positivity of v implies that if a class includes a type with an unconnected friend (Figure 1, graph 2), it will also include types where that friend has a further friend (Figure 1, graph 3).

Finally, we can use the preference classes to generate predictions about the measures or proportions of individuals of each network type. First, we need the probability of each preference class. This is the probability that ε falls into the region corresponding to that class. We define these probabilities conditionally on the characteristics of the ego, as those characteristics are fixed within a given class. Accordingly, the preference class probabilities are denoted as $P_{H|v_1} \equiv P(\varepsilon : \psi(X, \varepsilon) = H | X = v_1)$. These are direct functions of the utility specification and the distributions of the unobservables. Hence, given a parameterization of the model, these probabilities are known.

Then, to generate the predictions, we specify how many individuals from each preference class are assigned to each network type. This uses *allocation parameters*, defined as follows.

DEFINITION 4: An *allocation parameter* $\alpha_H(t) \in [0, 1]$ gives the proportion of individuals with preferences in preference class H that are of network type t .

The allocation parameters yield the measure of individuals of network type t as follows: $\mu_{v_1(t)} \sum_H P_{H|v_1(t)} \alpha_H(t)$, where $\mu_{v_1(t)}$ is the measure of individuals with characteristic $v_1(t)$ (the characteristic of the ego in type t). The proportion of individuals of network type t is this divided by the total measure μ . This provides the exact link between the data and the underlying preferences. The measures or proportions of individuals of each network type can be consistently estimated, and we will try to match these with predictions from the model.

5. IDENTIFICATION WITH NETWORK TYPES

In this section, we show how to use the model to map the observed proportions of individuals of each network type (or more succinctly, the *type shares*) into restrictions on the preference parameters. We develop two general conditions on the allocation parameters that are necessary for pairwise stability, which can then be used to collect preference parameters that could be compatible with the observed type shares. If, using allocation parameters that satisfy these conditions, a vector of structural preference parameters cannot predict the observed type shares, then that vector is not compatible with the observed network. Otherwise, if such a prediction can be made, the vector is included in the recovered set. The two conditions are as follows (their intuition is discussed after the theorem):

CONDITION 1—Existing Links: All existing links are pairwise stable. For any type t and preference class H , $t \notin H \implies \alpha_H(t) = 0$.

CONDITION 2—Nonexisting Links Between Distant Individuals: There are no mutually beneficial links to add between individuals who are at a distance greater than $2D$ ($d(i, j; G) > 2D$). For every pair of types t, s where the egos of both types have fewer than L links, and for the pair of types \tilde{t}, \tilde{s} that would result if a link were added between two individuals of these types who are greater than $2D$ from each other,

$$\left(\mu_{v_1(t)} \sum_{\tilde{H} \in \mathcal{H}} P_{\tilde{H}|v_1(t)} \alpha_{\tilde{H}}(t) 1_{\tilde{t} \in \tilde{H}} \right) \cdot \left(\mu_{v_1(s)} \sum_{\tilde{H} \in \mathcal{H}} P_{\tilde{H}|v_1(s)} \alpha_{\tilde{H}}(s) 1_{\tilde{s} \in \tilde{H}} \right) = 0.$$

The theorem below provides our general result on identification. It takes as given the predicted probabilities of the preference classes, $P_{\cdot|s}$, which are yielded by a parameterization of the utility function. The theorem provides necessary conditions for a pairwise stable network with specified type shares (i.e., the observed shares) to exist given this distribution of preference classes in the population. To state the theorem, we denote the vector of type shares as $\pi \equiv (\pi_t)_{t \in \mathcal{T}}$, where the element π_t is the proportion of individuals in the network who are of network type t . We maintain in this paper that our sample contains information on exactly these type shares.¹⁷ For a given vector π and a distribution of preference classes $\{P_{H|v_1}\}$, we have the following result. (The proof is in Appendix A.)

THEOREM 1: *Let the vector $(\pi_t)_{t \in \mathcal{T}}$ be known. Given a probability distribution of preference classes $\{P_{H|v_1}\}$, if there exists a pairwise stable network where the proportion of agents of type t is equal to π_t for each $t \in \mathcal{T}$, then there exists a vector of allocation parameters α satisfying Conditions 1 and 2 such that π_t is equal to $\frac{1}{\mu} \sum_H \mu_{v_1(t)} P_{H|v_1(t)} \alpha_H(t)$ for every $t \in \mathcal{T}$.*

¹⁷Our sampling approach to inference maintains that nodes are sampled randomly to collect information on types. This is common in the network statistics literature as it is common for networks to only be partially observed in practice (see, e.g., Chapter 5 in *Kolaczyk (2009)*).

This result can be extended to show that Conditions 1 and 2 are also sufficient for the existence of a pairwise stable network, under further restrictions on preferences. The main restriction is that the payoff from adding a link to a distant individual must be greater than that from adding a link to a nearby individual of the same type. Also, a separability condition is required in the utility function. Because these restrictions are more specialized, the result on sufficiency is presented in Supplemental Material Appendix B. Here, we provide the intuition for how Conditions 1 and 2 translate the pairwise stability of links into conditions on the allocation parameters.

Condition 1 relates to expression (i) in the definition of pairwise stability. It is, in fact, equivalent because it requires individuals to be allocated only to network types in their preference classes, and those types must satisfy the same inequalities as in expression (i). Although Condition 1 treats individuals separately, this nevertheless implies that the inequalities hold for any pair of linked individuals. If instead Condition 1 is violated ($\alpha_H(t) > 0$ for some H and $t \notin H$), then there is some positive measure of individuals who would like to drop one of their links, and so the network cannot be pairwise stable.

Condition 2 relates to expression (ii) in the definition of pairwise stability. It establishes that there would be no further mutually beneficial links to add between individuals *who are at a distance greater than $2D$ from each other* in the network. This limitation to distances greater than $2D$ makes our conditions necessary but not sufficient for pairwise stability (except under the restrictions discussed in Supplemental Material Appendix B).

To understand Condition 2, note that there is one such equation for every pair of types (t, s) , including pairs of the same type (t, t) , where the egos of both types have fewer than L links. The other pair of types referred to in the condition, (\bar{t}, \bar{s}) , would be obtained if a link were added between two individuals of types t and s who are greater than $2D$ from each other in the network. For a simple illustration, consider the types t and s shown in Figure 3, which follow Example 1 ($D = 2, L = 2$). If a link were added between two individuals of these types, who were at any distance greater than $2D$, they would be transformed to the types \bar{t} and \bar{s} shown in the figure. This is because the local adjacency matrices for the resulting types would not capture any differences in the resulting structure of the (global) network that depend on the exact distance between the two individuals. These differences would involve loops of lengths greater than $2D + 1$, which are not payoff-relevant and would not appear in the local adjacency matrices that extend only to distance D . Importantly, since the resulting types do not depend on the exact distance between the two individuals (so long as it is greater than $2D$), verifying Condition 2 does not require the global network G .

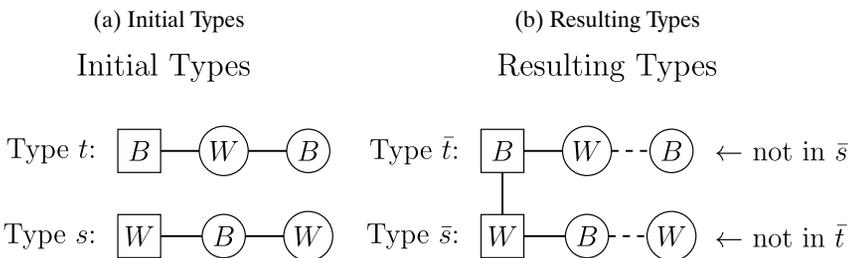


FIGURE 3.—Example of a link added between individuals who are initially distant from each other. *Notes:* The ego of each type is represented with a square node. Dashed lines in the resulting types indicate connections to alters that appear in only one of the new network types, because they are beyond distance $D = 2$ from the ego of the other type.

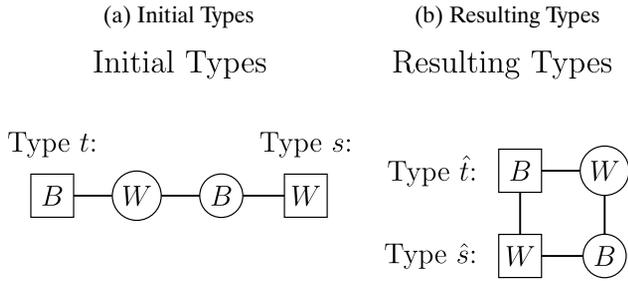


FIGURE 4.—Example of a link added between individuals who are already close in the network.

Nonexisting links between individuals who are $2D$ or less from each other are not considered in Condition 2 because different network types could result if a link were added. For example, if two individuals of the same types t and s from Figure 3 were initially connected at distance 3, as in Figure 4, adding a direct link would transform them to the types \hat{t} and \hat{s} shown in that figure. The utility of those types would potentially be different than for types \bar{t} and \bar{s} (e.g., there is only one friend of a friend in types \hat{t} and \hat{s}). Assessing the stability of nonexisting links such as these would be more complex in our framework because, among other reasons, only certain network types can be located near each other (they must be able to have overlapping alters), and because the placement of individuals with different preference classes becomes important.

Next, note that with a positive measure of individuals of type s in the network, there are infinitely many individuals of type s who are beyond $2D$ from any one individual of type t (this is a consequence of the bounded degree assumption). Any of them could feasibly link with this individual of type t and transform her to type \bar{t} . So if this individual of type t prefers \bar{t} , and a positive measure of individuals of type s prefer \bar{s} as well, the network is unstable. Accordingly, Condition 2 requires that if a positive measure of type t individuals prefer \bar{t} (i.e., $\alpha_{\bar{H}}(t) > 0$ for some \bar{H} where $\bar{t} \in \bar{H}$), there must be zero measure of type s individuals who prefer \bar{s} . Conversely, if a positive measure of type s individuals prefer \bar{s} , there must be zero measure of type t individuals who prefer \bar{t} . Notice that the expression $\mu_{v_1(t)} \sum_{\bar{H} \in \mathcal{H}} P_{\bar{H}|v_1(t)} \alpha_{\bar{H}}(t) 1_{\bar{t} \in \bar{H}}$ gives the total measure of type t individuals who prefer \bar{t} . Hence, this or the analogous expression for the measure of type s individuals who prefer \bar{s} must be zero. Because these measures cannot be negative, the condition that either one or the other measure must be zero is equivalent to requiring that their product be zero, as stated in Condition 2.

Finally, to complete the discussion, we show how these conditions relate to individual preferences in the context of Example 1.

EXAMPLE 1—Continued: Consider the types t and s illustrated in Figure 3. Condition 1 requires that all individuals who are type t in the network have this type in their preference class. This means that $f_{BW} + \varepsilon_1(W) + \nu \geq 0$, because the utility of type t is $f_{BW} + \varepsilon_1(W) + \nu$, and the utility of the type obtained by removing the one link in this type is zero. Hence, Condition 1 requires that all individuals who are type t prefer to have the link to their white friend (who has another friend). Similarly, all individuals who are type s must have preferences in a class where $f_{WB} + \varepsilon_1(B) + \nu \geq 0$, and so they prefer to have the link to their black friend (who has another friend).

Condition 2 checks for individuals of type t who would prefer to be type \bar{t} , or type s who would prefer to be type \bar{s} . The expression inside the first parentheses of the equation

for these types gives the measure of individuals of type t ($\alpha_{\tilde{H}}(t)$) are allocated from each preference class \tilde{H} who also have type \tilde{t} in their preference classes ($1_{\tilde{t} \in \tilde{H}}$). Such individuals would prefer to have a second white friend (who has another friend), because their preferences satisfy $f_{BW} + \nu + \varepsilon_2(W) \geq 0$.¹⁸ The expression inside the second parentheses gives the analogous measure of individuals of type s who also have \tilde{s} in their preference classes. One or the other of these measures must be zero. In other words, either none of the individuals of type t want a second white friend (who has another friend), or none of the individuals of type s want a second black friend (who has another friend), or both. If this condition is violated, there would be individuals of types t and s in the network who would prefer to be types \tilde{t} and \tilde{s} , respectively. These individuals would prefer to add a link with each other, and so the network would not be pairwise stable.

Taken together, Conditions 1 and 2 restrict the preference classes that can be used to generate the predicted type shares for types t and s . Hence, in order to match the observed type shares, a given parameterization must place sufficient probability on the allowable preference classes for these types (i.e., those with type t but not \tilde{t} or those with type s but not \tilde{s}).

The proof of Theorem 1 formalizes the preceding discussion by showing that, if expression (i) holds for all existing links and (ii) holds for all nonexistent links, then Conditions 1 and 2 must be satisfied. One note on this result is that it does not require the existence of equilibrium for every possible parameterization and realization of the variables (recall that nonexistence is possible under pairwise stability). If a particular parameterization cannot generate a pairwise stable network, then there may be no vector of allocation parameters satisfying Conditions 1 and 2. In that case, this parameterization would not be included in the identified set. If *no* parameterization can match the observed type shares while satisfying Conditions 1 and 2, then the identified set would be empty. We would conclude that the observation cannot be an equilibrium outcome under the model as specified, and so we might reject the model. Thus, our framework can be used with models where nonexistence is possible, for example, when links have negative externalities.

6. IMPLEMENTATION

We now describe how to use Theorem 1 to find values of the preference parameters that could be compatible with the observed network. First, we show that the necessary conditions in the theorem can be verified using a quadratic programming (QP) problem.¹⁹ Then, we show how consistency of the estimators for the type shares from a single, large network can be obtained under a sampling approach to inference.

6.1. Formulation as Quadratic Programming Problem

Condition 2 provides a quadratic function of the allocation parameters that must equal zero in equilibrium. Using this to develop an objective function, a QP problem based on Theorem 1 can be defined as follows. The variables in the problem are those allocation

¹⁸The utility of type \tilde{t} is $2f_{BW} + 2\nu + \varepsilon_1(W) + \varepsilon_2(W)$, and removing either link results in type t (note the symmetry of \tilde{t}) with utility $f_{BW} + \nu + \varepsilon_1(W)$. Hence, type \tilde{t} is in a preference class when $f_{BW} + \nu + \varepsilon_2(W) \geq 0$. Similarly, type \tilde{s} is in a preference class when $f_{WB} + \nu + \varepsilon_2(B) \geq 0$.

¹⁹A similar approach to “solving” for the identified set is seen in Honoré and Tamer’s (2006) use of a linear programming problem in a nonlinear panel data model. For a recent interesting example of the use of a quadratic programming problem in economics, see Kitamura and Stoye (2013).

parameters that are not set to zero by Condition 1: $\{\alpha_H(t) : t \in H\}$, or α for short. The objective function derived from Condition 2 is $\alpha^\top Q \alpha$, where the matrix Q is described in detail below. A set of linear constraints impose the requirement that the predicted type shares match the observed shares ($\frac{1}{\mu} \sum_H \mu_{v_1(t)} P_{H|v_1(t)}(\theta) \alpha_H(t) = \pi_t$). There are also adding-up and positivity constraints on the allocation parameters. The QP problem is thus

$$\begin{aligned} \min_{\{\alpha_H(t): t \in H\}} \quad & \alpha^\top Q \alpha \quad \text{subject to:} \\ & \frac{1}{\mu} \sum_H \mu_{v_1(t)} P_{H|v_1(t)}(\theta) \alpha_H(t) = \pi_t, \quad \forall t; \\ & \sum_{t \in H} \alpha_H(t) = 1, \quad \forall H; \quad \alpha_H(t) \geq 0. \end{aligned} \quad (3)$$

As we establish further below, this problem has an optimal value of zero if and only if the conditions of Theorem 1 are satisfied. Therefore, given a vector of preference parameters θ (which produces a probability distribution of preference classes, $\{P_{H|v_1}(\theta)\}$), if a solution can be found yielding a value of zero, that parameter vector belongs in the recovered set.

The assembly of the programming problem is straightforward except for the objective matrix Q , which encodes Condition 2. This is a square matrix that has one row (and column) for each variable in the problem, so the entries of Q correspond to pairs of allocation parameters such as $\alpha_H(t)$ and $\alpha_G(s)$. The entries equal 1 for those pairs that could yield a positive value in the expression for Condition 2, and otherwise equal 0. Specifically, $Q_{[\alpha_H(t), \alpha_G(s)]} = 1_{\bar{t} \in H} \cdot 1_{\bar{s} \in G}$, where \bar{t} and \bar{s} are the types that would result if a link were added between two individuals of types t and s (as defined in the condition). This entry yields the term $(\alpha_H(t) 1_{\bar{t} \in H}) \cdot (\alpha_G(s) 1_{\bar{s} \in G})$ in the objective function $\alpha^\top Q \alpha$. Similarly, the expression in Condition 2 for this pair of allocation parameters includes the term $(\mu_{v_1(t)} P_{H|v_1(t)} \alpha_H(t) 1_{\bar{t} \in H}) \cdot (\mu_{v_1(s)} P_{G|v_1(s)} \alpha_G(s) 1_{\bar{s} \in G})$. Hence, as long as $\mu_{v_1(\cdot)}$ and $P_{\cdot|\cdot}$ are strictly positive, the former can be used to assess whether the latter is nonzero.²⁰

The example below illustrates the matrix Q in the simple model from Section 3.

EXAMPLE 2: Let preferences be given by (1), with $D = 1$, $L = 1$, and $\mathcal{X} = \{B, W\}$. Here, network types can be described using just the vectors of characteristics $v = (v_1, v_2)$. Preference classes can be enumerated as follows: $H_1 = \{(B, 0)\}$, $H_2 = \{(B, 0), (B, B)\}$, $H_3 = \{(B, 0), (B, W)\}$, $H_4 = \{(B, 0), (B, B), (B, W)\}$, $H_5 = \{(W, 0)\}$, $H_6 = \{(W, 0), (W, W)\}$, $H_7 = \{(W, 0), (W, B)\}$, and $H_8 = \{(W, 0), (W, W), (W, B)\}$. After excluding allocation parameters set to zero by Condition 1, there are 16 remaining: $\alpha_1(B, 0)$; $\alpha_2(B, 0)$, $\alpha_2(B, B)$; $\alpha_3(B, 0)$, $\alpha_3(B, W)$; $\alpha_4(B, 0)$, $\alpha_4(B, B)$, $\alpha_4(B, W)$; $\alpha_5(W, 0)$; $\alpha_6(W, 0)$, $\alpha_6(W, W)$; $\alpha_7(W, 0)$, $\alpha_7(W, B)$; $\alpha_8(W, 0)$, $\alpha_8(W, W)$, and $\alpha_8(W, B)$ (the subscripts correspond to the subscripts of the preference classes). The vector α collects these allocation parameters.

The matrix Q is 16×16 , and its rows and columns correspond to the allocation parameters as listed in α (the matrix itself is shown in Figure D.1 in the Supplemental Material). The matrix is symmetric, sparse, and indefinite. The first row corresponds to $\alpha_1(B, 0)$.

²⁰An expression identical to Condition 2 would be obtained if the entries of Q included the measures and probabilities, that is, if $Q_{[\alpha_H(t), \alpha_G(s)]} = (\mu_{v_1(t)} P_{H|v_1(t)} 1_{\bar{t} \in H}) \cdot (\mu_{v_1(s)} P_{G|v_1(s)} 1_{\bar{s} \in G})$. However, having binary entries saves memory and potentially avoids recomputing Q at each candidate parameter vector.

All entries in that row are zero because the preference class associated with that allocation parameter, $H_1 = \{(B, 0)\}$, contains only the isolated type with a black ego. Hence, $1_{\bar{t} \in H_1} = 0$ for any type \bar{t} that could be obtained by adding a link. There are nonzero entries in six rows (and columns) of Q : those corresponding to allocation parameters $\alpha_2(B, 0)$, $\alpha_3(B, 0)$, $\alpha_4(B, 0)$, $\alpha_6(W, 0)$, $\alpha_7(W, 0)$, and $\alpha_8(W, 0)$. These parameters all indicate isolated individuals who would prefer to have a friend. (Because $L = 1$, only isolated individuals can add a link.) For example, in the row corresponding to $\alpha_3(B, 0)$, there are 1's in the columns corresponding to $\alpha_7(W, 0)$ and $\alpha_8(W, 0)$. The preference classes associated with these parameters are $H_3 = \{(B, 0), (B, W)\}$, $H_7 = \{(W, 0), (W, B)\}$, and $H_8 = \{(W, 0), (W, W), (W, B)\}$, respectively. To denote the types, let $t = (B, 0)$ (isolated black) and $s = (W, 0)$ (isolated white), and let $\bar{t} = (B, W)$ (black with white best friend) and $\bar{s} = (W, B)$ (white with black best friend), which are the types that would result if a link were added between two individuals of types t and s . Thus, we have $1_{\bar{t} \in H_3} \cdot 1_{\bar{s} \in H_7} = 1$ and $1_{\bar{t} \in H_3} \cdot 1_{\bar{s} \in H_8} = 1$.

The construction of Q in the example above is fairly simple because it is feasible to evaluate the expression $1_{\bar{t} \in H} \cdot 1_{\bar{s} \in G}$ for each entry individually. Given the corresponding pair of allocation parameters $\alpha_H(t)$ and $\alpha_G(s)$, the types t and s determine the types \bar{t} and \bar{s} that would result if a link were added, and it is then easy to check whether \bar{t} and \bar{s} are contained in H and G , respectively. However, for larger matrices it may be too burdensome to loop through the entries individually. Instead, we suggest first constructing a preliminary matrix S , with the same dimensions and organization as Q , whose entries $S_{[\alpha_H(t), \alpha_G(s)]}$ are defined as $1_{\bar{t} \in H}$ rather than $1_{\bar{t} \in H} \cdot 1_{\bar{s} \in G}$. Conceptually, the difference between S and Q is that the entries of S reflect only the preference class associated with the allocation parameter in the row, rather than those of both the row and the column. This makes it faster to construct S because operations can be applied row-by-row rather than element-by-element. The matrix Q then simply equals the Hadamard (i.e., entrywise) product of S with its transpose: $Q = S \circ S^T$. (See Supplemental Material Appendix D.7.2 for further details.)

With the objective matrix Q defined as above, we can establish the following result. (The proof appears in Appendix A.)

THEOREM 2: *Given a probability distribution of preference classes $\{P_{H|v_1}(\theta)\}$, there exists a vector of allocation parameters α yielding type shares $\{\pi_i\}$ while satisfying Conditions 1 and 2 if and only if the optimal value of QP problem (3) is zero.*

This theorem provides a computational avenue to implement our approach. Because this pertains to identification, however, the population type shares are assumed to be known. In order to accommodate data from a finite sample, the QP problem must be modified to allow for some error in the estimated shares. The approach we take is to add slack variables that define fixed “bands” around the type shares, the width of which are a function of the sample size.²¹ The modified QP problem then verifies whether a structural parameter vector can yield a prediction within these bands while satisfying Conditions 1 and 2.

²¹See Supplemental Material Appendix D.6. Given a sampling distribution for the vector of estimated type shares, one could instead define the bands to contain a 95% confidence set. This would be a computationally efficient means to incorporate statistical uncertainty, as discussed in more general terms in Supplemental Material Appendix C.

Finally, we note that the objective function $\alpha^\top Q\alpha$ may not be convex because, while the matrix Q is symmetric, it may be indefinite, as is the case in the example above. This rules out some standard QP solvers, but more general constrained nonlinear optimization routines can be used instead.²² In the simulation exercises in Section 7, we find that the problem solves quickly using an active set algorithm in the program KNITRO. Importantly, because the optimal value is known (i.e., $\alpha^\top Q\alpha = 0$), it is trivial to ascertain that a global rather than local optimum has been reached. On the other hand, one must exercise caution so that positive local minima do not erroneously lead to dismissal of a parameter value.

6.2. Consistency of Type Shares

The parameter of interest in our setup is the vector that characterizes the payoff structure, θ , for a given large network. The main insight from Theorems 1 and 2 is that, *knowing the type shares* π , we have a mapping that yields the identified set for θ . We now briefly discuss inference and show how it is possible to obtain consistent estimates of π for this mapping.

Assuming that we do not observe the full network (and hence do not know π), we use a sampling approach to inference whereby we maintain that the observed data are a simple random sample from this full network that records the types only. The central question in this sampling approach is how close the sampled quantities are to the true quantities that can be obtained if we had access to the full network. Here, the possibility of inverting from a sample to gain information on the population depends crucially on the sampling method employed. In our case, we maintain the assumption that we have a random sampling scheme where individuals are drawn independently (conditional on the realized network) and their types are recorded, which requires registering the features of their neighboring connections. The snowball sampling procedure defined in Goodman (1961), which starts with an initial random sample and interviews connections up to a specified distance, is an example of such a sampling scheme.²³ The only source of variation here is the random sampling, which then leads to statements about how close our estimate $\hat{\pi}$ is to π .²⁴

Hence, while we do not know π , we can estimate it from appropriately sampled individuals in the network. As indicated in Section 2, the equilibrium graphing used in our theoretical analysis is an approximation to a large but nevertheless finite network, so here we consider a population network that is finite. Suppose we have data recording the network types t_i of a sample of n individuals. A sample analogue of the proportion of types is then $\hat{\pi}_t = \frac{1}{n} \sum_i 1[t_i = t]$. Accordingly, let $\hat{\pi} \equiv (\hat{\pi}_t)_{t \in \mathcal{T}}$, which is the vector of estimated type shares. Given that the population is finite, as n increases (i.e., as we are drawing a larger sample), it is then simple to see that $\hat{\pi}$ will converge to π . We state this result in the proposition below.²⁵

²²Quadratic programs that are not convex are *NP*-hard, and our model may actually be an example of such a problem. Alternative relaxation schemes are nonetheless available in the numerical optimization literature to handle such cases. Another computationally simple approach that gets a superset of the identified set is to base inference on only Condition 1 for stability, which will yield a linear program.

²³The term “snowball sampling” has different uses. See Handcock and Gile (2011) for a helpful discussion.

²⁴This relates to what is sometimes referred to as a “design-based” paradigm in statistics and goes back to at least Neyman (1934). The randomness here obtains from the probability ascribed by the survey scheme to the sampling of the various individuals in the network. In applications where the full network is observed (e.g., with administrative data), this approach would imply there is no statistical uncertainty.

²⁵For a primer on the sampling approach to inference, see Cochran (1977). For sampling approaches with network data, see Kolaczyk (2009). Similar treatments using random sampling and relating finite and infinite

PROPOSITION 1: *Under random sampling (on the underlying, realized network) and as sample size increases, we have $\hat{\pi} \xrightarrow{p} \pi$.*

Supplemental Material Appendix C contains the proof and elaborates on our approach in a simple example.²⁶

7. SIMULATIONS

We now present two simulation exercises using examples of utility specification (1). The main purpose is to illustrate the performance of our approach, in terms of the parameter sets that are recovered and the computational burden that is involved. Additionally, the procedures described here and in Supplemental Material Appendix D provide some guidance on further aspects of the implementation, such as how to generate the sets of network types and preference classes, and details of the construction of the objective matrix Q .

7.1. Model of Best Friendships

The first exercise uses the simple model of best friendships from Section 3, where $D = 1$, $L = 1$, and $\mathcal{X} = \{B, W\}$. Here, network types can be represented using just the vector of characteristics $v = (x, y)$, where x is the race of the ego and y is the race of the alter (or 0 if no alter). The utility function simplifies to $u_i(x, y) = f_{xy} + \varepsilon_i(y)$, with four preference parameters, f_{xy} , $x, y \in \{B, W\}$, and two preference shocks, $\varepsilon_i(B)$ and $\varepsilon_i(W)$.

The parameters are set at $f_{BB} = 0.4$, $f_{BW} = 0.2$, $f_{WB} = 0.15$, $f_{WW} = 0.5$, with population sizes $\mu_B = 1.0$, $\mu_W = 1.2$. The shocks have a uniform distribution ($U[-1, 0]$) so that the parameters can be interpreted as probabilities. We then use a microsimulation procedure to generate a single pairwise stable network with a large number of individuals (500 blacks and 600 whites).²⁷ From this network we extract the type shares, which are plotted in Figure 5. The figure also shows *all* possible type shares from pairwise stable networks under this parameterization, which demonstrates the full range of equilibria.²⁸

Given the type shares from the simulated network, we then search for parameter vectors where the QP problem can be minimized to zero. Figure 6 shows the resulting identified set. The values of the cross-race preference parameters (f_{BW} and f_{WB}) are unbounded from above, and together they display a Leontief pattern: if tastes on one side of the

populations that we follow also here are given in Imbens and Rubin (2015) and references therein (see Chapter 6). Also, ARD data (Aggregated Relational Data) is an approach to collecting exactly sampled information (types) from a subset of nodes in a network. This ARD sampling approach to recovering full network properties from a random sample of nodes fits with our sampling approach (for the latest in econometrics on ARD work, see Breza, Chandrasekhar, McCormick, and Pan (2017)).

²⁶There are alternative approaches to inference that take the repeated sampling approach whereby one observes one large network that is a draw from a population distribution of networks. There, large sample limit theory (where randomness is due to repeated sampling from this population distribution) is subtle due to the built-in correlations that result from interactions. For some recent advances, see Leung (2015) and Lee and Song (2017).

²⁷This procedure is described in detail for the second exercise, in Supplemental Material Appendix D.4.

²⁸In this example, the QP problem simplifies so that it is trivial to verify whether a vector of preference parameters and a vector of type shares are compatible with each other. This makes it easy to find *all* the equilibrium vectors of type shares for a given vector of structural parameters. See Supplemental Material Appendix D.2 for details.

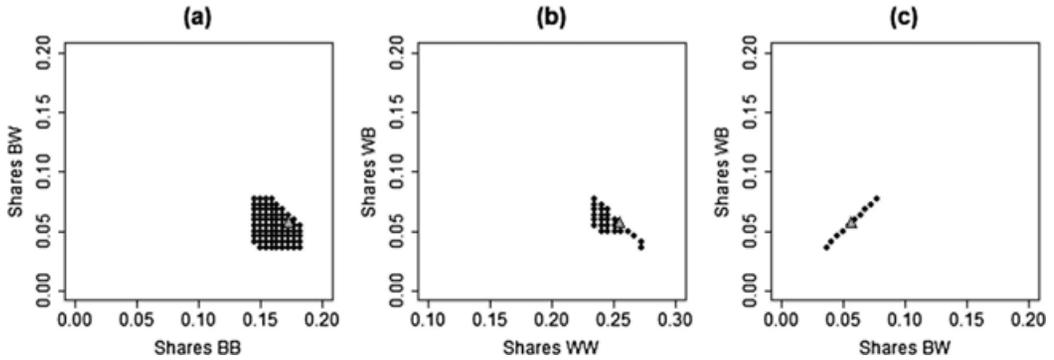


FIGURE 5.—Equilibrium type shares in first simulation exercise. *Notes:* Points illustrate the full set of equilibrium type shares under parameter values $f_{BB} = 0.40$, $f_{BW} = 0.20$, $f_{WB} = 0.15$, and $f_{WW} = 0.50$. Triangles indicate the type shares from one finite network simulation, with values $\pi_{(B,B)} = 0.173$, $\pi_{(B,W)} = 0.056$, $\pi_{(W,B)} = 0.056$, and $\pi_{(W,W)} = 0.255$.

market constrain the number of cross-race linkages, tastes on the other side could be unbounded. By contrast, the identified ranges for the own-race preference parameters (f_{BB} and f_{WW}) are small and informative: 0.38–0.51 for blacks and 0.46–0.57 for whites. However, the set would not provide conclusive evidence on the ranking of these parameters (i.e., $f_{WW} > f_{BB}$).

Finally, for this model, we are able to confirm the results obtained above and explore how the identified set changes with additional observations, using a simplification of the QP problem that can be derived in this particular case (see Supplemental Material Appendix D.2). With the simplification, it is trivial to verify that the optimal value is zero, and so we can evaluate a large grid of parameter vectors almost instantaneously. Figure D.2 in the Supplemental Material shows that the identified set matches what we obtained with the search procedure above. We then evaluate four other vectors of type shares, selected at random from the set of all equilibria shown in Figure 5. Qualitatively similar identified sets are recovered from each of these observations. (See Figure D.3 for the vectors of type shares, and Figures D.4 and D.5 for the identified sets.) If we were to observe all four of the randomly selected networks, the identified set would be the intersection of the sets recovered with each network because the parameters must be able to predict all four vectors of type shares as equilibrium outcomes. The resulting identified set is quite precise (see Figure D.6). Only one parameter vector in the grid we evaluate is included ($f_{BB} = 0.40$, $f_{BW} = 0.20$, $f_{WB} = 0.16$, $f_{WW} = 0.50$). Since the grid uses intervals of 0.02, this means that the identified set spans less than 0.04 in each dimension. This indicates substantial identifying power from the observation of multiple large networks.

7.2. General Friendship Model

Next, we consider a more complete version of specification (1), where $D = 2$ and $L = 3$. To motivate the latter, we note that fewer than five percent of the students in the Add Health study have more than three same-sex friends, based on reciprocated nominations. The predetermined characteristics are the same as before: black or white race, $\mathcal{X} = \{B, W\}$.

The preference parameters ($\theta \equiv (f_{BB}, f_{BW}, f_{WB}, f_{WW}, \nu, \omega)'$) are chosen to generate a degree distribution that is similar to the observed distribution in Add Health, and the values of ν and ω also satisfy Assumption 3 so that the sharp identified set is obtained

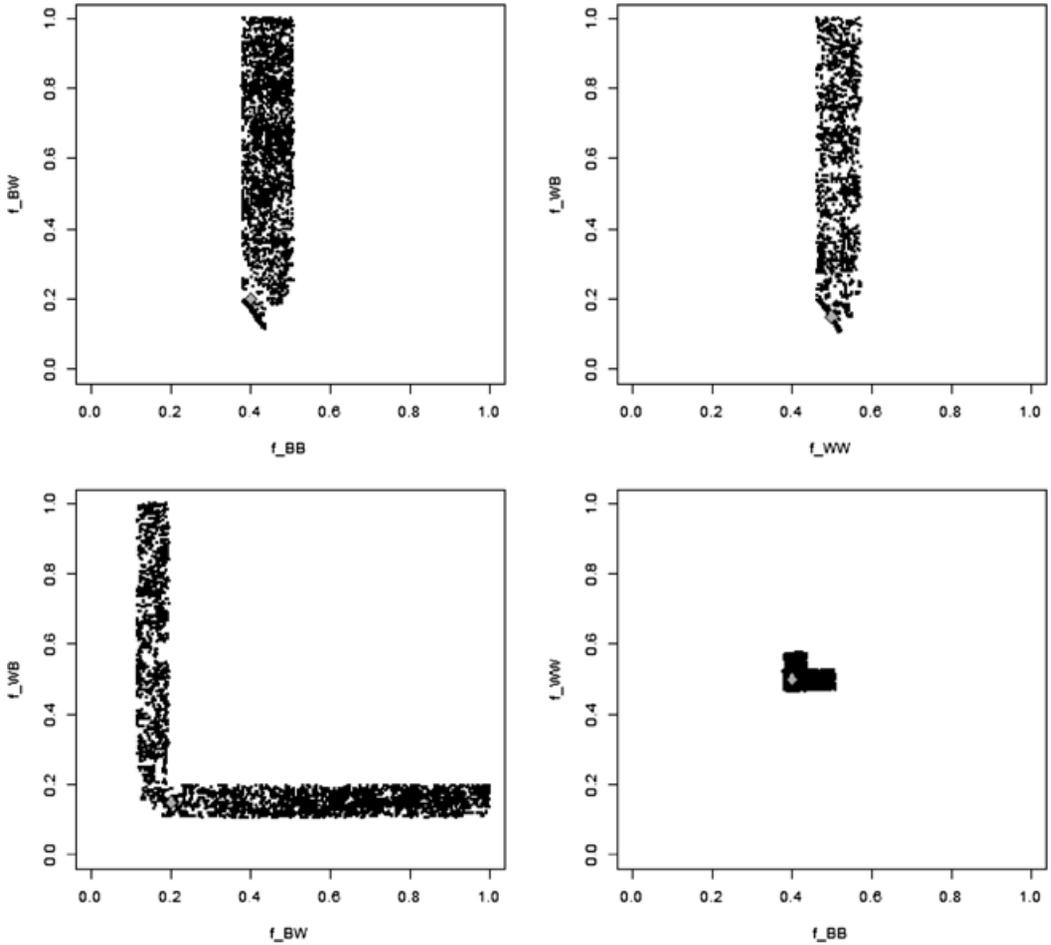


FIGURE 6.—Identified set from one finite network in first exercise. *Notes:* Points illustrate the identified set obtained using type shares from one finite network (shown with triangles in Figure 5). Diamonds indicate true parameter values: $f_{BB} = 0.40$, $f_{BW} = 0.20$, $f_{WB} = 0.15$, and $f_{WW} = 0.50$.

here (see Supplemental Material Appendix B).²⁹ The six preference shocks are drawn independently from a standard normal distribution, but they are assigned to links in descending order (within alter characteristic). That way, the particular convention we use to select the representation, (A, v) , for each network type does not impact their utility (see Supplemental Material Appendix D.3).

To generate the data, we simulate a number of finite networks with $n = 500$ individuals (100 blacks and 400 whites, reflecting $\mu_B/\mu_W = 1/4$), using a procedure described in Supplemental Material Appendix D.4. Figure 7 plots the shares of certain combinations of types in these simulated networks, to illustrate the variation that can arise for a fixed vector of preference parameters but with different realizations of shocks and equilibria. One network, selected at random, serves as the observation we use to recover the identified

²⁹The parameter values are $(f_{BB}, f_{BW}, f_{WB}, f_{WW}) = (-0.9, -1.5, -1.7, -0.7)$, $v = 0.2$, and $\omega = 0.2$. Figure D.7 in the Supplemental Material shows the degree distributions. The average degree in the simulated networks matches the average degree for same-sex friendships in Add Health.

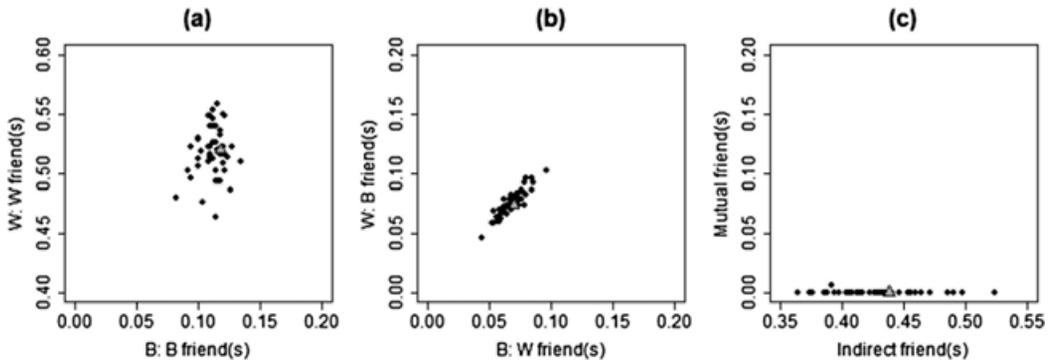


FIGURE 7.—Equilibrium type shares in second simulation exercise. *Notes:* Figure plots shares of certain combinations of types: types with any own-race friends (x : x friend(s)), any opposite-race friends (x : y friend(s)), any indirect friends, and any mutual friends. Points represent the shares from different simulated networks, and triangles indicate the shares from the network randomly selected to use as the observation.

set (indicated with triangles in Figure 7). Supplemental Material Appendices D.6 to D.8 describe the specific procedures we then use to formulate and solve the QP problem (for a given parameter vector) and to search through the parameter space. One key point is that, to save memory and improve computational speed, we only consider network types that are either observed in the data or adjacent to an observed type (i.e., they can be reached via addition or deletion of one link). The search procedure uses Markov Chain Monte Carlo (MCMC) algorithms (Supplemental Material Appendix D.8).

Projections of the identified set are shown in Figure 8. The identified range for f_{WB} appears to be unbounded from above, as in the previous example, while f_{BW} is bounded in both directions. Also, for both blacks and whites, we would not be able to conclude that there is a preference for same-race over different-race friends (i.e., the fact that $f_{BB} > f_{BW}$ and $f_{WW} > f_{WB}$), as there are points in the identified set where the opposite holds. On the other hand, the values of same-race friendships (f_{BB} and f_{WW}) again have fairly tight ranges. Furthermore, if more than one network were observed, the identified set would be even smaller, as seen in the previous example.³⁰ The identified range of ν (friends of friends) seems reasonably informative, as it spans less than one standard deviation of the preference shocks, and its sign could be correctly inferred (the minimum identified value is 0.055). The parameter ω (mutual friends) is not shown. The data do not provide any information on its value, mainly because there are no mutual friendships in the observed network. The identified range for ω is consequently unbounded from below, and the upper bound ($\omega \leq \frac{3}{2}\nu$) follows from Assumption 3 (see Supplemental Material Appendix B).

Importantly, the computational burden involved in this exercise indicates that our approach is feasible for empirically realistic models. The average time required to evaluate a candidate vector of structural parameters was just less than 30 seconds (see Supplemental Material Appendix D.8). Furthermore, the MCMC search process is trivial to parallelize by running multiple chains simultaneously, so richer models with potentially longer evaluation times (say 5 to 15 minutes) would remain tractable. Most of the computational burden (80% of the total compute time) comes from solving the QP problem, which also

³⁰In preliminary results from a simulation using two observations, the identified range for f_{WB} is bounded and is roughly similar in size to the identified range for f_{BW} , for example.

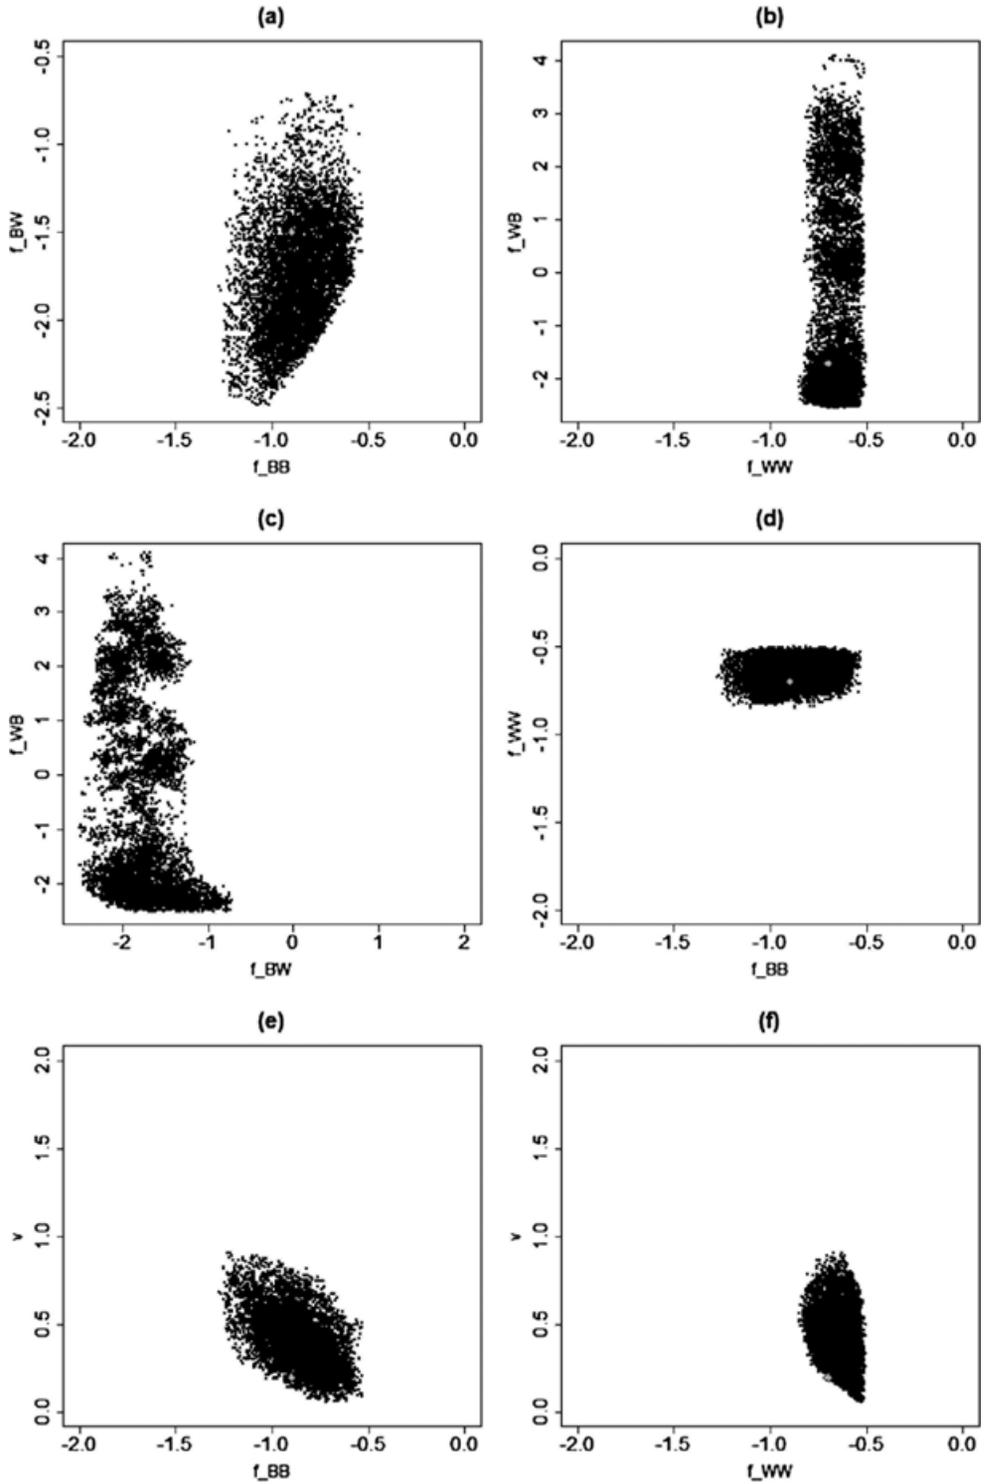


FIGURE 8.—Identified set in second exercise. *Notes:* Diamonds indicate true parameter values: $f_{BB} = -0.9$, $f_{BW} = -1.5$, $f_{WB} = -1.7$, $f_{WW} = -0.7$, and $\nu = 0.2$.

suggests that additional gains in performance may be possible with further advances in solution algorithms.

8. CONCLUSION

We conclude with a brief discussion about empirical settings where our framework might apply, and where it might not. The main consideration is whether the assumption of bounded degree is reasonable in a particular context. The maximum degree and the degree distribution are, of course, observable in a given data set, and, depending on the sampling scheme (and the model), one can determine whether a specific bound is reasonable to impose. For example, in the social and financial networks studied in [Jackson, Rodriguez-Barraquer, and Tan \(2012\)](#), a negligible portion of the sample reached the caps of 5 or 8 nominations for each type of relationship. This suggests that the bounded degree assumption may be reasonable in analyses of informal insurance as well as close friendships. On the other hand, bounded degree would not apply in online social networks. Nor would it apply in certain markets for intermediate goods or financial networks where firms with many connections (i.e., “hubs”) are observed. Economically, the difference between settings with bounded degree distributions versus approximately power law distributions may relate to the cost of links. If there are substantial fixed costs, relative to a finite budget of time, for example, individuals would be limited in the number of links they can maintain.

APPENDIX A: PROOFS OF THEOREMS 1 AND 2

A.1. *Theorem 1*

PROOF: Given a pairwise stable network G , predetermined characteristics \mathbf{X} , and preference shocks ε (for all the individuals in the network), partition individuals by their preference class H and define $\alpha_H(t)$ as the fraction of individuals in this pairwise stable network in preference class H who are of network type t . This allocation yields the observed proportions of network types; that is, $(\mu_{v_1(t)}/\mu) \sum_H P_{H|v_1(t)} \alpha_H(t) = \pi_t$.

If the network G is pairwise stable, then all existing links satisfy expression (i) in Definition 1. Hence, for any individual i , whose network type is characterized as (A, v) , we have $u(A, v; \varepsilon_i) \geq u(A_{-l}, v; \varepsilon_i), \forall l = 1, \dots, L$. By definition, this type is in the individual’s preference class. Hence, $\alpha_H(t) > 0$ only if $t \in H$, and so Condition 1 is satisfied.

To show that Condition 2 is satisfied, consider an arbitrary pair of types (t, s) where both types have fewer than the maximum L links. Let (\bar{t}, \bar{s}) be the pair of types that would result if a link were added between two individuals of types t and s who are greater than $2D$ from each other. Suppose that the measure of individuals of type t who prefer \bar{t} is positive: $\mu_{v_1(t)} \sum_{\bar{H} \in \mathcal{H}} P_{\bar{H}|v_1(t)} \alpha_{\bar{H}}(t) 1_{i \in \bar{H}} > 0$. Hence, there is an infinite number of individuals of type t who would prefer to add links to some individuals of type s that are greater than $2D$ away from them. These individuals of type t would indeed prefer to add links to *any* individuals of type s that are greater than $2D$ away, because adding a link to any distant individual of type s would yield the same utility for a given individual of type t . Also, because degree is bounded, only a finite number of individuals of type t can be within distance $2D$ from any given individual of type s . Hence, every individual of type s faces individuals of type t who are greater than $2D$ away and who want to add a link. Because all nonexisting links satisfy expression (ii) in a pairwise stable network, the utility of type \bar{s} must therefore be lower than the utility of type s for *all* individuals of type s . Because type s can be obtained by dropping one of the links in type \bar{s} , this means that \bar{s} is not contained

in the preference class of any individuals of type s in the network. Therefore, the measure $\mu_{v_1(s)} \sum_{\check{H} \in \mathcal{H}} P_{\check{H}|v_1(s)} \alpha_{\check{H}}(s) 1_{\check{s} \in \check{H}}$ is zero. Thus, at least one of the measures expressed in the equation for types (t, s) in Condition 2 must be zero, which gives us the condition. *Q.E.D.*

A.2. Theorem 2

PROOF: Condition 2 is satisfied if and only if the objective function is equal to zero. This is because, as long as $\mu_{v_1(\cdot)}$ and $P_{\cdot|}$ are strictly positive,

$$\begin{aligned} \mu_{v_1(t)} \mu_{v_1(s)} \sum_{\check{H} \in \mathcal{H}} \sum_{\check{H} \in \mathcal{H}} P_{\check{H}|v_1(t)} P_{\check{H}|v_1(s)} \alpha_{\check{H}}(t) \alpha_{\check{H}}(s) 1_{\check{t} \in \check{H}} 1_{\check{s} \in \check{H}} &= 0 \\ \Leftrightarrow \sum_{\check{H} \in \mathcal{H}} \sum_{\check{H} \in \mathcal{H}} \alpha_{\check{H}}(t) \alpha_{\check{H}}(s) 1_{\check{t} \in \check{H}} 1_{\check{s} \in \check{H}} &= 0. \end{aligned}$$

The first set of constraints in (3) is to match the observed proportions of network types (π_t) . The second and third sets of constraints simply require that allocations from a given preference class add up to 1 and are nonnegative. Finally, Condition 1 is encoded by the fact that allocation parameters are only defined for the types in each preference class (i.e., the variables in the problem are $\{\alpha_H(t) : t \in H\}$, not all $\{\alpha_H(t)\}$). Hence, from each preference class, there are no allocations made to types not in that preference class. *Q.E.D.*

REFERENCES

- BADEV, A. (2013): “Discrete Games in Endogenous Networks: Theory and Policy,” Working Paper, University of Pennsylvania. [265]
- BLOCH, F., AND M. JACKSON (2006): “Definitions of Equilibrium in Network Formation Games,” *International Journal of Game Theory*, 34, 305–318. [268]
- BOUCHER, V., AND MOURIFIÉ (2017): “My Friend Far Far Away: Asymptotic Properties of Pairwise Stable Networks,” *Econometrics Journal*, 20 (3), S14–S46. [265]
- BREZA, E., A. G. CHANDRASEKHAR, T. H. MCCORMICK, AND M. PAN (2017): “Using Aggregated Relational Data to Feasibly Identify Network Structure Without Network Data,” Technical Report, National Bureau of Economic Research. [281]
- CHANDRASEKHAR, A. G., AND M. O. JACKSON (2014): “A Network Formation Model Based on Subgraphs,” Working Paper 20276, NBER. [265]
- CHARBONNEAU, K. (2017): “Multiple Fixed Effects in Binary Response Panel Data Models,” *Econometrics Journal*, 20 (3), S1–S13. [265]
- CHIAPPORI, P.-A., A. GALICHON, AND B. SALANIÉ (2016): “On Human Capital and Team Stability,” Working Paper, Columbia University. [265]
- CHOO, E., AND A. SIOW (2006): “Who Marries Whom and Why,” *Journal of Political Economy*, 114 (1), 175–201. [265,266]
- CHRISTAKIS, N., J. FOWLER, G. IMBENS, AND K. KALIANARAMAN (2010): “An Empirical Model of Strategic Network Formation,” Working Paper, Harvard University. [264,265,267]
- COCHRAN, W. G. (1977): *Sampling Techniques*. New York: Wiley. [280]
- CURRARINI, S., M. JACKSON, AND P. PIN (2009): “An Economic Model of Friendship: Homophily, Minorities and Segregation,” *Econometrica*, 77 (4), 1003–1045. [264,266]
- DE PAULA, A. (2017): “Econometrics of Network Models,” in *Advances in Economics and Econometrics: Eleventh World Congress*, ed. by B. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson. Econometric Society Monographs. Cambridge: Cambridge University Press, 268–323. [264]
- DE PAULA, A., S. RICHARDS-SHUBIK, AND E. TAMER (2017): “Supplement to ‘Identifying Preferences in Networks With Bounded Degree,’” *Econometrica Supplemental Material*, 85, <http://dx.doi.org/10.3982/ECTA13564>. [269]
- DZEMSKI, A. (2014): “An Empirical Model of Dyadic Link Formation in a Network With Unobserved Heterogeneity,” Working Paper, University of Manheim. [265]

- ECHENIQUE, F., S. LEE, AND M. SHUM (2010): "Aggregate Matchings," Working Paper, Caltech. [265]
- FOX, J. T. (2010): "Identification in Matching Games," *Quantitative Economics*, 1 (2), 203–254. [265]
- (2018): "Estimating Matching Games With Transfers," *Quantitative Economics* (forthcoming). [265]
- GALICHON, A., AND B. SALANIE (2009): "Matching With Trade-Offs: Revealed Preferences Over Competing Characteristics," Working Paper, Ecole Polytechnique. [265,266]
- GOLDSMITH-PINKHAM, P., AND G. IMBENS (2013): "Social Networks and the Identification of Peer Effects," *Journal of Business and Economic Statistics*, 31 (3), 253–264. [267]
- GOODMAN, L. A. (1961): "Snowball Sampling," *The Annals of Mathematical Statistics*, 32 (1), 148–170. [280]
- GRAHAM, B. S. (2015): "Methods of Identification in Social Networks," *Annual Review of Economics*, 7 (1), 465–485. [264]
- (2017): "An Econometric Model of Network Formation With Degree Heterogeneity," *Econometrica*, 85 (4), 1033–1063. [265]
- HANDCOCK, M. S., AND K. J. GILE (2011): "Comment: On the Concept of Snowball Sampling," *Sociological Methodology*, 41 (1), 367–371. [280]
- HONORÉ, B. E., AND E. TAMER (2006): "Bounds on Parameters in Panel Dynamic Discrete Choice Models," *Econometrica*, 74 (3), 611–629. [277]
- IMBENS, G. W., AND D. B. RUBIN (2015): *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge: Cambridge University Press. [281]
- JACKSON, M. (2009): *Social and Economic Networks*. Princeton, NJ: Princeton University Press. [265,268]
- JACKSON, M. O., AND A. WOLINSKY (1996): "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 71 (1), 44–74. [263,265,268]
- JACKSON, M. O., T. RODRIGUEZ-BARRAQUER, AND X. TAN (2012): "Social Capital and Social Quilts: Network Patterns of Favor Exchange," *American Economic Review*, 102 (5), 1857–1897. [266,286]
- KALAI, E. (2004): "Large Robust Games," *Econometrica*, 72 (6), 1631–1665. [266]
- KHAN, M. A., AND Y. SUN (2002): "Non-Cooperative Games With Many Players," in *Handbook of Game Theory With Economic Applications*, Vol. 3, ed. by R. Aumann and S. Hart. Amsterdam: Elsevier, Chapter 46, 1761–1808. [265]
- KITAMURA, Y., AND J. STOYE (2013): "Nonparametric Analysis of Random Utility Models: Testing," Working Paper CWP36/13, CeMMAP. [277]
- KOLACZYK, E. (2009): *Statistical Analysis of Network Data*. New York: Springer. [274,280]
- LEE, J. H., AND K. SONG (2017): "Stable Limit Theorems for Empirical Processes Under Conditional Neighborhood Dependence," Working Paper, UBC. [281]
- LEUNG, M. (2015): "A Random-Field Approach to Inference in Large Models of Network Formation," Working Paper, Stanford University. [265,281]
- LOVASZ, L. (2012): *Large Networks and Graph Limits*. Providence, RI: American Mathematical Society. [265]
- MELE, A. (2017): "A Structural Model of Dense Network Formation," *Econometrica*, 85 (3), 825–850. [264, 265,267]
- MENZEL, K. (2015): "Large Matching Markets as Two-Sided Demand Systems," *Econometrica*, 83 (3), 897–941. [265]
- (2016): "Inference for Games With Many Players," *Review of Economic Studies*, 83, 306–337. [266]
- MIYAUCHI, Y. (2016): "Structural Estimation of a Pairwise Stable Network With Nonnegative Externality," *Journal of Econometrics*, 195 (2), 224–235. [265,267]
- NEYMAN, J. (1934): "On the Two Different Aspects of the Representative Method: The Method of Stratified Sampling and the Method of Purposive Selection," *Journal of the Royal Statistical Society*, 97, 558–625. [280]
- SHENG, S. (2014): "A Structural Econometric Analysis of Network Formation Games," Working Paper, UCLA. [265,267]

Co-editor Liran Einav handled this manuscript.

Manuscript received 23 June, 2015; final version accepted 3 August, 2017; available online 21 August, 2017.