

Identifying Network Ties from Panel Data: Theory and an Application to Tax Competition*

Áureo de Paula

Imran Rasul

Pedro CL Souza[†]

October 2023

Abstract

Social interactions determine many economic behaviors, but information on social ties does not exist in most publicly available and widely used datasets. We present results on the identification of social networks from observational panel data that contains no information on social ties between agents. In the context of a canonical social interactions model, we provide sufficient conditions under which the social interactions matrix, endogenous and exogenous social effect parameters are globally identified if networks are constant over time. We also provide an extension of the method for time-varying networks. We then describe how high-dimensional estimation techniques can be used to estimate the interactions model based on the Adaptive Elastic Net Generalized Method of Moments. We employ the method to study tax competition across US states. The identified social interactions matrix implies that tax competition differs markedly from the common assumption of competition between geographically neighboring states, providing further insights into the long-standing debate on the relative roles of factor mobility and yardstick competition in driving tax setting behavior across states. Most broadly, our identification and application show that the analysis of social interactions can be extended to economic realms where no network data exists.

JEL Classification: C31, D85, H71.

*We gratefully acknowledge financial support from the ESRC through the Centre for the Microeconomic Analysis of Public Policy (ES/T014334/1), the Centre for Microdata Methods and Practice (RES-589-28-0001) and the Large Research Grant ES/P008909/1 and from the ERC (SG338187). We thank Edo Airolidi, Luis Alvarez, Michele Aquaro, Oriana Bandiera, Larry Blume, Yann Bramoullé, Stephane Bonhomme, Vasco Carvalho, Gary Chamberlain, Andrew Chesher, Christian Dustmann, Sérgio Firpo, Jean-Pierre Florens, Eric Gautier, Giacomo de Giorgi, Matthew Gentzkow, Stefan Hoderlein, Bo Honoré, Matt Jackson, Dale Jorgensen, Christian Juliard, Maximilian Kasy, Miles Kimball, Thibaut Lamadon, Simon Sokbae Lee, Arthur Lewbel, Tong Li, Xiadong Liu, Elena Manresa, Charles Manski, Marcelo Medeiros, Angelo Mele, Francesca Molinari, Pepe Montiel, Andrea Moro, Whitney Newey, Ariel Pakes, Eleonora Pattachini, Michele Pelizzari, Martin Pesendorfer, Christiern Rose, Adam Rosen, Bernard Salanie, Olivier Scaillet, Sebastien Sieglösch, Pasquale Schiraldi, Tymon Sloczynski, Kevin Song, John Sutton, Adam Szeidl, Thiago Tachibana, Elie Tamer, and seminar and conference participants for valuable comments. We also thank Tim Besley and Anne Case for comments and sharing data. Daniel Barbosa provided outstanding research assistance. A previous version of this paper was circulated as “Recovering Social Networks from Panel Data: Identification, Simulations and an Application.” All errors remain our own. Codes are available on Zenodo (<https://zenodo.org/XXX>) and github (<https://github.com/YYY>) repositories.

[†]de Paula: University College London, CeMMAP and IFS, a.paula@ucl.ac.uk; Rasul: University College London and IFS, i.rasul@ucl.ac.uk; Souza: Queen Mary University, p.souza@qmul.ac.uk.

1 Introduction

In many economic environments, behavior is shaped by social interactions between agents. In individual decision problems, social interactions have been key to understanding outcomes as diverse as educational test scores, the demand for financial assets, and technology adoption (Sacerdote, 2001; Bursztyn *et al.*, 2014; Conley and Udry, 2010). In macroeconomics, the structure of firms' production and credit networks propagate shocks, or help firms to learn (Acemoglu *et al.*, 2012; Chaney, 2014). In political economy and public economics, ties between jurisdictions are key to understanding tax setting behavior (Tiebout, 1956; Shleifer, 1985; Besley and Case, 1994).

Underpinning all these bodies of research is some measurement of the underlying social ties between agents. However, information on social ties does not exist in most publicly available and widely used datasets. To overcome this limitation, studies of social interaction either *postulate* ties based on common observables or homophily, or *elicit* data on networks. However, it is increasingly recognized that postulated and elicited networks remain imperfect solutions to the fundamental problem of missing data on social ties, because of econometric concerns that arise with either method, or simply because of the cost of collecting network data.¹

Two consequences are that (i) the classes of problems in which social interactions occur are understudied, because social networks data is missing or too costly to collect; and (ii) there is no way to validate social interactions analysis in contexts where ties are postulated. In this paper, we tackle this challenge by deriving sufficient conditions under which global identification of the *entire structure* of social networks is obtained, using only observational panel data that itself contains *no* information on network ties. Our identification results allow the study of social interactions without data on social networks, and the validation of structures of social interaction where social ties have hitherto been postulated. The recovered networks are economically meaningful to explain the effects under study, since they are entirely estimated from the data itself, and not driven by *ex ante* assumptions on how individuals interact.

A researcher is assumed to have panel data on individuals $i = 1, \dots, N$ for instances $t = 1, \dots, T$. An instance refers to a specific observation for i and need not correspond to a time period (for example, if i refers to a firm, t could refer to market t). The outcome of interest for individual i in

¹As detailed in de Paula (2017), elicited networks are often self-reported and can introduce error to the outcome of interest. Network data can be censored if only a limited number of links can feasibly be reported. Incomplete survey coverage of nodes in a network may lead to biased aggregate network statistics. Chandrasekhar and Lewis (2016) show that even when nodes are randomly sampled from a network, partial sampling leads to non-classical measurement error and biased estimation. Collecting social network data is also a time- and resource-intensive process. In response to these concerns, a nascent strand of literature explores cost-effective alternatives to full elicitation to recover aggregate network statistics (Breza *et al.*, 2020).

instance t is y_{it} and is generated according to a canonical structural model of social interactions:²

$$y_{it} = \rho_0 \sum_{j=1}^N W_{0,ij} y_{jt} + \beta_0 x_{it} + \gamma_0 \sum_{j=1}^N W_{0,ij} x_{jt} + \alpha_i + \alpha_t + \epsilon_{it}. \quad (1)$$

Outcome y_{it} depends on the outcomes of other individuals to whom i is socially tied, y_{jt} , and x_{jt} includes characteristics of those individuals.³ $W_{0,ij}$ measures how the outcome and characteristics of j causally impact the outcome for i . The network is initially assumed to be fixed over time, and we later provide an extension of the method for time-varying networks. As outcomes for all individuals obey equations analogous to (1), the system of equations can be written in matrix notation, where the structure of interactions is captured by the adjacency matrix, denoted by W_0 . Our approach allows for unobserved heterogeneity across individuals α_i and common shocks to individuals α_t . This framework encompasses a classic linear-in-means specification as in Manski (1993). In his terminology, ρ_0 and γ_0 capture endogenous and exogenous social effects, and α_t captures correlated effects. The distinction between endogenous and exogenous peer effects is critical, as only the former generates social multiplier effects. In line with the literature, we maintain that the same W_0 governs the structure of both endogenous and exogenous effects. We later discuss relaxing this assumption when more than one regressor is used.

Manski’s seminal contribution set out the reflection problem of separately identifying endogenous, exogenous, and correlated effects in linear models. However, it has been somewhat overlooked that he also set out another challenge in the identification of the social network in the first place.⁴ This is the problem we tackle, and thus, we expand the scope of identification beyond ρ_0 , β_0 , and γ_0 . Our point of departure from much of the literature is therefore to presume W_0 is *entirely unknown* to the researcher. We derive sufficient conditions under which all the entries in W_0 , and the endogenous and exogenous social effect parameters, ρ_0 and γ_0 , are globally identified from “reduced form” parameters. By identifying the social interactions matrix W_0 , our results allow the recovery of aggregate network characteristics, such as the degree distribution and patterns of homophily, as well as node-level statistics such as the strength of social interactions between nodes, and the centrality of nodes. Such aggregate and node-level statistics often map back to underlying models of social interaction (Ballester *et al.*, 2006; Jackson *et al.*, 2017; de Paula, 2017).

²Blume *et al.* (2015) present micro-foundations for this estimating equation based on non-cooperative games of incomplete information for individual choice problems.

³In the case in which t is considered to be a time period, x_{it} may also include lagged values of y_{it} .

⁴Manski (1993) highlights difficulties (and potential restrictions) in identifying ρ_0 , β_0 and γ_0 when *all* individuals interact with each other, and when this is observed by the researcher. In (1), this corresponds to $W_{0,ij} = N^{-1}$, for $i, j = 1, \dots, N$. At the same time, he states (p. 536), “I have presumed that researchers know how individuals form reference groups and that individuals correctly perceive the mean outcomes experienced by their supposed reference groups. There is substantial reason to question these assumptions (...) If researchers do not know how individuals form reference groups and perceive reference-group outcomes, then it is reasonable to ask whether observed behavior can be used to infer these unknowns (...) The conclusion to be drawn is that informed specification of reference groups is a necessary prelude to analysis of social effects.”

Our identification strategy is new and fundamentally different from those employed elsewhere in the literature and does not rely on requirements about network sparsity. However, it delivers sufficient conditions that are mild and relate to existing results on the identification of social effects parameters when W_0 is known (Bramoullé *et al.*, 2009; De Giorgi *et al.*, 2010; Blume *et al.*, 2015). The intuition for our identification result is simple: model (1) has N^2 reduced-form parameters, and there are $N(N - 1) + 3$ structural unknowns (as no unit affects itself, so $W_{0,ii} = 0$). So there are more equations than unknowns if $N \geq 2$, and we demonstrate those can be solved for the parameters of interest under the assumptions we invoke. Our identification result is also useful in other estimation contexts, such as when a researcher has partial knowledge of W_0 ,⁵ or in navigating between priors on reduced-form and structural parameters in a Bayesian framework (see, e.g., Gefang *et al.*, 2023), thus avoiding issues the raised by Kline and Tamer (2016).

Global identification is a necessary requirement for consistency of extremum estimators such as those based on the GMM (Hansen 1982; Newey and McFadden 1994). Our identification analysis provides primitives for this condition. To estimate the model, we employ the Adaptive Elastic Net GMM method (Caner and Zhang, 2014), as this allows us to deal with a potentially high-dimensional parameter vector (in comparison to the time dimension in the data) including all the entries of the social interactions matrix W_0 , although other estimation protocols may also be entertained (e.g. using Bayesian methods or *a priori* information).⁶

We showcase the method using Monte Carlo simulations based on stylized random network structures as well as real-world networks. In each case, we take a fixed network structure W_0 and simulate panel data as if the data generating process were given by (1). We then apply the method to the simulated panel data to recover estimates of all elements in W_0 , as well as the endogenous and exogenous social effect parameters (ρ_0, γ_0) . The networks considered vary in size, complexity, and their aggregate and node-level features. In small samples, we find that the majority of links are identified even for $T = 5$, and the proportion of true non-links (zeros in W_0) captured correctly as zeros is over 85% even when $T = 5$. Of course, there are important limitations to the use of the method in small- T cases. Biases are expected and manifest themselves in two ways. First, weak links can be shrunk to zero, and the strength of strong edges can be overestimated. Second, the estimates of ρ and γ can suffer from small-sample bias, being analogous to well-known results for autoregressive time series models. Both properties rapidly improve with T . For instance, biases in the estimation of endogenous and exogenous effects parameters $(\hat{\rho}, \hat{\gamma})$ fall quickly with T and are close to zero for large sample sizes. The endogenous and exogenous social effects are also

⁵One such example is the nascent literature of Aggregate Relational Data (ARD) as in Breza *et al.* (2020). Another possibility is that individuals are known to belong to subgroups, so W_0 is block diagonal.

⁶The Elastic Net was introduced by Zou and Hastie (2005) in part to circumvent difficulties faced by alternative estimation protocols (e.g., LASSO) when the number of parameters, p , exceeds the number of observations, n (where p and n follow the notation in that paper). Whereas the theoretical results on the large-sample properties of elastic net estimators usually have not exploited sparsity, several articles have demonstrated their performance in data scenarios where this occurs. In Section 3, we provide an informal discussion on the performance in our context.

correctly captured as T increases. *A fortiori*, we estimate aggregate and node-level statistics of each network, demonstrating the accurate recovery of key players in networks, for example.

In the final part of our analysis, we apply the method to shed new light on a classic real-world social interactions problem: tax competition between US states. The literatures in political economy and public economics have long recognized the behavior of state governors might be influenced by decisions made in “neighboring” states. The typical empirical approach has been to postulate the relevant neighbors as being geographically contiguous states. Our approach allows us to infer the set of “economic” neighbors determining social interactions in tax setting behavior from panel data on outcomes and covariates alone. In this application, the panel data dimensions cover mainland US states, $N = 48$, for the years 1962-2015, $T = 53$.

The identified network structure of tax competition differs markedly from the assumption of competition between geographic neighbors. The identified economic network has fewer edges, and we identify non-adjacent states that influence tax setting behaviors. Differences in the structure of the identified economic and geography-based networks are reflected in the far lower clustering coefficient in the former (.042 versus .419). With the recovered social interactions matrix we establish, beyond geography, which covariates correlate to the existence of ties between states and so shed new light on hypotheses for social interactions in tax setting: factor mobility and yardstick competition (Tiebout, 1956; Shleifer, 1985; Besley and Case, 1994). The identified network highlights significant predictors of tax competition between states beyond distance: political homophily *reduces* the likelihood of a link, suggesting any yardstick competition driving social interactions occurs when voters compare their governor to those of the opposing party in other states. Tax haven states appear to be less influential in tax setting behaviors, easing concerns over a race-to-the-bottom in tax setting. Labor mobility between states does not robustly predict the existence of economic ties between states in tax setting behavior.

Given the relatively long study period in this application, at a final stage of analysis we extend our method to allow the strength of social interactions in tax competition (ρ_0, γ_0) and the structure of links in the economic network (W_0) to vary over time as we change the weight placed on observations from any given time period. We document the gradual increase in strength of social interactions over time, and the changing nature of the network of interactions. We utilize these findings to conduct counterfactual simulations of the general equilibrium propagation of tax shocks from a given state to all other mainland US states, and how these general equilibrium effects of the same policy shock vary as we place weight on observations later in our study period.

Our paper contributes to the literature on the identification of social interactions models. The first generation of papers studied the case where W_0 is known, so only the endogenous and exogenous social effects parameters needed to be identified. It is now established that if the known W_0 differs from the linear-in-means example where all units are linked with equal weights, ρ_0 and γ_0 can be identified (Bramoullé *et al.*, 2009; De Giorgi *et al.*, 2010). Intuitively, identification in those cases can use peers-of-peers, are not necessarily connected to individual i and can be used

to leverage variation from exclusion restrictions in (1), or can use groups of different sizes within which all individuals interact with each other (Lee, 2007). Bramoullé *et al.* (2009) show these conditions are met if I , W_0 , and W_0^2 are linearly independent, which is shown to hold generically by Blume *et al.* (2015). However, as made precise in Section 2, the linear algebraic arguments employed by Bramoullé *et al.* (2009) or Blume *et al.* (2015) do not apply when W_0 is unobserved, and other arguments have to be used instead.⁷

Blume *et al.* (2015) investigate the case when W_0 is *partially* observed and show that if two individuals are *known* not to be directly connected, the parameters of interest in a model related to (1) can be identified. Blume *et al.* (2011) take an alternative approach: suggesting a parameterization of W_0 according to a pre-specified distance between nodes. We do not impose such restrictions, but note that partial observability of W_0 or placing additional structure on W_0 is complementary to our approach, as it reduces the number of parameters in W_0 to be retrieved. Bonaldi *et al.* (2015) and Manresa (2016) estimate models like (1) when W_0 is not observed, but where ρ_0 is set to zero so there are no endogenous social effects. They use sparsity-inducing methods from the statistics literature, but the presence of ρ_0 in our case complicates identification because it introduces issues of simultaneity that we address.⁸

Rose (2015) also presents related identification results for linear models like (1), assuming the sparsity of the neighborhood structure. Intuitively, given two observationally equivalent systems, sparsity guarantees the existence of pairs that are not connected in either. Since observationally equivalent systems are linked via the reduced-form coefficient matrix, this pair allows one to identify certain parameters in the model. Having identified those parameters, Rose (2015) shows that one can proceed to identify other aspects of the structure (see also Gautier and Rose, 2016). This is related to the ideas in Blume *et al.* (2015), who show identification results can be leveraged if individuals are *known* not to be connected. Our main identification results do *not* rely on properties of sparse networks, and make use of plausible and intuitive conditions, whereas the auxiliary rank conditions necessary may be computationally complex to verify. More recently, Lewbel *et al.* (2022) propose an estimation strategy for the parameters ρ_0 , β_0 , and γ_0 of model (1) in the absence of network links if many different groups can be observed. Battaglini *et al.* (2022) estimate a structural model specifically for the case of unobserved social connections in the US Congress.

Finally, in the statistics literature, Lam and Souza (2020) study the penalized estimation of model (1) when W_0 is not observed, assuming the model and social interactions are identified. The statistical literature on graphical models has investigated the estimation of neighborhoods

⁷Alternative identification approaches when W_0 is known focus on higher moments (variances and covariances across individuals) of outcomes (de Paula, 2017) and rely on additional restrictions on higher moments of ϵ_{it} . Note that (1) is a spatial autoregressive model. In that literature, W_0 is also typically assumed to be known (Anselin, 2010).

⁸Manresa (2016) allows for unit-specific β_0 parameters. While in many applications those are taken to be homogeneous, we also discuss extensions on how heterogeneity in those parameters can be handled when $\rho_0 \neq 0$ in Appendix B.

defined by the covariance structure of the random variables at hand (Meinshausen and Buhlmann, 2006). This corresponds to a model where $y_t = (I - \rho_0 W_0)^{-1} \epsilon_t$ is jointly normal (abstracting from covariates). On a graph with N nodes corresponding to the variables in the model, an edge between two nodes (variables) i and j is absent when these two variables are conditionally independent given the other nodes. In the model above, the inverse covariance matrix is $(I - \rho_0 W_0)^\top \Sigma_\epsilon^{-1} (I - \rho_0 W_0)$, where Σ_ϵ is the variance covariance structure for ϵ_t . The discovery of zero entries in this matrix is not equivalent to the identification of W_0 and involves Σ_ϵ (as do identification strategies using higher moments when W_0 is known).⁹

We build on these papers by studying the problem where W_0 is potentially entirely unknown to the researcher. In so doing, we open up the study of social interactions to realms where social network data does not exist. In our case, we consider the definition of the network as the one that mediates, together with the variables x_{it} , the outcome process y_{it} according to Equation (1). The identified network may be a combination of elicited types of social interactions – such as friendship formation, lending and borrowing relations, links with relatives – or different from elicited data, as long as the links are relevant in determining the outcomes. In our case, and in line with the literature, the network ties W_{ij} are considered to be deterministic parameters or predetermined. Alternatively, the networks are assumed to be the outcome of a stochastic process, such as the latent space model (Hoff *et al.*, 2002; Breza *et al.*, 2020) or Exponential Random Graphs models (Holland and Leinhardt, 1981).

Our conclusions discuss how our approach can be modified, and assumptions weakened, to integrate partial knowledge of W_0 . We discuss further applications and the steps required to simultaneously identify models of network formation and the structure of social interactions. The practical use of our proposed method has already been demonstrated in applications. For example, Fetzer *et al.* (2021) study the impact on conflict of the transition of security responsibilities between international and Afghan forces. Our proposed method is used to control for violation of SUTVA-type hypotheses that might occur because of spillover and displacement effects of insurgent forces across districts. Since the pattern of displacement is unobserved – and, in fact, insurgents have incentives to obfuscate their strategy – the current method is applied to fully recover the network and bound the effects of the end of the military occupation on conflict.¹⁰

We proceed as follows. Section 2 presents our core result: the sufficient conditions under which the social interactions matrix, endogenous and exogenous social effects are globally identified. Section 3 describes the high-dimensional techniques used for estimation based on the Adaptive Elastic Net GMM method and presents simulation results from stylized and real-world networks. Section 4 applies our methods to study tax competition between US states. Section 5 concludes. The Appendix provides proofs and further details on estimation and simulations.

⁹Meinshausen and Buhlmann (2006)’s and Lam and Souza (2020)’s neighborhood estimates rely on (penalized) regressions of y_{it} on $y_{1t}, \dots, y_{i-1,t}, y_{i+1,t}, \dots, y_{N,t}$, which do not address the endogeneity in estimating W_0 .

¹⁰Zhou (2019) applies our identification results, focusing on unobserved networks with grouped heterogeneity, to suggest a nonlinear least squares procedure for estimation on a single network observation.

2 Identification

2.1 Setup

Consider a researcher with panel data covering $i = 1, \dots, N$ individuals repeatedly observed over $t = 1, \dots, T$ instances. The number of individuals N in the network is fixed but potentially large. The aim is to use this data to identify a social interactions model with no data on actual social ties. For expositional ease, we first consider identification in a simpler version of the canonical model in (1), where we drop individual-specific (α_i) and time-constant fixed effects (α_t) and assume x_{it} is a one-dimensional regressor for individual i and instance t . We later extend the analysis to include individual-specific, time-constant fixed effects and allow for multidimensional covariates $x_{k,it}$, $k = 1, \dots, K$. We adopt the subscript “0” to denote parameters generating the data, and non-subscripted parameters are generic values in the parameter space:

$$y_{it} = \rho_0 \sum_{j=1}^N W_{0,ij} y_{jt} + \beta_0 x_{it} + \gamma_0 \sum_{j=1}^N W_{0,ij} x_{jt} + \epsilon_{it}. \quad (2)$$

As the outcomes for all individuals $i = 1, \dots, N$ obey equations analogous to (2), the system of equations can be more compactly written in matrix notation as:

$$y_t = \rho_0 W_0 y_t + \beta_0 x_t + \gamma_0 W_0 x_t + \epsilon_t. \quad (3)$$

The vector of outcomes $y_t = (y_{1t}, \dots, y_{Nt})'$ assembles the individual outcomes in instance t ; the vector x_t does the same with individual characteristics. y_t , x_t , and ϵ_t have dimension $N \times 1$, the social interactions matrix W_0 is $N \times N$, and ρ_0 , β_0 , and γ_0 are scalar parameters. We do not make any distributional assumptions on ϵ_t beyond $\mathbb{E}(\epsilon_t | x_t) = 0$ (or $\mathbb{E}(\epsilon_t | z_t) = 0$ for an appropriate instrumental variable z_t if x_t is endogenous). We assume the network structure is predetermined and constant, and that the number of individuals N is fixed and repeated. In reality, networks may evolve over time. We thus later expand the method for dynamic network cases. The network structure W_0 is a parameter to be identified and estimated.

The social interaction model (3) has been widely studied (Manski, 1993; Manresa, 2016; and Blume *et al.*, 2015, among many others), but it is also restrictive in at least two senses. First, we consider W_{ij} to be fixed and predetermined, and not through models of strategic network formation (Jackson and Wolinsky, 1996; de Paula *et al.*, 2018) or of stochastic nature, as in the class of Exponential Random Graphs (Holland and Leinhardt, 1981) or Latent Distance models (Hoff *et al.*, 2002; Breza *et al.*, 2020). If there is feedback between outcome determination and link formation, and especially if this involves unobservables, it would be important to model network formation more explicitly.

A regression of outcomes on covariates corresponds, then, to the reduced form for (3),

$$y_t = \Pi_0 x_t + \nu_t, \quad (4)$$

with $\Pi_0 = (I - \rho_0 W_0)^{-1}(\beta_0 I + \gamma_0 W_0)$ and $\nu_t \equiv (I - \rho_0 W_0)^{-1} \epsilon_t$. If W_0 is observed, Bramoullé *et al.* (2009) note that a structure (ρ, β, γ) that is observationally equivalent to $(\rho_0, \beta_0, \gamma_0)$ is such that $(I - \rho_0 W_0)^{-1}(\beta_0 I + \gamma_0 W_0) = (I - \rho W_0)^{-1}(\beta I + \gamma W_0)$. This can be written as a linear equation in I, W_0 , and W_0^2 , and identification is established if those matrices are linearly independent. If W_0 is not observed, the putative unobserved structure comprises W_0 , and an observationally equivalent parameter vector will instead satisfy $(I - \rho_0 W_0)^{-1}(\beta_0 I + \gamma_0 W_0) = (I - \rho W)^{-1}(\beta I + \gamma W)$. Following the strategy in Bramoullé *et al.* (2009) would lead to an equation in I, W, W_0 , and WW_0 , so the insights obtained in that paper do *not* carry over to the case we study when W_0 is unknown.

We establish identification of the structural parameters of the model, including the social interactions matrix W_0 , from the coefficients matrix Π_0 . Without data on the network W_0 , we treat it as an additional parameter in an otherwise standard model relating outcomes and covariates. Our identification strategy relies on how changes in covariates x_{it} reverberate through the system and impact y_{it} , as well as outcomes for other individuals. These are summarized by the entries of the coefficient matrix Π_0 , which, in turn, encode information about W_0 and $(\rho_0, \beta_0, \gamma_0)$. A non-zero partial effect of x_{it} on y_{jt} indicates the existence of direct *or* indirect links between i and j . When $\rho_0 = 0$ (and $\Pi_0 = \beta_0 I + \gamma_0 W_0$), only direct links produce such a correlation. When $\rho \neq 0$, both direct and indirect connections may generate a non-zero response, but distant connections will lead to a lower response. Our results formally determine sufficient conditions to precisely disentangle these forces.

We set out six assumptions underpinning our main identification results. Three of these are entirely standard. A fourth is a normalization required to separately identify (ρ_0, γ_0) from W_0 , and the fifth is closely related to known results on the identification of (ρ_0, γ_0) when W_0 is known (Bramoullé *et al.*, 2009). The sixth assumption pertains to the relation between the nature of repeated multiple observations of the outcome and covariates and restrictions on the stability of W . These Assumptions (A1-A6) deliver an identified set of up to two points.

Our first assumption explicitly states that no individuals affect themselves and is a standard condition in social interaction models:

$$(A1) \quad (W_0)_{ii} = 0, \quad i = 1, \dots, N.$$

Assumption (A1) rules out applications with self-influence. For example, Input-Output matrices typically feature $(W_0)_{ii} > 0$, as firms tend to source from other firms in the same industry. With Assumption (A1), we can omit elements on the diagonal of W_0 from the parameter space. We thus can denote a generic parameter vector as $\theta = (W_{12}, \dots, W_{N,N-1}, \rho, \gamma, \beta)' \in \mathbb{R}^m$, where $m = N(N-1) + 3$, and W_{ij} is the (i, j) -th element of W . Reduced-form parameters can be tied

back to the structural model (3) by letting $\Pi : \mathbb{R}^m \rightarrow \mathbb{R}^{N^2}$ define the relation between structural and reduced-form parameters:

$$\Pi(\theta) = (I - \rho W)^{-1} (\beta I + \gamma W),$$

where $\theta \in \mathbb{R}^m$, and $\Pi_0 \equiv \Pi(\theta_0)$.

As ϵ_t (and, consequently, ν_t) is mean-independent from x_t , $\mathbb{E}[\epsilon_t|x_t] = 0$, the matrix Π_0 can be identified as the linear projection of y_t on x_t . We do not impose additional distributional assumptions on the disturbance term, except for conditions that allow us to identify the reduced-form parameters in (4). If x_t is endogenous, i.e., $\mathbb{E}[\epsilon_t|x_t] \neq 0$, a vector of instrumental variables z_t may still be used to identify Π_0 . In either case, identification of Π_0 requires variation of the regressor across individuals i and through instances t . In other words, either $\mathbb{E}[x_t x_t']$ (if exogeneity holds) or $\mathbb{E}[x_t z_t']$ (otherwise) is full-rank.

Our next assumption controls the propagation of shocks and guarantees that they die as they reverberate through the network. This provides adequate stability and is related to the concept of stationarity in network models. It implies the maximum eigenvalue norm of $\rho_0 W_0$ is less than one and ensures $(I - \rho_0 W_0)$ is a non-singular matrix. As the variance of y_t exists, the transformation $\Pi(\theta_0)$ is well-defined, and the Neumann expansion $(I - \rho_0 W_0)^{-1} = \sum_{j=0}^{\infty} (\rho_0 W_0)^j$ is appropriate.

(A2) $\sum_{j=1}^N |\rho_0(W_0)_{ij}| < 1$ for every $i = 1, \dots, N$, $\|W_0\| < C$ for some positive $C \in \mathbb{R}$ and $|\rho_0| < 1$.

We next assume that network effects do not cancel out, another standard assumption. As we will show, this assumption rules out the pathological case in which endogenous and exogenous effects exactly cancel each other out:

(A3) $\beta_0 \rho_0 + \gamma_0 \neq 0$.

The need for this assumption can be shown by expanding the expression for $\Pi(\theta_0)$, which is possible by (A2):

$$\Pi(\theta_0) = \beta_0 I + (\rho_0 \beta_0 + \gamma_0) \sum_{k=1}^{\infty} \rho_0^{k-1} W_0^k. \quad (5)$$

If Assumption (A3) were violated, $\beta_0 \rho_0 + \gamma_0 = 0$ and $\Pi_0 = \beta_0 I$, so the endogenous and exogenous effects would balance each other out, and network effects would be altogether eliminated in the reduced form.¹¹

Identification of the social effects parameters (ρ_0, γ_0) requires that at least one row of W_0 adds to a fixed and known number. Otherwise, ρ_0 and γ_0 cannot be separately identified from W_0 . Clearly, no such condition would be required if W_0 were observed.

¹¹One important case is when networks do not determine outcomes, which we interpret as $\rho_0 = \gamma_0 = 0$ or with W_0 representing the empty network. From equation (5), it is clear that if $\Pi(\theta_0)$ is *not* diagonal with constant entries, then it must be that $(\rho_0 \beta_0 + \gamma_0) \neq 0$, which implies that $\rho_0 \neq 0$ or $\gamma_0 \neq 0$, and also that W_0 is non-empty. Taken together, this suggests that the observation that $\Pi(\theta_0)$ is not diagonal is sufficient to ensure that network effects are present and Assumption (A3) is not violated.

(A4) There is an i such that $\sum_{j=1,\dots,N}(W_0)_{ij} = 1$.

Letting $W_y \equiv \rho_0 W_0$ and $W_x \equiv \gamma_0 W_0$ denote the matrices that summarize the influence of peers' outcomes (the endogenous social effects) and characteristics on one's outcome (the exogenous social effects), respectively, the assumption above can be seen as a normalization. In this case, ρ_0 and γ_0 represent the row-sum for individual i in W_y and W_x , respectively.¹²

The fifth assumption allows for a specific kind of network asymmetry. We require the diagonal of W_0^2 not to be constant as one of our sufficient conditions for identification.

(A5) There exists l, k such that $(W_0^2)_{ll} \neq (W_0^2)_{kk}$, i.e., the diagonal of W_0^2 is not proportional to ι , where ι is the $N \times 1$ vector of ones.

In unweighted networks, the diagonal of the square of the social interactions matrix captures the number of reciprocated links for each individual or, in the case of undirected networks, the popularity of those individuals. Assumption (A5) hence intuitively suggests differential popularity across individuals in the social network.

This assumption is related to the network asymmetry condition proposed elsewhere, such as in Bramoullé *et al.* (2009). They show that when W_0 is known, the structural model (2) is identified if I , W_0 , and W_0^2 are linearly independent. Given the remaining assumptions, this condition is satisfied if (A5) is satisfied, but the converse is *not* true: one can construct examples in which I , W_0 , and W_0^2 are linearly independent when W_0^2 has a constant diagonal, so Π_0 does not pin down θ_0 . See Example 1 in Appendix A. The strengthening of this hypothesis is the formal price to pay for the social interactions matrix W_0 being unknown to the researcher.

Before proceeding to our formal results, we provide a very simple illustration to shed light on how the assumptions above come together to provide identification. Suppose the observed reduced-form matrix is,

$$\Pi_0 = \frac{1}{455} \begin{bmatrix} 275 & 310 & 0 \\ 310 & 275 & 0 \\ 0 & 0 & 182 \end{bmatrix},$$

and that, following (A4), the first row is normalized to one. From the third row and column of Π_0 , we see there is no path of any length connecting the individual in row 3 to or from those in rows 1 or 2, since her outcome is not affected by their covariates and their outcomes are not affected by her covariates. In other words, individual 3, is isolated and $(W_0)_{13} = (W_0)_{23} = (W_0)_{31} = (W_0)_{32} = 0$. On the other hand, individuals 1 and 2 cannot be isolated, as their covariates are correlated with the other individual's outcome, reflecting (A5).¹³ Due to the row-sum normalization of the first

¹²Alternatively, one could normalize $\rho^* = 1$ and rescale the network accordingly. In this case, $W^* = \rho_0 W_0$ would be identified instead. Also, $W_x = \frac{\gamma_0}{\rho_0} W^*$ so γ_0 would be identified relative to ρ_0 . W_y and W_x would be unchanged.

¹³If on the other hand, $(W_0)_{ij} = 0.5, i \neq j$ in violation of (A5), and all agents were connected, the model would not be identified.

row, $(W_0)_{12} = 1$. Using (A3), it can be seen that W_0 is symmetric if Π_0 is symmetric. We thus find that $(W_0)_{21} = 1$. This and (A1) map all elements of W_0 , and thus,

$$W_0 = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

As the third individual is isolated, she will only be affected by her exogenous x_i and not by endogenous or exogenous peer effects. Hence, the (3,3) element of Π_0 is equal to $\beta_0 = \frac{182}{455} = .4$. To find ρ_0 , note that $(I - \rho_0 W_0)\Pi_0 = \beta_0 I + \gamma_0 W_0$. Hence, focusing on the (1,1) elements of the matrices above, we find that $\frac{275}{455} - \rho_0 \frac{310}{455} = .4$, implying $\rho_0 = .3$ (complying with (A2)). Finally, γ_0 is identified from entry (1,2), giving $\gamma_0 = \frac{310}{455} - .3 \frac{275}{455} = .5$.

Our final assumption articulates the need for a constant network W_0 observed over multiple instances of y_t and x_t :

(A6) y_t and x_t are observed for individuals $i = 1, \dots, N$, and instances $t = 1, \dots, T$, and the network W_0 does not depend on t

Here, “instances” can refer to time but also to settings in which the same units are observed over multiple episodes. For example, if i are firms, then t can be segmented markets in which they operate. For simplicity, we refer to an instance as a time period. If Π_0 is known, the main identification result we articulate below will state that W_0 , ρ_0 , β_0 and γ_0 are globally identified. However, in practice, Π_0 is rarely observed and thus all quantities need to be estimated. For this purpose, when Π_0 is not known, multiple observations of y_t and x_t with a constant W_0 are required to implement the estimator. We expand on estimation requirements in Section 3.1.

Importantly, the main identification results (for a given Π_0) could, in principle, be applied for each time period t . That is, one can write a version of Equation (3) as

$$y_t = \rho_0 W_{0t} y_t + \beta_0 x_t + \gamma_0 W_{0t} x_t + \epsilon_t,$$

where W_{0t} is time-varying and, consequently, the reduced-form interaction matrix $\Pi_{0t} = (I - \rho_0 W_{0t})^{-1}(I\beta_0 + W_{0t}\gamma_0)$ is also time-varying. If the reduced-form matrices were known, the identification results we develop below could be applied to the reduced-form element by element for each Π_{0t} . Again, one rarely observes Π_{0t} , for all $t \in [1, T]$. This observation will motivate an extension of the method, presented in Section 2.3.3, where Assumption (A6) is relaxed and W_0 is allowed to vary with t .

In an extension, we allow the network W_{0t} to vary over time and introduce kernel weights. Akin to the nonparametric regression $Y = f(X) + \epsilon$, f is identified if $\mathbb{E}(\epsilon|X) = 0$, and it is possible to estimate f using neighboring observations if f is sufficiently smooth or varies slowly. Similar considerations extend to varying-coefficient models and, in particular, time-varying coefficient

models where local stability conditions as those discussed in Dahlhaus (2012) are usually invoked (see also Hastie and Tibshirani (1993) and their Example (e)).

2.2 Main Identification Results

Under the assumptions above, we can begin to identify parameters related to the network. These results are then useful for our main identification theorems. Let λ_{0j} denote an eigenvalue of W_0 with corresponding eigenvector $v_{0,j}$ for $j = 1, \dots, N$. Assumptions (A2) and (A3) allow us to identify the eigenvectors of W_0 directly from the reduced form. As $|\rho_0| < 1$:

$$\begin{aligned} \Pi_0 v_{0,j} &= \beta_0 v_{0,j} + (\rho_0 \beta_0 + \gamma_0) \sum_{k=1}^{\infty} \rho_0^{k-1} W_0^k v_{0,j} \\ &= \left[\beta_0 + (\rho_0 \beta_0 + \gamma_0) \sum_{k=1}^{\infty} \rho_0^{k-1} \lambda_{0,j}^k \right] v_{0,j} \\ &= \frac{\beta_0 + \gamma_0 \lambda_{0,j}}{1 - \rho_0 \lambda_{0,j}} v_{0,j}. \end{aligned} \tag{6}$$

The infinite sum converges as $|\rho_0 \lambda_{0,j}| < 1$ by (A2). The equation above implies that $v_{0,j}$ is also an eigenvector of Π_0 with the associated eigenvalue $\lambda_{\Pi,j} = \frac{\beta_0 + \gamma_0 \lambda_{0,j}}{1 - \rho_0 \lambda_{0,j}}$. The fact that eigenvectors of W_0 are also eigenvectors of Π_0 has a useful implication: eigencentralities may be identified from the reduced form, even when W_0 is not identified. As detailed in de Paula (2017) and Jackson *et al.* (2017), such eigencentralities often play an important role in empirical work as they allow a mapping back to underlying models of social interaction.¹⁴

Now let $\Theta \equiv \{\theta \in \mathbb{R}^m : \text{Assumptions (A1)-(A6) are satisfied}\}$ be the structural parameter space of interest. Our identification argument is structured as follows: a) we first establish local identification of the mapping $\Pi(\theta)$ using classical results on the rank of the gradient of Rothenberg, 1971 (Theorem 1); b) we then show that $\Pi(\theta)$ is proper (Corollary 1); and c) has a connected image (Lemma 2, in the Appendix); d) allowing us to state the cardinality of the pre-image $\Pi^{-1}(\bar{\Pi})$ is constant for any $\bar{\Pi}$ in the image of $\Pi(\cdot)$, and that the cardinality is at most 2 (Theorem 2). We then provide additional conditions to narrow the identified set to a singleton (Corollaries 2-4).

We now formally present our results. Our first theorem establishes local identification of the mapping. A parameter point θ_0 is locally identifiable if there exists a neighborhood of θ_0 containing no other θ which is observationally equivalent. Using classical results in Rothenberg (1971), we show that our assumptions are sufficient to ensure that the Jacobian of Π relative to θ is non-singular, which, in turn, suffices to establish local identification.

Theorem 1. *Assume (A1)-(A6). $\theta_0 \in \Theta$ is locally identified.*

¹⁴To identify the eigencentralities, we identify the eigenvector that corresponds to the dominant eigenvalue. If W_0 is non-negative and irreducible, this is the (unique) eigenvector with strictly positive entries, by the Perron-Frobenius theorem for non-negative matrices (see Horn and Johnson, 2013, p. 534).

An immediate consequence of local identification is that the set $\{\theta \in \Theta : \Pi(\theta) = \Pi(\theta_0)\}$ is discrete (i.e., its elements are isolated points). The following corollary establishes that Π is a proper function, i.e. the inverse image $\Pi^{-1}(K)$ of any compact set $K \subset \mathbb{R}^{N^2}$ is also compact (Krantz and Parks, 2013, p. 124). Since it is discrete, the identified set must be finite.

Corollary 1. *Assume (A1)-(A6). Then $\Pi(\cdot)$ is a proper mapping. Moreover, the set $\{\theta : \Pi(\theta) = \Pi(\theta_0)\}$ has a finite number of elements.*

Under additional assumptions, the identified set is at most a singleton in each of the partitioning sets $\Theta_- \equiv \Theta \cap \{\rho\beta + \gamma < 0\}$ and $\Theta_+ \equiv \Theta \cap \{\rho\beta + \gamma > 0\}$.¹⁵

Since $\Theta = \Theta_- \cup \Theta_+$, if the sign of $\rho_0\beta_0 + \gamma_0$ is unknown, the identified set contains, at most, two elements. In the theorem that follows, we show global identification only for $\theta \in \Theta_+$, since arguments are mirrored for $\theta \in \Theta_-$.

Theorem 2. *Assume (A1)-(A6). Then for every $\theta \in \Theta_+$, we have $\Pi(\theta) = \Pi(\theta_0) \Rightarrow \theta = \theta_0$. That is, θ_0 is globally identified with respect to the set Θ_+ .*

Similar arguments apply if Theorem 2 instead were to be restricted to $\theta \in \Theta_-$. The proof of the corollary below is immediate and therefore omitted.

Corollary 2. *Assume (A1)-(A6). If $\rho_0\beta_0 + \gamma_0 > 0$, then the identified set contains at most one element, and similarly if $\rho_0\beta_0 + \gamma_0 < 0$. Hence, if the sign of $\rho_0\beta_0 + \gamma_0$ is unknown, the identified set contains, at most, two elements.*¹⁶

We now turn our attention to the problem of identifying the sign of $\rho_0\beta_0 + \gamma_0$ from the observation of Π_0 . This would then allow us to establish global identification using Theorem 2. It is apparent from (5) that if $\rho_0 > 0$ and $(W_0)_{ij} \geq 0$, for all $i, j = \{1, \dots, N\}$, the off-diagonal elements of Π_0 identify the sign of $\rho_0\beta_0 + \gamma_0$.

Corollary 3. *Assume (A1)-(A6). If $\rho_0 > 0$ and $(W_0)_{ij} \geq 0$, the model is globally identified.*

Real-world applications often suggest endogenous social interactions are positive ($\rho_0 > 0$), in which case global identification is fully established by Corollary 3. On the other hand, if $\rho_0 < 0$ (e.g., if outcomes are strategic substitutes), ρ_0^k in (5) alternates signs with k , and the off-diagonal elements no longer carry the sign of $\rho_0\beta_0 + \gamma_0$. Nonetheless, if W_0 is non-negative and irreducible (i.e., not permutable into a block-triangular matrix or, equivalently, a strongly connected social network), the model is also identifiable without further restrictions on ρ_0 :

¹⁵The global inversion results we use are related to, but different from, variations on a classic inversion result of Hadamard that has been used in the literature. In contrast, we employ results on the cardinality of the pre-image of a function, relying on less stringent assumptions. While the Hadamard result requires the image of the function to be simply-connected (Theorem 6.2.8 of Krantz and Parks, 2013), the results we rely on do not.

¹⁶Under some special conditions, the mirror image of θ_0 can be characterized from equation (5). If $-W_0$ satisfies Assumption (A4), we may set $\rho^* = -\rho_0$, $\beta^* = \beta_0$, $\gamma^* = -\gamma_0$ and $W^* = -W_0$. Then, $\rho_0\beta_0 + \gamma_0 = -(\rho^*\beta^* + \gamma^*)$. Also note that $\sum_{k=1}^{\infty} \rho_0^{k-1} W_0^k = -\sum_{k=1}^{\infty} (\rho^*)^{k-1} (W^*)^k$, so $(\rho_0\beta_0 + \gamma_0) \sum_{k=1}^{\infty} \rho_0^{k-1} W_0^k = (\rho^*\beta^* + \gamma^*) \sum_{k=1}^{\infty} (\rho^*)^{k-1} (W^*)^k$. It follows that $\Pi(\theta_0) = \Pi(\theta^*)$, where $\theta^* = (\rho^*, \beta^*, \gamma^*, W^*)$.

Corollary 4. *Assume (A1)-(A6), $(W_0)_{ij} \geq 0$ and W_0 is irreducible. If W_0 has at least two real eigenvalues or $|\rho_0| < \sqrt{2}/2$, then the model is globally identified.*

Corollary 4 requires that W_0 be irreducible, i.e., that it is not permutable into a block upper-triangular matrix. In the context of directed graphs, this is similar to requiring that the matrix be strongly connected, that is, that any node can be reached from any other node. The corollary then rules out cases when the network is not connected, for example, if there are two disjoint groups (with no connection across groups), or a star network pointing from the center towards the edges. The corollary holds if there are at least two real eigenvalues, or if ρ_0 is appropriately bounded. Since W_0 is non-negative, it has at least one real eigenvalue by the Perron-Frobenius theorem. If W_0 is symmetric, for example, its eigenvalues are all real, and Corollary 4 holds. It also holds if $(W_0)_{ij} \leq 0$, as we can rewrite the model as $\rho W_0 = -\rho|W_0|$, where $|W_0|$ is the matrix whose entries are the absolute values of the entries in W_0 . However, Corollary 4 rules out cases that mix positive $(W_0)_{ij} \geq 0$ and negative interactions $(W_0)_{ij} \leq 0$. In any case, the bound on $|\rho_0|$ is sufficient and holds in most (if not all) empirical estimates we are aware of obtained from either elicited or postulated networks, and in our application on tax competition.

2.3 Extensions

We present three extensions of the method for individual fixed effects, common shocks, and time-varying W . Appendix B describes extensions for multivariate covariates and heterogeneous β_0 .

2.3.1 Individual Fixed Effects

We observe outcomes for $i = 1, \dots, N$ individuals repeatedly through $t = 1, \dots, T$ instances. If t corresponds to time, it is natural to think of there being unobserved heterogeneity across individuals, α_i , to be accounted for when estimating Π_0 . The structural model (2) is then,

$$y_{it} = \rho_0 \sum_{j=1}^N W_{0,ij} y_{jt} + \beta_0 x_{it} + \gamma_0 \sum_{j=1}^N W_{0,ij} x_{jt} + \alpha_i + \epsilon_{it},$$

which can be written in matrix form as

$$y_t = \rho_0 W_0 y_t + x_t \beta_0 + W_0 x_t \gamma_0 + \alpha^* + \epsilon_t,$$

where α^* is the vector of fixed effects. Individual-specific and time-constant fixed effects can be eliminated using the standard subtraction of individual time averages. Defining $\bar{y}_t = T^{-1} \sum_{t=1}^T y_t$, $\bar{x}_t = T^{-1} \sum_{t=1}^T x_t$, and $\bar{\epsilon}_t = T^{-1} \sum_{t=1}^T \epsilon_t$,

$$y_t - \bar{y}_t = \rho_0 W_0 (y_t - \bar{y}_t) + (x_t - \bar{x}_t) \beta_0 + W_0 (x_t - \bar{x}_t) \gamma_0 + \epsilon_t - \bar{\epsilon}_t,$$

if W_0 does not change with time. Identification from the reduced form follows from previous theorems, since Π_0 is unchanged when regressing $y_t - \bar{y}_t$ on $x_t - \bar{x}_t$.¹⁷

2.3.2 Common Shocks

We next allow for unobserved common shocks to all individuals in the network in the same instance t . Such correlated effects α_t can confound the identification of social interactions. As we have not placed any distributional assumption on the covariance matrix of the disturbance term, our analysis readily incorporates correlated effects that are orthogonal to x_t . When this is not the case, one possibility is to model the correlated effects α_t explicitly. The model then is,

$$y_t = \rho_0 W_0 y_t + x_t \beta_0 + \gamma_0 W_0 x_t + \alpha_t \iota + \epsilon_t,$$

where α_t is a scalar capturing shocks in the network common to all individuals. Let $\Pi_{01} = (I - \rho_0 W_0)^{-1}$ and $\Pi_{02} = (\beta_0 I + \gamma_0 W_0)$ such that $\Pi_0 = \Pi_{01} \Pi_{02}$. The reduced-form model is

$$y_t = \Pi_0 x_t + \alpha_t \Pi_{01} \iota + v_t.$$

We propose a transformation to eliminate the correlated effects: exclude the individual-invariant α_t , subtracting the mean of the variables in a given period (global differencing). For this purpose, define $H = \frac{1}{n} \iota \iota'$. We note that in empirical and theoretical work, it is customary to strengthen Assumption (A4) and require that *all* rows of W_0 sum to one if no individual is isolated (see for example Blume *et al.*, 2015). This strengthened assumption is usually referred to as row-sum normalization, and is stated below:

(A4') For all $i = 1, \dots, N$, we have that $\sum_{j=1, \dots, N} (W_0)_{ij} = 1$.

This can be written compactly as $W_0 \iota = \iota$. In this case, W_0 can be interpreted as the normalized adjacency matrix. Under row-sum normalization we have that,

$$\begin{aligned} (I - H) y_t &= (I - H) (I - \rho_0 W_0)^{-1} (\beta_0 I + \gamma_0 W_0) x_t + (I - H) (I - \rho_0 W_0)^{-1} \epsilon_t \\ &= (I - H) \Pi_0 x_t + (I - H) v_t, \end{aligned}$$

because $(I - H) (I - \rho_0 W_0)^{-1} \alpha_t \iota = 0$ if Assumption (A4') holds. It then follows that $\tilde{\Pi}_0 = (I - H) \Pi_0$ is identified. The next proposition shows that, under row-sum normalization of W_0 , Π_0 is identified from $\tilde{\Pi}_0$ (and, as a consequence, the previous results immediately apply).

Proposition 1. *If W_0 is non-negative, irreducible, and row-sum normalized, Π_0 is identified from $\tilde{\Pi}_0$.*

¹⁷As is the case in panel data, this would require strict exogeneity ($\mathbb{E}[\epsilon_s | x_t] = 0$ for any s and t) or predetermined errors ($\mathbb{E}[\epsilon_s | x_t] = 0$ for $s \geq t$) so that the matrix Π_0 can be consistently estimated.

Under row-sum normalization of W_0 , a common group-level shock affects individuals homogeneously since $(I - \rho_0 W_0)^{-1} \alpha_t \iota = \alpha_t (I + \rho_0 W_0 + \rho_0^2 W_0^2 + \dots) \iota = \frac{\alpha_t}{1 - \rho_0} \iota$, which is a vector with no variation across entries. Consequently, global differencing eliminates correlated effects and $(I - H)(I - \rho_0 W_0)^{-1} \alpha_t \iota = (I - \rho_0 W_0)^{-1} \alpha_t (I - H) \iota = 0$. Absent row-sum normalization, global differencing does not ensure correlated effects are eliminated. To see this, note that $(I - \rho_0 W_0)^{-1}$ is no longer row-sum normalized and $\alpha_t (I - \rho_0 W_0)^{-1} \iota$ does not have constant entries.

The next proposition makes this point formally: that the stronger Assumption (A4') is *necessary* to eliminate group-level shocks by showing it is not possible to construct a data transformation that eliminates group effects in the absence of row-sum normalization.

Proposition 2. *Define $r_{W_0} = (I - \rho_0 W_0)^{-1} \iota$. If in space $\Theta = \{\theta \in \mathbb{R}^m : \text{Assumptions (A1)-(A6) are satisfied}\}$, there are N matrices $W_0^{(1)}, \dots, W_0^{(N)}$ such that $[r_{W_0^{(1)}} \dots r_{W_0^{(N)}}]$ has rank N , then the only transformation such that $(I - \tilde{H})(I - \rho_0 W_0)^{-1} \iota = 0$ is $\tilde{H} = I$.*

It is useful to be able to test for row-sum normalization (A4') as it enables common shocks to be accounted for in the social interactions model. This is possible as

$$\begin{aligned} \Pi_0 \iota &= \beta_0 \iota + (\rho_0 \beta_0 + \gamma_0) \sum_{k=1}^{\infty} \rho_0^{k-1} W_0^k \iota \\ &= \left[\beta_0 + (\rho_0 \beta_0 + \gamma_0) \sum_{k=1}^{\infty} \rho_0^{k-1} \right] \iota \\ &= \frac{\beta_0 + \gamma_0}{1 - \rho_0} \iota. \end{aligned} \tag{7}$$

The last equality follows from the observation that, under row-normalization of W_0 , $W_0^k \iota = W_0 \iota = \iota$, $k > 0$. This implies Π_0 has constant row-sums, which suggests row-sum normalization is testable. In the Appendix, we derive a Wald test statistic to do so.¹⁸

2.3.3 Time-varying W

We now relax Assumption (A6), which states that W_0 does not vary across the time periods $t = 1, \dots, T$. The version of Equation (3) with time-varying network is

$$y_t = \rho_0 W_{0t} y_t + \beta_0 x_t + \gamma_0 W_{0t} x_t + \epsilon_t,$$

with the reduced-form matrix $\Pi_{0t} = (I - \rho_0 W_{0t})^{-1} (I \beta_0 + W_{0t} \gamma_0)$. We note that the identification results developed in Subsection 2.2 can, in principle, be applied element by element to each Π_{0t} , leading to the identification of a time-varying W_{0t} (and, potentially, of the parameters ρ_0 , β_0 and γ_0).

¹⁸For ease of explanation, in the Appendix, we derive the test under the asymptotic distribution of the OLS estimator. The test generally holds with minor adjustments for estimators with known asymptotic distributions.

In practice, implementing any estimation strategy with a time-varying Π_{0t} (or W_{0t}) is not feasible using only observation from the single time period t . We instead adopt a kernel-weighted version. Define period-specific weights ω_t , and consider the transformed data $\tilde{y}_s = \omega_s(t)y_s$ and $\tilde{x}_s = \omega_s(t)x_s$, $s = 1, \dots, T$. Evidently, uniform weights $\omega_s = 1$, $s = 1, \dots, T$ are equivalent to the strategy not considering time-varying networks, and assuming that the networks are fixed within those windows. Alternatively, one could estimate W_t in time windows by setting $\omega_t = 1[\underline{t} \leq t \leq \bar{t}]$, where \underline{t} and \bar{t} are the start and end of the time window for which W_t is estimated. In this case, the minimum effective window length $\bar{t} - \underline{t}$ can be computed as we discuss in Section 3.1. In the context of DSGE models with time-varying parameters, Kapetanios *et al.* (2019) suggests a Gaussian kernel with positive weights throughout the entire sample. As in nonparametric regression with smooth kernel weights, it also assumes that the network evolves slowly over time. We further discuss this strategy in the estimation section, and it is implemented in the empirical application section below.

3 Implementation

We now transition from our core identification results to their practical implementation. In practice, Ordinary Least Squares (OLS) can only be used to estimate θ if $T \gg N$, which is in practice unlikely to be met, as this is a high-dimensional problem. Our preferred approach makes use of penalized estimation techniques that can be used for any given T . More specifically, we make use of the Adaptive Elastic Net GMM (Caner and Zhang, 2014), which is based on the penalized GMM objective function. Given the identification results presented in Section 2, the population moments used in forming the GMM objective function will be satisfied at the true parameter vector.

After setting out the estimation procedure, we showcase the method using Monte Carlo simulations based on stylized and real-world network structures. In each case, we take a fixed network structure W_0 , and simulate panel data as if the data generating process were given by the model in (1). We apply the method to the simulated panel data to recover estimates of all elements in W_0 , as well as the endogenous and exogenous social effect parameters.

3.1 Estimation

The parameter vector to be estimated is high-dimensional: $\theta = (W_{12}, \dots, W_{N,N-1}, \rho, \gamma, \beta)' \in \mathbb{R}^m$, where $m = N(N-1) + 3$ and W_{ij} is the (i, j) -th element of the $N \times N$ social interactions matrix W_0 . To be clear, in a network with N individuals, there are $N(N-1)$ potential interactions because an individual could interact with everyone else but herself (which would violate Assumption A1). As a consequence, even with a modest N , there are many more parameters to estimate, and m is large. For example, a network with $N = 50$ implies more than 2,000 parameters to estimate. While we consider N (and thus m) fixed, we still refer to θ as high-dimensional. OLS estimation requires $m \ll NT (\Rightarrow N \ll T)$, so many more time periods than individuals: a requirement often

met in finance data sets (van Vliet, 2018) or in other fields (see, e.g., Section 4.2 in Rothenhäusler *et al.*, 2015). Instead, to estimate a large number of parameters with limited data, we utilize high-dimensional estimation methods, which are the focus of a rapidly growing literature.

Sparsity is a key assumption underlying many high-dimensional estimation techniques. In the context of social interactions, we say that W_0 is sparse if \tilde{m} , the number of non-zero elements of W_0 , is such that $\tilde{m} \ll NT$. The notion of sparsity thus depends on the number of time periods. Sparsity corresponds to assuming that individuals influence or are influenced by a small number of others, relative to the overall size of the potential network and the time horizon in the data. As such, sparsity is typically *not* a binding constraint in social networks analysis.¹⁹

In the estimation of sparse models, the “effective number of parameters” (or “effective degrees of freedom”) relates to the number of variables with non-zero estimated coefficients (Tibshirani and Taylor, 2012). In the context of the current social network model, this is equivalent to m parameters, where $m = dN(N - 1) + K$ and d is the network density defined as $\tilde{m}/(N(N - 1))$. The Adaptive Elastic Net GMM estimator presented by Caner and Zhang (2014) converges at a rate of $\sqrt{NT/\tilde{m}} = \sqrt{NT/[dN(N - 1) + K]} = O(\sqrt{T/(dN)})$ (see remark 7 in Caner and Zhang, 2014). Hence the quality of the large sample results relies on a comparison between T and dN . In line with this, we thus require $NT \gg dN(N - 1) + K$. For example, in the high-school network of Coleman (1964) that is part of our simulation exercise, $N = 70$ and $d = 0.076$. Assuming $K = 3$, $N(N - 1) \times d + 3 = 370.1$.²⁰

Finally, to reiterate, our identification results themselves do *not* depend on the sparsity of networks. In particular, Assumptions (A1)-(A6) *do not* impose restrictions on the number of links in W_0 , or \tilde{m} .²¹ The identification results presented in Section 2 apply more broadly and irrespective of the estimation procedure.

Our preferred approach estimates the interaction matrix in the reduced form while penalizing and imposing sparsity on the structural object W_0 . We impose sparsity and penalization in the structural-form matrix W_0 because this is a weaker requirement than imposing sparsity and penalization in the reduced-form matrix Π_0 .²² To do so, we use the Adaptive Elastic Net GMM (Caner

¹⁹Common stylized networks are sparse, such as the star, lattice (each individual is a source of spillover only to one other individual), or interactions in pairs, triads or small groups (De Giorgi *et al.*, 2010). Real-world economic networks are also sparse. The sparsity in *AddHealth* friendship network is around 98%. Sparsity of the production networks in the US is above 99% (Atalay *et al.*, 2011).

²⁰As pointed out by a referee, variation in x will also matter for estimation precision. This is reflected in the asymptotic distribution for this estimator, shown later in this subsection.

²¹If $N \rightarrow \infty$, Assumption (A2) would imply vanishing $(W_0)_{ij}$ entries. As highlighted previously, we consider N to be fixed, in line with many practical applications. Furthermore, Assumption (A2) is used to represent inverse matrices as Neumann series in our identification results. What is necessary for this to hold is that a sub-multiplicative norm on ρW be less than one. Here we use a specific norm (i.e., the maximum row-sum norm), but other (induced) norms are also possible (i.e., the 2-norm or the 1-norm) (see Horn and Johnson, 2013, Chapter 5.6).

²²Note that even if W is sparse, Π may not be sparse. In Appendix C.1, we show that $[\Pi_0]_{ij} = 0$ if, and only if, there are no paths between i and j in W_0 , so the pair is not connected. So, sparsity in Π_0 is understood as W_0 being “sparsely connected”, which is a stronger assumption than sparsity in W_0 .

and Zhang, 2014), which is based on the penalized GMM objective function,

$$G_{NT}(\theta, p) \equiv g_{NT}(\theta)' M_T g_{NT}(\theta) + p_1 \sum_{\substack{i,j=1 \\ i \neq j}}^N |W_{i,j}| + p_2 \sum_{\substack{i,j=1 \\ i \neq j}}^N |W_{i,j}|^2 \quad (8)$$

where $\theta = (W_{1,2}, \dots, W_{N,N-1}, \rho, \gamma, \beta)'$ with dimension $m = N(N-1) + 3$, and p_1 and p_2 are the penalization terms. The term $g_{NT}(\theta)' M_T g_{NT}(\theta)$ is the unpenalized GMM objective function with moment conditions based on orthogonality between the structural disturbance term and the covariates: $g_{NT}(\theta) = \sum_{t=1}^T [x_{1t}e_t(\theta)' \cdots x_{Nt}e_t(\theta)']'$, $e_t(\theta) = y_t - \rho W y_t - \beta x_t - \gamma W x_t$. There are $q \equiv N^2$ moment conditions since x_{it} is orthogonal to e_{jt} for each $i, j = 1, \dots, N$. Hence, the GMM weight matrix M_T is of dimension $N^2 \times N^2$, symmetric, and positive definite. For simplicity, we use $M_T = I_{N^2 \times N^2}$. Note that if x_t is econometrically endogenous, one can also exploit moment conditions with respect to available instrumental variables.²³ Given the identification results presented in Section 2, if $\theta \neq \theta_0$ and does not belong to the identified set, then $\Pi(\theta) \neq \Pi(\theta_0)$. Consequently, the populational version of the GMM objective function is uniquely minimized at the true parameter vector θ_0 .

The penalization terms in Equation (8) are what makes this different from a standard GMM problem. The first term, $p_1 \sum_{i,j=1, i \neq j}^N |W_{i,j}|$, penalizes the sum of the absolute values of $W_{i,j}$, i.e., the sum of the strength of links, for all node-pairs. Depending on the choice of p_1 , some $W_{i,j}$'s will be estimated as exact zeros. A larger share of parameters will be estimated as zeros if p_1 increases. The second term, $p_2 \sum_{i,j=1, i \neq j}^N |W_{i,j}|^2$, penalizes the sum of the square of the parameters. This term has been shown to provide better model-selection properties, especially when explanatory variables are correlated (Zou and Zhang, 2009). The first-stage estimate is

$$\tilde{\theta}(p) = (1 + p_2/T) \cdot \arg \min_{\theta \in \mathbb{R}^m} G_{NT}(\theta, p) \quad (9)$$

where $(1 + p_2/T)$ is a bias-correction term also used by Caner and Zhang (2014).

Implementing the numerical optimization embedded in Equation (9) is computationally challenging, as $m = N(N-1) + 3$ may entail a large number of function arguments. We instead implement the following modification to use fast Least-Angle Regression (LARS) algorithms (Efron *et al.*, 2004). For any given ρ , β , and γ , the expression for $e_t(\theta)$ is linear in W :

$$e_t(\theta) = y_t - x_t \beta - W(\rho y_t + x_t \gamma) = \tilde{y}_{it}(\beta) - W \tilde{x}_t(\rho, \gamma)$$

where $\tilde{y}_{it}(\beta) \equiv y_t - x_t \beta$ and $\tilde{x}_t(\rho, \gamma) \equiv \rho y_t + x_t \gamma$ and, following the strategy above, is instrumented

²³For expositional ease, we describe estimation in the context of the reduced-form model (4), thereby abstaining from individual fixed or correlated effects. As the GMM estimator uses moments between the structural disturbance terms and covariates, this endogeneity is built into the estimation procedure.

with x_t . This motivates a two-step optimization routine:

$$\min_{\theta \in \Theta = \Theta_1 \times \Theta_2} G_{NT}(\theta, p) = \min_{(\rho, \beta, \gamma) \in \Theta_1} \left[\min_{W_{ij} \in \Theta_2} G_{NT}(\theta, p) \right],$$

where the expression in brackets has a computationally efficient solution through the LARS algorithm. The numerical optimization is then subsequently conducted over the parameter space of (ρ, β, γ) only. We also impose row-sum normalization. Details of the implementation are expanded in Appendix Subsection C.2.

A second (adaptive) step provides improvements by re-weighting the penalization by the inverse of the first-step estimates (Zou, 2006):

$$\tilde{\theta}^*(p) = (1 + p_2/T) \cdot \arg \min_{\theta \in \Theta} \left\{ g_{NT}(\theta)' M_T g_{NT}(\theta) + p_1^* \sum_{\substack{\{i,j:\tilde{W}_{ij} \neq 0, \\ i,j=1,\dots,N, \\ i \neq j\}}} \frac{|W_{i,j}|}{|\tilde{W}_{i,j}|^c} + p_2 \sum_{\substack{\{i,j:\tilde{W}_{ij} \neq 0, \\ i,j=1,\dots,N, \\ i \neq j\}}} |W_{i,j}|^2 \right\}, \quad (10)$$

where $\tilde{W}_{i,j}$ is the (i, j) -th element of the first-step estimate of W . We follow Caner and Zhang (2014) and set $c = 2.5$. If $\tilde{W}_{i,j} < 0.05$, we set $\tilde{W}_{i,j} = 0.05$. This ensures that the second-stage estimates can be non-zero even if the first-stage estimates were zero or small. The computational improvement – described above for the first-stage estimator – is also applied in the adaptive stage.

As a third and final step, we fix the support of $\tilde{\theta}^*(p)$, $\mathcal{S} = \{\rho, \beta, \gamma\} \cup \{W_{ij} : \tilde{W}_{ij}^* \neq 0\}$ and estimate the final parameters without penalization. This takes as arguments only the elements of $\tilde{\theta}^*(p)$ that were estimated as non-zero in the adaptive step. In essence, this step boils down to a standard GMM approach,

$$\hat{\theta}_{\mathcal{S}}(p) = \arg \min_{\theta \in \mathcal{S}} \{g_{NT}(\theta)' M_T g_{NT}(\theta)\}. \quad (11)$$

Importantly, Caner and Zhang (2014) show that the third-step estimator is asymptotically normal, with a known and easy-to-compute distribution,

$$\delta' \left[(\hat{G}' M_T \hat{G})^{-1} \cdot (\hat{G}' M_T \Omega M_T \hat{G}) \cdot (\hat{G}' M_T \hat{G})^{-1} \right]^{1/2} \cdot \sqrt{NT/\tilde{m}} \cdot (\hat{\theta}_{\mathcal{S}} - \theta_0) \xrightarrow{d} N(0, 1),$$

where $\hat{G} \equiv \hat{G}(\hat{\theta}) = \nabla g_{NT}(\theta)$ and $\Omega \equiv E[g_{NT}(\theta)g_{NT}(\theta)']$.²⁴ This allows us to conduct hypothesis testing and inference on the ρ, β, γ and the non-zero elements of W .

We write $p = (p_1, p_1^*, p_2)$ as the final set of penalization parameters. Conditional on p , the estimate of the procedure is $\hat{\theta}(p)$. As in Caner and Zhang (2014, p. 35), the penalization parameters

²⁴This applies in the case of small p_2 . In the case of large p_2 , the asymptotic distribution is pre-multiplied by $K_n = \frac{I + p_2 [\hat{G}(\hat{\theta})' \hat{\Omega}^{-1} \hat{G}(\hat{\theta})]^{-1}}{1 + p_2/NT}$. See Theorem 4 of Caner and Zhang (2014).

p are chosen by the BIC criterion. This balances model fit with the number of parameters included in the model.²⁵

In Appendix C.2, we provide further implementation details, including the choice of initial conditions. Of course, other estimation methods are available, and our identification results do not hinge on any particular estimator. Our aim is to demonstrate the practical feasibility of using the Adaptive Elastic Net estimator rather than claim it is the optimal estimator.²⁶

3.2 Simulations

We showcase the method using Monte Carlo simulations. We describe the simulation procedures, results, and robustness checks in more detail in Appendix D.1. Here, we just provide a brief overview to highlight how well the method works to recover social networks even in relatively short panels.

For each simulated network, we take a fixed network structure W_0 and simulate panel data as if the data generating process were given by (1). We then apply the method to the simulated panel data to recover estimates of all elements in W_0 , as well as the endogenous and exogenous social effect parameters (ρ_0, γ_0) . Our result identifies entries in W_0 and so naturally recovers links of varying strength. It is long recognized that link strength might play an important role in social interactions (Granovetter, 1973). Data limitations often force researchers to postulate some ties to be weaker than others (say, based on interaction frequency). In contrast, our approach identifies the continuous strength of ties, $W_{0,ij}$, where $W_{0,ij} > 0$ implies node j influences node i .

The stylized networks we consider are a random network and a political party network in which two groups of nodes each cluster around a central node. The real-world networks we consider are the high-school friendship network in Coleman (1964) from a small high school in Illinois, and one of the village networks elicited in Banerjee *et al.* (2013) from rural Karnataka, India.

Summary statistics for each network are presented in Panel A of Table A1. The four networks differ in their size, complexity, and the relative importance of strong and weak ties. For example, the Erdős-Renyi network only has strong ties, while the political party network has twice as many strong as weak ties. For the real-world networks, the mean out-degree distributions are higher, so the majority of ties are weak, with the high school network having around 80% of its edges being weak ties. All four networks are also sparse.

²⁵Following Caner and Zhang, 2014, the choice of p , which we denote as \hat{p} , is the one that minimizes

$$\text{BIC}(p) = \log \left[g_{NT}(\hat{\theta}(p))' M_T g_{NT}(\hat{\theta}(p)) \right] + A(\hat{\theta}(p)) \cdot \frac{\log T}{T}$$

where $A(\hat{\theta}(p))$ counts the number of non-zero coefficients among $\{W_{1,2}, \dots, W_{N,N-1}\}$, and larger than a numerical tolerance, which we set at 10^{-5} . See also Zou *et al.* (2007).

²⁶See the alternative approaches of Gautier and Tsybakov (2014), Manresa (2016), Lam and Souza (2016), and Gautier and Rose (2016).

For the stylized networks, we assess the performance of the estimator for a fixed network size, $N = 30$. We simulate the real-world networks using non-isolated nodes in each (so $N = 70$ and 65 respectively).²⁷

We evaluate the procedure over varying panel lengths (starting from short panels with $T = 5$), using various metrics. Given our core contribution is to identify the social interactions matrix, we first examine the proportion of true zero entries in W_0 estimated as zeros and the proportion of true non-zero entries estimated as non-zeros. A global perspective of the proximity between the true and estimated networks can be inferred from their average absolute distance between elements. This is the mean absolute deviation of \hat{W} and $\hat{\Pi}$ relative to their true values, defined as $MAD(\hat{W}) = \frac{1}{N(N-1)} \sum_{i,j,i \neq j} |\hat{W}_{ij} - W_{ij,0}|$ and $MAD(\hat{\Pi}) = \frac{1}{N(N-1)} \sum_{i,j,i \neq j} |\hat{\Pi}_{ij} - \Pi_{ij,0}|$. As these metrics are closer to zero, more of the elements in the true matrix are correctly estimated. Finally, we evaluate the procedure’s performance using averaged estimates of the endogenous and exogenous social effect parameters, $\hat{\rho}$ and $\hat{\gamma}$. In keeping with the estimation strategy in our empirical application, we report unpenalised GMM.

3.3 Results

Figure A1 shows the simulation results as evaluated using the six metrics described above. Panel A shows that for each network, the proportion of zero entries in W_0 correctly estimated as zeros is above 95% even when $T = 5$. The proportion approaches 100% as T grows. Conversely, Panel B shows the proportion of strong non-zero entries estimated as non-zeros (defined as larger than 0.3) is also high for a small T . It is above 70% from $T = 5$ for the Erdős-Renyi network, being at least 85% across networks for $T = 25$, and increasing as T grows. As discussed above, the Adaptive Elastic Net estimator may only recover strong edges well, and not necessarily the weaker ones, due to the well-known issue with shrinkage estimators that they tend to shrink small parameters to zero. We return to this issue below.

Panels C and D show that for each simulated network, the mean absolute deviation between estimated and true networks for \hat{W} and $\hat{\Pi}$ falls quickly with T and is close to zero for large sample sizes. Finally, Panels E and F show that biases in the endogenous and exogenous social effects parameters, $\hat{\rho}$ and $\hat{\gamma}$, also fall in T (we do not report the bias in $\hat{\beta}$ since it is close to zero for all T). The fact that biases are not zero is as expected for a small T , being analogous to well-known results for autoregressive time series models.²⁸

Figure A2 shows that, as T increases, the procedure detects weaker links. The figure also shows that, with low sample sizes, weak edges are generally not detected. This pattern is consistent with the well-known fact that small parameters are likely shrunk to zero due to the penalization (Belloni and Chernozhukov, 2011). The absence of weak edges also implies that the strength of strong edges

²⁷Like Bramoullé *et al.* (2009), we exclude isolated nodes because they do not conform to row-sum normalization.

²⁸The bias in spatial auto-regressive models with a small number of observations *even when the network is observed* is similarly documented by Smith (2009), Neuman and Mizruchi (2010), Wang *et al.* (2014), and others.

may also be over-estimated, since rows are normalized to one. In Panel A, we show the distribution of the estimates of \hat{W}_{ij} , with $T = 25$ and for the high-school network. We show the distribution for the five most common values of $W_{0,ij}$. We find that most edges weaker than .5 are not detected; edges with a strength of .75 are substantially more likely to be estimated as non-zeros. When they are detected as non-zeros, they are more likely to be over-estimated. When we estimate W with $T = 150$, Panel B shows that virtually all edges with strength greater than .5 are estimated as non-zeros, and most edges with strength .375 are also detected. We further see a more continuous distribution of estimates of edge strength. Only edges smaller than .25 are not detected. Panels C and D show a similar conclusion for the village network.

Figure A3 shows the simulated and actual networks under $T = 100$ time periods. The network size is set to $N = 30$ in the two stylized networks, $N = 70$ for the high school network, and $N = 65$ for the village household network. In comparing the simulated and true networks, Figure A3 distinguishes between kept edges, added edges, and removed edges. Kept edges are depicted in blue: these links are estimated as non-zero in at least 5% of the iterations, and are also non-zero in the true network. Added edges are depicted in green: these links are estimated as non-zero in at least 5% of the iterations but the edge is zero in the true network. Removed edges are depicted in red: these links are estimated as zero in at least 5% of the iterations but are non-zero in the true network. Figure A3 further distinguishes between strong and weak links: strong links are shown as solid edges ($W_{0,ij} > .3$), and weak links are shown as dashed edges.

Panel A of Figure A3 compares the simulated and true Erdős-Renyi networks. All links are recovered. For the political party network, Panel B shows that all strong edges are correctly estimated. However, around half the weak edges are recovered (blue dashed edges), with the others being missed (red dashed edges). As discussed above, this is not surprising given that shrinkage estimators force small non-zero parameters to zero. Hence, a larger T is needed to achieve similar performance to the other simulated networks in terms of detecting weak links. For the more complex and larger real-world networks, Panel C shows that in the high-school network, the strong edges are all recovered. However, around half the weak edges are missing (red dashed edges), and there are a relatively small number of added edges (green edges): these amount to 87 edges, or approximately 1.9% of the 4,534 zero entries in the true high-school network. A similar pattern of results is seen in the village network in Panel D: the strong edges are all recovered, and here the majority of weak edges are also recovered.

Panel B of Table A1 compares the network- and node-level statistics calculated from the recovered social interactions matrix \hat{W} to those in Panel A from the true interactions matrix W_0 . The random Erdős-Renyi network is perfectly recovered. For the political party network, the number of recovered edges is slightly lower than in the true network (41 vs. 45), and all edges are classified as strong. The mean of the in- and out-degree distributions are slightly lower in the recovered network, and all three nodes with the highest out-degree are correctly captured (nodes 1, 11, and 28), include both party leaders (individuals 1 and 11). We then move to discussing the performance in

the two real-world networks. In the high-school network, 30% of all edges are correctly recovered, and they are all strong edges. As already noted in Figure A2, weak edges are not well estimated in the high-school network. This draws two main consequences. First, the average in- and out-degrees are smaller in the recovered network relative to the true network. Second, we over-estimate the number of strong edges (61 vs. 113). This is a downside of row-sum normalization: because some weak edges get estimated as zeros, the non-zeros are over-estimated so that the row adds to one. We do, however, recover all three individuals with the highest out-degree. Finally, in the village network, half the edges are recovered. The same phenomena of underestimating weak and overestimating strong edges are again observed. We again recover the three households with the highest out-degree (nodes 16, 35, and 57).

In the Appendix, we show the robustness of the simulation results to (i) varying network sizes and (ii) alternative parameter choices and richening the structure of shocks across nodes. We also demonstrate the gains from using the Adaptive Elastic Net GMM estimator over alternative estimators, such as the Adaptive Lasso and OLS.

4 Application: Tax Competition between US States

We apply our results to shed new light on a classic social interactions problem: tax competition between US states (Wilson, 1999). Defining competing “neighbors” remains the central empirical challenge in this literature. Theory provides some guidance on the issue through two mechanisms driving interactions across jurisdictions: factor mobility and yardstick competition.

On factor mobility, Tiebout (1956) first argued that labor and capital can move in response to differential tax rates across jurisdictions. Factor mobility leads naturally to the postulated social interactions matrix being (i) geographic neighbors, given labor mobility, and (ii) jurisdictions with similar economic or demographic characteristics, given capital mobility (Case *et al.*, 1989).

Yardstick competition is driven by voters making comparisons between states to learn about their own politician’s quality (Shleifer, 1985). Besley and Case (1995a) formalize the idea in a model where voters use taxes set by governors in other states to infer their own governor’s quality. Yardstick competition leads naturally to the postulated interactions matrix being “political neighbors”: states that voters make comparisons to.

In this application, the number of nodes and time periods is relatively low: the data covers mainland US states, $N = 48$, for the years 1962-2015, $T = 53$. Our approach identifies the structure of social interactions among “economic neighbors”, denoted W_{econ} . We contrast this against a null that state taxes are influenced by geographic neighbors, W_{geo} , as shown in Panel A of Figure 1A. With W_{econ} recovered, we can establish, beyond geography, what predicts the strength of ties between states and provide fresh insights on drivers of tax competition.

Before using the real data, we confirm the estimator’s performance when the true network is W_{geo} in simulated settings. In line with the findings of the previous section, Panel B of Figure 1A

shows that (i) the procedure recovers strong edges frequently (more specifically, 89% of the true strong edges are recovered) and (ii) performance deteriorates when recovering weak edges (72%). In all cases, the estimator does not add edges not in the true network. This suggests recovered economic links that deviate from geographic links may indeed carry signal, while weak links may not get detected. Finally, the estimator for ρ and γ may show some downward bias with the sample sizes in the application, consistent with the simulations in Appendix Figure A1.

4.1 Data and Empirical Specification

We denote state tax liabilities for state i in year t as τ_{it} , covering state taxes collected from real per capita income, sales, and corporate taxes. We extend the sample used by Besley and Case (1995a), that runs from 1962-1988 ($T = 26$).²⁹ The outcome considered, $\Delta\tau_{it}$, is the change in tax liabilities between years t and $(t - 2)$ because it might take a governor more than a year to implement a tax program. Their model implies a standard social interactions specification for the tax setting behavior of state governors:

$$\Delta\tau_{it} = \rho_0 \sum_{j=1}^N W_{0,ij} \Delta\tau_{jt} + \sum_{k=1}^K \sum_{j=1}^N W_{0,ij} x_{jkt} \gamma_{0,k} + \sum_{k=1}^K \beta_{0,k} x_{ikt} + \alpha_i + \alpha_t + \epsilon_{it}, \quad (12)$$

where $k = 1, \dots, K$ are the covariates for state i in period t . Tax setting behavior is determined by (i) endogenous social effects arising through neighbors' tax changes ($\sum_{j=1}^N W_{0,ij} \Delta\tau_{jt}$); (ii) exogenous social effects arising through neighbors' characteristics ($\sum_{j=1}^N W_{0,ij} x_{jkt}$); and (iii) state i 's characteristics (x_{ikt}), including income per capita, the unemployment rate, and the proportions of young and elderly in the state's population. All specifications include state and time effects (α_i , α_t). Due to the inclusion of the time effects α_t , we normalize the rows of W_{econ} to one. Table A6 presents descriptive statistics for the Besley and Case (1995a) sample and our extended sample.

Much of the earlier literature on tax competition has focused on endogenous social effects and ignored exogenous social effects by setting $\gamma = 0$. Our identification result allows us to relax this restriction and estimate the full typology of social effects described by Manski (1993). This is important because only endogenous social effects lead to social multipliers from tax competition, and they are crucial to identify as they can lead to a race-to-the-bottom or suboptimal public goods provision (Brennan and Buchanan, 1980; Wilson, 1986; Oates and Schwab, 1988).

²⁹Besley and Case (1995a) test their political agency model using a two-equation set-up: (i) on gubernatorial re-election probabilities and (ii) on tax setting. Our application focuses on the latter because this represents a social interaction problem. They use two tax series: (i) TAXSIM data (from the NBER), which runs from 1977-1988 and (ii) state tax liabilities series constructed from data published annually in the Statistical Abstract of the US, that runs from 1962-1988. All their results are robust to either series. We extend the second series.

4.2 Preliminary Findings

Table 1 presents our preliminary findings and comparison to Besley and Case (1995a). Throughout this section, we refer to “OLS estimates” as the estimates of the main equation (12) when W_0 is postulated as W_{geo} or W_{econ} and ρ_0 , $\gamma_{0,k}$, and $\beta_{0,k}$ are estimated by OLS.³⁰ Column 1 shows those estimates where the postulated social interactions matrix is based on geographic neighbors, exogenous social effects are ignored and the panel includes all 48 mainland states but runs only from 1962-1988 as in Besley and Case (1995a). Social interactions influence gubernatorial tax setting behavior: $\hat{\rho}_{OLS} = .375$. Column 2 shows this to be robust to instrumenting neighbors’ tax changes using the instrument set proposed by Besley and Case (1995a): namely, instrumenting for $\Delta\tau_{jt}$ using geographic neighbors’ lagged changes in per capita income and unemployment rates. These instruments are in the spirit of using exogenous social effects to instrument for neighbors’ tax changes. $\hat{\rho}_{2SLS}$ is more than double the magnitude of $\hat{\rho}_{OLS}$, suggesting tax setting behaviors across jurisdictions are strategic complements.

Columns 3 and 4 replicate both specifications over the longer sample, confirming Besley and Case’s (1995) finding on social interactions to be robust. $\hat{\rho}_{2SLS}$ is again more than double $\hat{\rho}_{OLS}$. The result in Column 4 implies that for every dollar increase in the average tax rates among geographic neighbors, a state increases its own taxes by 64 cents. This is similar to the headline estimate of Besley and Case (1995a).³¹

4.3 Endogenous and Exogenous Social Interactions

We now move beyond much of the earlier literature to first establish whether there are endogenous and exogenous social interactions in tax setting. We first focus on the endogenous and exogenous social interaction parameters, and in the next subsection, we detail the identified social interactions matrix, \hat{W}_{econ} . To do so, we need to modify slightly how we instrument for neighbors’ tax changes: the instrument set proposed by Besley and Case (1995a) based on geographic neighbors’ characteristics will generally be weaker when estimating the full specification in (12) because the instruments are now directly controlled for in (12). We use an Adaptive Elastic Net GMM approach, which instruments neighbors’ tax changes with the characteristics of all other states. With the inclusion of endogenous and exogenous social effects, this represents our preferred approach.

Columns 1 and 2 of Table 2 show OLS and GMM estimates for ρ obtained from the Adaptive Elastic Net procedure, where we still set $\gamma = 0$ but use our preferred instrument set: $\hat{\rho}_{GMM} = .709 >$

³⁰We postulate that W is W_{econ} obtained by running the procedure in Section 3, retrieving \hat{W} , and re-running model (12) with $W = \hat{W}$. For such OLS estimates, we use robust standard errors and ignore the sampling uncertainty in the estimated W_{econ} .

³¹Nor is the magnitude very different from earlier work examining fiscal expenditure spillovers. For example, Case *et al.* (1989) find that US state governments’ levels of per-capita expenditure are significantly impacted by the expenditures of their neighbors, with a one-dollar increase in neighbors’ expenditures leading to a seventy-cent increase in own-state expenditures.

$\hat{\rho}_{OLS} = .649$. Columns 3 and 4 estimate the full model in (12). Relative to when exogenous social effects are assumed away ($\gamma = 0$), the OLS and GMM estimates of ρ are smaller, but we continue to find robust evidence of endogenous social interactions in tax setting. The specification in Column 4 represents our preferred one: $\hat{\rho}_{GMM} = .452$ (with a standard error of .132). This value meets the requirements on ρ in Corollaries 3 and 4 for global identification.³²

4.4 Identified Social Interactions Matrix

Figure 1B shows how the structures of economic (\hat{W}_{econ}) and geographic networks (W_{geo}) differ, where connected edges imply that two states are linked in at least one direction (state i causally impacts state taxes in j , and/or *vice versa*). This comparison makes clear whether all states geographically adjacent to i matter for its tax setting behavior and whether there are non-adjacent states that influence its tax rate.

The left-hand panel of Figure 1B shows the network of geographic neighbors (whose edges are colored blue), onto which we superimpose edges *not* identified as links in W_{econ} ; dropped edges are in red. The vast majority of geographically adjacent states are irrelevant for tax setting behavior. The right-hand panel of Figure 1B adds new edges identified in \hat{W}_{econ} that are *not* part of W_{geo} ; these added edges are in green and represent non-geographically adjacent states through which social interactions occur. For tax-setting behavior, economic distance is imperfectly measured if we simply assume interactions depend only on physical distance.

Table 3 summarizes the comparison between W_{geo} and \hat{W}_{econ} . W_{geo} has 214 edges, while \hat{W}_{econ} has only 49. \hat{W}_{econ} and W_{geo} have 9 edges in common. Hence, the vast majority of geographical neighbors ($205/214 = 96\%$) are not relevant for tax setting. \hat{W}_{econ} has 40 edges that are absent in W_{geo} , and the identified social interactions are more spatially dispersed than under the assumption of geographic networks. This is reflected in the far lower clustering coefficient in \hat{W}_{econ} than in W_{geo} (.042 versus .419).³³

4.5 Links and Reciprocity

Our estimation strategy identifies the continuous strength of links, $W_{0,ij}$, where $W_{0,ij} > 0$ is interpreted as state j influencing outcomes in state i . This is useful because recent developments in tax competition theory, using insights from the social networks literature, suggest links need not be reciprocal (Janeba and Osterleh, 2013).

Table 3 reveals that only 12.2% of edges in \hat{W}_{econ} are reciprocal (all edges in W_{geo} are reciprocal by definition). Hence, tax competition is both spatially disperse and asymmetric. In most cases

³²Table A7 shows the full set of exogenous social effects (so Columns 1 and 2 refer to the same specifications as Columns 3 and 4 in Table 2). Exogenous social effects operate through economic neighbors' unemployment rate, demographic characteristics, and their governor's age.

³³The clustering coefficient is the frequency of the number of fully connected triplets over the total number of triplets.

where tax setting in state i is influenced by taxes in state j , the opposite is not true.

Given common time shocks α_t in (12), row-sum normalization is required and ensures $\sum_j W_{0,ij} = 1$. Hence, for every state i , there will be at least one economic neighbor state j^* that impacts it, so $W_{0,ij^*} > 0$. This just reiterates that social interactions matter. On the other hand, our procedure imposes no restriction on the derived columns of \hat{W}_{econ} . It could be that a state does not affect any other state. To see this in more detail, the final rows of Table 3 report the degree distribution across states, splitting for in-networks and out-networks. In W_{geo} , the in-degree is by construction equal to the out-degree, as all ties are reciprocal. The greater sparsity of the network of economic neighbors is reflected in the degree distribution being lower for \hat{W}_{econ} than W_{geo} . In \hat{W}_{econ} , the dispersion of in- and out-degree networks is very different (as measured by the standard deviation), being nearly nine times higher for the in-degree. Hence one reason for so few reciprocal ties being in the economic network is that out-degree network ties are rarely also in-degree ties.

This asymmetry in \hat{W}_{econ} further suggests that some highly influential states drive tax setting behavior in other states. To see which states these are, Figure 2 shows a histogram for the number of out-degree links from states. Twenty states have an out-degree of zero, so their tax rates have no direct impact on any other state's tax setting behavior. The most influential states in terms of the highest out-degree are Alabama (directly impacts tax setting behavior in five other states) and South Carolina, Pennsylvania, and Montana (which each directly impact tax setting behavior in four other states). Taking South Carolina as an example, the four states that it directly impacts include its geographic neighbor, Georgia, as well as non-geographic neighbors Missouri, Montana, and Virginia.

4.6 Factor Mobility or Yardstick Competition?

We use \hat{W}_{econ} to shed light on the roles of factor mobility and yardstick competition in driving tax competition. To do so, we estimate the factors correlated with the existence of links between states i and j in \hat{W}_{econ} relative to \hat{W}_{geo} . For state pairs with non-zero links in either \hat{W}_{econ} or \hat{W}_{geo} , we define a dummy outcome $\hat{W}_{econ,ij} = 1$ if a link between states i and j is estimated under \hat{W}_{econ} and $\hat{W}_{econ,ij} = 0$ if a link between states i and j exists under \hat{W}_{geo} but not under \hat{W}_{econ} . We examine correlates of links using the following dyadic regression:

$$\hat{W}_{econ,ij} = \lambda_0 + \lambda_1 X_{ij} + \lambda_2 X_i + \lambda_3 X_j + u_{ij}, \quad (13)$$

estimated using a linear probability model. The elements X_{ij} , X_i , and X_j correspond to characteristics of the pair of states (i, j) , state i , and state j , respectively. Covariates are time-averaged over the sample period, and robust standard errors are reported.

Table 4 presents the dyadic regression results. Column 1 controls only for the distance between states i and j : this is highly predictive of an economic link between them. This reflects that the

economic network of state i often comprises states are in the same region, but not necessarily contiguous to state i . Column 2 adds two X_{ij} covariates to capture the economic and demographic homophily between states i and j . GDP homophily is the absolute difference in the states' GDP per capita. Demographic homophily is the absolute difference in the share of young people (aged 5-17) plus the absolute difference of the share of elderly people (aged 65+) across the states. GDP homophily does not predict economic ties, whereas demographic homophily does.

Columns 3-5 then sequentially add in several sets of controls. For labor mobility, we use net state-to-state migration data to control for the net migration flow of individuals from state i to state i (defined as the flow from i to j minus the flow from j to i).³⁴ We then add a political homophily variable between states. For any given year, this is set to one if a pair of states have governors of the same political party. As this is time-averaged over our sample, this element captures the share of the sample in which the states have governors of the same party. Lastly, we include whether state j is considered a tax haven (and so might have disproportionate influence on other states). Based on Findley *et al.* (2012), the following states are coded as tax havens: Nevada, Delaware, Montana, South Dakota, Wyoming, and New York.

Column 5 shows that with this full set of controls, distance remains a robust predictor of the existence of economic links between states. However, the identified economic network highlights additional significant predictors of tax competition between states: political homophily *reduces* the likelihood of a link, suggesting any yardstick competition driving social interactions occurs when voters compare their governor to those of the opposing party in other states. Tax haven states appear to be especially less influential in the tax setting behaviors of other states. This mirrors what was observed in Figure 2, where some of the prominent tax havens – Nevada, Delaware, and New York – were all identified to have zero out-degree links. The relatively weak influence of tax haven states eases concerns over a race-to-the-bottom in tax setting behaviors.

Column 6 controls for state i and state j fixed effects. This reinforces the idea that distance and political homophily correlate to the strength of influence states tax setting has on others (the tax haven dummy cannot be separately identified in this specification). Labor mobility between states does not robustly predict the existence of economic ties.

4.7 Dynamics

As in our identification result, our empirical approach has taken the network structure as fixed over the entire sample period. In the context of tax competition over our study period, this might be a strong assumption. We examine the issue in more detail by allowing the estimated W_{econ} matrix

³⁴We also experimented with alternative measures of labor migration, and the results were qualitatively the same. State-to-state migration data are based on year-to-year address changes reported on individual income tax returns filed with the IRS. The data cover filing years 1991 through 2015 and include the number of returns filed, which approximates the number of households that migrated, and the number of personal exemptions claimed, which approximates the number of individuals who migrated. The data are available at <https://www.irs.gov/statistics/soi-tax-stats-migration-data> (accessed September 2017).

to vary over time by changing the weight placed on observations from any given time period. More precisely, for any given time period t , we weight observations using a Gaussian kernel with its center varying period-by-period from 1962 to 2015. The variance of the kernel is set such that 75% of the weight is given to the first half of the data (i.e., pre-1988) when the kernel is centered in 1962. Figure A4 shows the kernel employed as we vary its center: the solid kernel is centered in 1962, the start of our sample – when we place the most weight on observations from 1962. The static case considered previously is akin to using a uniform kernel over all periods. This kernel weighting approach is outlined in Section 2.3.3. We fully describe the algorithm in Appendix Section C.2.

We begin by considering time-varying estimates of the endogenous and exogenous social interaction parameters from the full model in (12). The results for the endogenous social interactions parameter are shown in Figure 3, where the shaded areas are the 95% confidence intervals of the period-by-period estimates. Panel A shows that OLS estimates of $\hat{\rho}$ drift up over time, so the strength of endogenous interactions increases from around .35 in the late 1960s to .50 by the 2010s. In all periods, we can reject the null that the endogenous social effect is zero. Recall the earlier static estimate was $\hat{\rho}_{OLS} = .375$.

Panel B shows the estimated endogenous social effect when we use GMM based on the characteristics of all other states as IVs. This also drifts up from around .35 in the late 1960s to .50 by the 2010s. In the majority of periods, we can reject the null of no endogenous social effect.³⁵

Figure 4 shows the evolution of \hat{W}_{econ} over time as we center the kernel on different periods, following the same color-coding as in Figure 1. In all periods, geography-based edges play little role, and over time, the economic network becomes denser. This highlights not only that economic networks for tax competition always differ starkly from geography-based networks, but that the nature of economic networks relevant to tax competition has changed steadily over time.

Figure 5 shows how the features of \hat{W}_{econ} evolve as we place greater weight from early to later periods. For each statistic, we plot the period-by-period estimate when we center the kernel in any given period. The resulting smoothed estimates are then shown. To ease exposition, networks edges with $W_{0,ij} < 1/47$ are removed. This cutoff is chosen as, in theory, states can only link at maximum with 47 other states. Panel A shows the share of edges that are kept from the previous estimate (centered in the previous period). We see relatively high stability in \hat{W}_{econ} with the smoothed estimate suggesting more than 60% of edges always being kept from one estimate to the next, with this stability increasing from the late 1980s.

Panel B shows how the overlap between \hat{W}_{econ} and W_{geo} varies over time, as measured by the share of edges that are only present in \hat{W}_{econ} . There is little overlap between the two networks over the entire sample. The smoothed estimate suggests that at least 80% of identified edges in \hat{W}_{econ} are never in W_{geo} . The divergence between economic and geographic neighbors becomes starker from the mid-1980s onwards.

³⁵Standard errors are estimated without imposing the restriction that parameters vary slowly over time, and fluctuations across periods reflect variations in the network across periods.

Panels C and D show how the clustering and reciprocity of links in \hat{W}_{econ} vary as we shift the weight to later observations. Clustering of \hat{W}_{econ} increases from the 1960s through to the early 2000s. Thereafter, social interactions in tax competition become sparser. We also observe a reversal in the extent to which social interactions are reciprocal, with reciprocity rising to a peak in the early 1980s – when 20% of ties were reciprocal – and slowly falling thereafter.

Taken together, the results suggest the nature of tax competition between US states has changed over time through two mechanisms: (i) the strength of endogenous social interactions ($\hat{\rho}$) has increased over time and (ii) the network of states interacted with (\hat{W}_{econ}) varies over time. This has important implications for policy evaluation: the same intervention might have different spillover effects if implemented at different moments in time due to the evolution of $\hat{\rho}$ and \hat{W}_{econ} . We consider this next using counterfactual simulations.

4.8 Counterfactuals

We use a counterfactual exercise to contrast how shocks to tax setting in a given state propagate under \hat{W}_{econ} , relative to what would have been predicted under W_{geo} . We do so for both static and dynamic estimates of \hat{W}_{econ} . We focus on South Carolina (SC), a state with one of the highest out-degree, as shown in Figure 2. We consider a scenario in which SC exogenously increases its taxes per capita by 10%. We measure the differential change in equilibrium state taxes in state j under the two network structures using the following statistic:

$$\Upsilon_j = \log(\Delta\tau_{jt}|\hat{W}_{econ}) - \log(\Delta\tau_{jt}|W_{geo}), \quad (14)$$

so that positive (negative) values imply equilibrium taxes being higher (lower) under \hat{W}_{econ} .³⁶

Starting with the static case, Panel A of Figure 6 shows for each mainland US state the spillover effects through the economic network of tax competition. This highlights positive spillovers on tax rates in many states that are not geographic neighbors of SC. Panel B graphs Υ_j to make precise how spillovers derived from \hat{W}_{econ} diverge from those predicted under W_{geo} . In 26 states, Υ_j is smaller than .01% because both networks predict negligible spillovers to those states. In the remaining 22 mainland states, there is a wide discrepancy between the equilibrium state tax rates predicted under \hat{W}_{econ} relative to W_{geo} : Υ_j varies from -1 to 4.03 . The long-run effect in SC itself is also higher under \hat{W}_{econ} than under W_{geo} . The former states that given feedback effects, the long-run increase in tax rates in SC from a 10% increase is 11.4%, while the geographically based network implies a smaller equilibrium increase of 10.3%.

As \hat{W}_{econ} is spatially more dispersed than W_{geo} , the general equilibrium effects are different under the two network structures. Table 5 summarizes the general equilibrium implications for tax inequality under \hat{W}_{econ} and the W_{geo} counterfactual. The average tax rate increase under

³⁶For W_{geo} , we calculate the counterfactual at $\hat{\rho}_{GMM} = .452$, the endogenous effect parameter estimated in our preferred specification, Column 4 of Table 2.

\hat{W}_{econ} is three times that estimated under W_{geo} . Moreover, the dispersion of tax rates across states increases under \hat{W}_{econ} relative to W_{geo} . Finally, assuming interactions are based solely on geographic neighbors, we miss the fact that many states have relatively small tax increases.

We can repeat the exercise using the dynamically estimated economic network. Throughout, we calculate the general equilibrium effects of the *same* policy experiment: SC increasing its taxes per capita by 10%. These general equilibrium effects vary over the sample period because the strength of social interactions in tax competition vary ($\hat{\rho}_{GMM}$), as shown in Figure 3, and identified economic neighbors vary over time (\hat{W}_{econ}), as shown in Figure 4. The results are summarized in Figure 7. Placing weight on the early or later part of the sample generates similar changes in average tax rates and their variance in general equilibrium – with both being lower than simulated under the static model. Placing more weight on the middle of the sample period generates higher changes in average tax rates and their variance in general equilibrium.

The differential general equilibrium impacts found as we place different weights across sample observations links to recent discussions on the external validity of internally valid causal impacts based on micro-evidence. While the earlier literature has emphasized the potential interaction of treatment effects with aggregate shocks (Rosenzweig and Udry, 2020) or how behavioral responses to social insurance policies vary over the business cycle (Kroft and Notowidigdo, 2016), our analysis provides another explanation for the changing impacts of policies where social interactions determine behavior: changes in the strength of social interactions and the network of economic interactions.

5 Discussion

In a canonical social interactions model, we provide sufficient conditions under which the social interactions matrix, and endogenous and exogenous social effects are all globally identified, even absent information on social links. Our identification strategy is novel and may bear fruit in other areas. The method is immediately applicable to other classic social interactions problems, but where data on social links is either missing or partial. In fields such as macroeconomics, political economy, and trade, there are core areas of research where social interactions across jurisdictions/countries etc. drive key outcomes, panel data exist over many periods, and the number of nodes is relatively fixed. Moreover, while our discussion and application have focused on a continuous policy response (state taxes), our methods can also be applied to the extensive margin of policy adoption and diffusion. Such diffusion models might generate network interactions where some states influence the later adoption of economic and social policies in other jurisdictions. This issue is studied by DellaVigna and Kim (2022) in the context of US state policies – they examine the diffusion of over 700 policies in the past 70 years. Their work also suggests the nature of interactions across states has changed: while geographic proximity is a good predictor of policy diffusion, they also find that since 2000, political alignment across states has become the strongest

predictor of diffusion.

In finance, high-frequency panel data is readily available and relevant for the study of core research questions. For example, a long-standing question has been whether CEOs are subject to relative performance evaluation, and if so, what is the comparison set of firms/CEOs used (Edmans and Gabaix, 2016). More generally, our method can be readily applied to a large class of economic questions around contagion, risk, and the fragility of economic and environmental systems. For example, since the financial crisis of 2008, it has become clear that linkages between actors such as firms or banks are complex and often hidden, yet because endogenous network interactions cause feedback loops and have multiplier effects, they can have enormous implications for the evolution of a financial crisis or the propagation of supply shocks in aggregate. Identifying such synchronicity is a critical first step to putting in place policies to reduce the fragility of economic systems (van Vliet, 2018; Elliott and Golub, 2022; Goldstein *et al.*, 2022).

Advances in the availability of administrative data, data from social media, mobile technologies, and online economic transactions all offer new possibilities to identify social interactions with long panels or high-frequency data collection, where data on social ties will typically be missing.

Three further directions for future research are of note. First, under partial observability of W_0 (as in Blume *et al.*, 2015), the number of parameters in W_0 to be retrieved falls quickly. Our approach can then still be applied to complete knowledge of W_0 , such as if Aggregate Relational Data is available, and this could be achieved with potentially weaker assumptions for identification, and in even shorter panels. To illustrate possibilities, Figure A5 shows results from a final simulation exercise in which we assume the researcher starts with partial knowledge of W_0 . We do so for Banerjee *et al.* (2013) village family network, showing simulation results for scenarios in which the researcher knows the social ties of the three (five, ten) households with the highest out-degree. For comparison, we also show the earlier simulation results when W_0 is entirely unknown. This clearly illustrates that with partial knowledge of the social network, performance on all metrics improves rapidly for any given T .

Second, we have developed our approach in the context of the canonical linear social interactions model (1). This builds on Manski (1993) when W_0 is known to the researcher, and the reflection problem is the main challenge in identifying endogenous and exogenous social effects. However, the reflection problem is functional-form dependent and may not apply to many non-linear models (Blume *et al.*, 2011, Blume *et al.*, 2015). An important topic for future research is to extend the analysis to non-linear social interaction settings. Relatedly, the canonical social interaction model assumes that the same W_0 governs the endogenous and exogenous channels. Despite the relaxation we propose in Section 2.3.3, we see this as a limitation of the current method, and future research is needed to allow for a fully flexible approach.

Finally, an important part of the social networks literature examines endogenous network formation (Jackson *et al.*, 2017; de Paula, 2017). Our analysis allows us to begin probing the issue in two ways. First, the kind of dyadic regression analysis in Section 4 on the correlates of entries

in $W_{0,ij}$ suggests factors driving link formation and dissolution. Second, this leads naturally to a broad agenda going forward, to address the challenge of simultaneously identifying and estimating time varying models of network formation and social interaction, all in cases where data on social networks is not required.

Data Availability Statement

The data and code underlying this research is available on Zenodo at <https://dx.doi.org/10.5281/zenodo.10160768>.

References

- ACEMOGLU, D., V. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012). The Network Origins of Aggregate Fluctuations. *Econometrica*, 80, 1977–2016.
- ANSELIN, L. (2010). Thirty Years of Spatial Econometrics. *Papers in Regional Science*, 89, 3–25.
- ATALAY, E., A. HORTACSU, J. ROBERTS, AND C. SYVERSON (2011). Network Structure of Production. *Proceedings of the American Mathematical Society*, 108, 5199–202.
- BALLESTER, C., A. CALVO-ARMENDOL, AND Y. ZENOU (2006). Who’s Who in Networks. Wanted: The Key Player. *Econometrica*, 74, 1403–17.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013). The Diffusion of Microfinance. *Science*, 341, 1236498.
- BATTAGLINI, M., E. PATACCHINI, AND E. RAINONE (2022). Endogenous social interactions with unobserved networks. *The Review of Economic Studies*, 89, 1694–1747.
- BELLONI, A. AND V. CHERNOZHUKOV (2011). High Dimensional Sparse Models: an Introduction. Tech. Rep. 1106.5242, arXiv.org Collection, <http://arxiv.org/abs/1106.5242>.
- BESLEY, T. AND A. CASE (1994). Unnatural Experiments? Estimating the Incidence of Endogenous Policies. *NBER Working Paper 4956*.
- (1995a). Incumbent Behavior: Vote-seeking, Tax-setting, and Yardstick Competition. *American Economic Review*, 85, 25–45.
- (1995b). Incumbent Behavior: Vote-seeking, Tax-setting, and Yardstick Competition. Unpublished data.
- BLUME, L., W. A. BROCK, S. N. DURLAUF, AND Y. IOANNIDES (2011). Identification of Social Interactions. in *Handbook of Social Economics*, ed. by J. Behabib, A. Bisin, and M. O. Jackson, North-Holland, vol. 1B.
- BLUME, L. E., W. A. BROCK, S. N. DURLAUF, AND R. JAYARAMAN (2015). Linear Social Interactions Models. *Journal of Political Economy*, 123, 444–96.

- BONALDI, P., A. HORTACSU, AND J. KASTL (2015). An Empirical Analysis of Funding Costs Spillovers in the EURO-zone with Application to Systemic Risk. Princeton University Working Paper.
- BRAMOULLÉ, Y., H. DJEBBARI, AND B. FORTIN (2009). Identification of Peer Effects Through Social Networks. *Journal of Econometrics*, 150, 41–55.
- BRENNAN, G. AND J. BUCHANAN (1980). *The Power to Tax: Analytical Foundations of a Fiscal Constitution*, Cambridge University Press.
- BREZA, E., A. G. CHANDRASEKHAR, T. H. MCCORMICK, AND M. PAN (2020). Using aggregated relational data to feasibly identify network structure without network data. *American Economic Review*, 110, 2454–84.
- BURSZTYN, L., F. EDERER, B. FERMAN, AND N. YUCHTMAN (2014). Understanding Mechanisms Underlying Peer Effects: Evidence From a Field Experiment on Financial Decisions. *Econometrica*, 82, 1273–301.
- CANER, M. AND H. H. ZHANG (2014). Adaptive Elastic Net for Generalized Method of Moments. *Journal of Business and Economic Statistics*, 32, 30–47.
- CASE, A., J. R. HINES, AND H. S. ROSEN (1989). Copycatting: Fiscal Policies of States and Their Neighbors. *NBER Working Paper 3032*.
- CHANDRASEKHAR, A. AND R. LEWIS (2016). Econometrics of Sampled Networks. Working Paper.
- CHANEY, T. (2014). The Network Structure of International Trade. *American Economic Review*, 104, 3600–34.
- COLEMAN, J. S. (1964). *Introduction to Mathematical Sociology*, London Free Press Glencoe.
- CONLEY, T. G. AND C. R. UDRY (2010). Learning About a New Technology: Pineapple in Ghana. *American Economic Review*, 100, 35–69.
- DAHLHAUS, R. (2012). 13 - Locally Stationary Processes. in *Time Series Analysis: Methods and Applications*, ed. by T. Subba Rao, S. Subba Rao, and C. Rao, Elsevier, vol. 30 of *Handbook of Statistics*, 351–413.
- DE GIORGI, G., M. PELLIZZARI, AND S. REDAELLI (2010). Identification of Social Interactions through Partially Overlapping Peer Groups. *American Economic Journal: Applied Economics*, 2, 241–75.
- DE PAULA, A. (2017). Econometrics of Network Models. in *Advances in Economics and Econometrics: Theory and Applications*, ed. by B. Honore, A. Pakes, M. Piazzesi, and L. Samuelson, Cambridge University Press.
- DE PAULA, A., E. TAMER, AND S. RICHARDS-SHUBIK (2018). Identifying Preferences in Networks with Bounded Degree. *Econometrica*, 86, 263–288.
- DELLAVIGNA, S. AND W. KIM (2022). Policy Diffusion and Polarization across US States. *National Bureau of Economic Research*.

- EDMANS, A. AND X. GABAIX (2016). Executive Compensation: a modern primer. *Journal of Economic literature*, 54, 1232–87.
- EFRON, B., T. HASTIE, I. JOHNSTONE, AND R. TIBSHIRANI (2004). Least angle regression. *The Annals of statistics*, 32, 407–499.
- ELLIOTT, M. AND B. GOLUB (2022). Networks and Economic Fragility. *Annual Review of Economics*, 14, 665–696.
- FETZER, T., P. C. SOUZA, O. VANDEN EYNDE, AND A. L. WRIGHT (2021). Security transitions. *American Economic Review*, 111, 2275–2308.
- FINDLEY, M., D. NIELSON, AND J. SHARMAN (2012). Global Shell Games: Testing Money Launderers’ and Terrorist Financiers’ Access to Shell Companies. Griffith University Working Paper.
- GAUTIER, E. AND C. ROSE (2016). Inference in Social Effects when the Network is Sparse and Unknown. (in preparation).
- GAUTIER, E. AND A. TSYBAKOV (2014). High-Dimensional Instrumental Variables Regression and Confidence Sets. Working Paper CREST.
- GEFANG, D., S. G. HALL, AND G. S. TAVLAS (2023). Identifying Spatial Interdependence in Panel Data with Large N and Small T. .
- GOLDSTEIN, I., A. KOPYTOV, L. SHEN, AND H. XIANG (2022). Synchronicity and Fragility. *Manuscript*.
- GRANOVETTER, M. (1973). The Strength of Weak Ties. *American Journal of Sociology*, 6, 1360–80.
- HASTIE, T. AND R. TIBSHIRANI (1993). Varying-Coefficient Models. *Journal of the Royal Statistical Society. Series B (Methodological)*, 55, 757–796.
- HOFF, P. D., A. E. RAFTERY, AND M. S. HANDCOCK (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97, 1090–1098.
- HOLLAND, P. W. AND S. LEINHARDT (1981). An Exponential Family of Probability Distributions for Directed Graphs. *Journal of the American Statistical Association*, 76, 33–50.
- HORN, R. A. AND C. R. JOHNSON (2013). *Matrix Analysis*, Cambridge University Press.
- JACKSON, M., B. ROGERS, AND Y. ZENOU (2017). The Economic Consequences of Social Network Structure. *Journal of Economic Literature*, 55, 49–95.
- JACKSON, M. AND A. WOLINSKY (1996). A Strategic Model of Social and Economic Networks. *Journal of Economic Theory*, 71, 44–74.
- JANEBA, E. AND S. OSTERLEH (2013). Tax and the City – A Theory of Local Tax Competition. *Journal of Public Economics*, 106, 89–100.

- KAPETANIOS, G., R. M. MASOLO, K. PETROVA, AND M. WALDRON (2019). A time-varying parameter structural model of the UK economy. *Journal of Economic Dynamics and Control*, 106, 103705.
- KLINE, B. AND E. TAMER (2016). Bayesian Inference in a Class of Partially Identified Models. *Quantitative Economics*, 7, 329–366.
- KRANTZ, S. G. AND H. R. PARKS (2013). *The Implicit Function Theorem*, Birkhauser.
- KROFT, K. AND M. J. NOTOWIDIGDO (2016). Should unemployment insurance vary with the unemployment rate? Theory and evidence. *The Review of Economic Studies*, 83, 1092–1124.
- LAM, C. AND P. C. SOUZA (2016). Detection and Estimation of Block Structure in Spatial Weight Matrix. *Econometric Reviews*, 35, 1347–1376.
- (2020). Estimation and selection of spatial weight matrix in a spatial lag model. *Journal of Business & Economic Statistics*, 38, 693–710.
- LEE, L.-F. (2007). Identification and Estimation of Econometric Models with Group Interactions, Contextual Factors and Fixed Effects. *Journal of Econometrics*, 60, 531–42.
- LEWBEL, A., X. QU, AND X. TAN (2022). Social Networks with Unobserved Links. *Journal of Political Economy* (forthcoming).
- MANRESA, E. (2016). Estimating the Structure of Social Interactions Using Panel Data. Manuscript.
- MANSKI, C. F. (1993). Identification of Endogenous Social Effects: the reflection problem. *The Review of Economic Studies*, 60, 531–42.
- MEINSHAUSEN, N. AND P. BUHLMANN (2006). High-Dimensional Graphs and Variable Selection with the Lasso. *The Annals of Statistics*, 34, 1436–1462.
- NEUMAN, E. J. AND M. S. MIZRUCHI (2010). Structure and bias in the network autocorrelation model. *Social Networks*, 32, 290–300.
- OATES, W. AND R. SCHWAB (1988). Economic Competition Among Jurisdictions: Efficiency-enhancing or Distortion-inducing?. *Journal of Public Economics*, 35, 333–54.
- ROSE, C. (2015). Essays in Applied Microeconometrics. Ph.D. thesis, University of Bristol.
- ROSENZWEIG, M. R. AND C. UDRY (2020). External validity in a stochastic world: Evidence from low-income countries. *The Review of Economic Studies*, 87, 343–381.
- ROTHENBERG, T. (1971). Identification in Parametric Models. *Econometrica*, 39, 577–91.
- ROTHENHÄUSLER, D., C. HEINZE, J. PETERS, AND N. MEINSHAUSEN (2015). BACKSHIFT: Learning causal cyclic graphs from unknown shift interventions. in *Advances in Neural Information Processing Systems*, 1513–1521.
- SACERDOTE, B. (2001). Peer Effects with Random Assingment: Results for Dartmouth Roommates. *The Quarterly Journal of Economics*, 116, 681–704.

- SHLEIFER, A. (1985). A Theory of Yardstick Competition. *Rand Journal of Economics*, 16.
- SMITH, T. E. (2009). Estimation bias in spatial models with strongly connected weight matrices. *Geographical Analysis*, 41, 307–332.
- TIBSHIRANI, R. J. AND J. TAYLOR (2012). Degrees of freedom in lasso problems. *The Annals of Statistics*, 40, 1198–1232.
- TIEBOUT, C. (1956). A Pure Theory of Local Expenditures. *Journal of Political Economy*, 64, 416–24.
- U.S. BUREAU OF ECONOMIC ANALYSIS (2017). State Annual Personal Income Summary. Accessed April 17, 2017.
- U.S. BUREAU OF LABOR STATISTICS (2017). State’s Unemployment Rate. Retrieved from www.usgovernmentdebt.us. Accessed April 17, 2017.
- U.S. CENSUS BUREAU (2016a). Census Population & Housing Data. Accessed April 17, 2017.
- (2016b). State Government Tax Tables, 1992-2015. Accessed April 17, 2017.
- VAN VLIET, W. (2018). Connections as Jumps: Estimating Financial Interconnectedness from Market Data. Manuscript.
- WANG, W., E. J. NEUMAN, AND D. A. NEWMAN (2014). Statistical power of the social network autocorrelation model. *Social Networks*, 38, 88–99.
- WIKIPEDIA (2017). List of State Governors of the United States. Accessed April 17, 2017.
- WILSON, J. (1986). A Theory of Interregional Tax Competition. *Journal of Urban Economics*, 19, 296–315.
- (1999). Theories of Tax Competition. *National Tax Journal*, 52, 269–304.
- ZHOU, W. (2019). A Network Social Interaction Model with Heterogeneous Links. *Economics Letters*, 180, 50–53.
- ZOU, H. (2006). The Adaptive Lasso and Its Oracle Properties. *Journal of the American Statistical Association*, 101, 1418–29.
- ZOU, H. AND T. HASTIE (2005). Regularization and Variable Selection via the Elastic Net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67, 301–320.
- ZOU, H., T. HASTIE, AND R. TIBSHIRANI (2007). On the “Degrees of Freedom” of the LASSO. *The Annals of Statistics*, 35, 2173–2192.
- ZOU, H. AND H. H. ZHANG (2009). On the Adaptive Elastic-net with a Diverging Number of Parameters. *Ann. Statist.*, 37, 1733–51.