

NBER WORKING PAPER SERIES

SEEING THE FOREST FOR THE TREES? AN INVESTIGATION OF NETWORK
KNOWLEDGE

Emily Breza
Arun G. Chandrasekhar
Alireza Tahbaz-Salehi

Working Paper 24359
<http://www.nber.org/papers/w24359>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2018

We thank Abhijit Banerjee, Joshua Blumenstock, Steven Durlauf, Paul Goldsmith-Pinkham, Ben Golub, Matt Jackson, Cynthia Kinnan, Pooya Molavi, Melanie Morten, Suresh Naidu, Xu Tan, and Leeat Yariv for helpful comments. Financial support from the NSF under grants SES-1156182 and SES-1326661 is gratefully acknowledged. We thank Shobha Dundi, Devika Lakhote, Tithee Mukhopadhyay, Gowri Nagraj, and Ambika Sharma for excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Emily Breza, Arun G. Chandrasekhar, and Alireza Tahbaz-Salehi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Seeing the Forest for the Trees? An Investigation of Network Knowledge
Emily Breza, Arun G. Chandrasekhar, and Alireza Tahbaz-Salehi
NBER Working Paper No. 24359
February 2018
JEL No. C8,D8,D85,L14,O1

ABSTRACT

This paper assesses the empirical content of one of the most prevalent assumptions in the economics of networks literature, namely the assumption that decision makers have full knowledge about the networks they interact on. Using network data from 75 villages, we ask 4,554 individuals to assess whether five randomly chosen pairs of households in their village are linked through financial, social, and informational relationships. We find that network knowledge is low and highly localized, declining steeply with the pair's network distance to the respondent. 46% of respondents are not even able to offer a guess about the status of a potential link between a given pair of individuals. Even when willing to offer a guess, respondents can only correctly identify the links 37% of the time. We also find that a one-step increase in the social distance to the pair corresponds to a 10pp increase in the probability of misidentifying the link. We then investigate the theoretical implications of this assumption by showing that the predictions of various models change substantially if agents behave under the more realistic assumption of incomplete knowledge about the network. Taken together, our results suggest that the assumption of full network knowledge (i) may serve as a poor approximation to the real world and (ii) is not innocuous: allowing for incomplete network knowledge may have first-order implications for a range of qualitative and quantitative results in various contexts.

Emily Breza
Harvard University
Littauer Center, M28
1805 Cambridge Street
Cambridge, MA 02138
and NBER
ebreza@fas.harvard.edu

Alireza Tahbaz-Salehi
Kellogg School of Management
Northwestern University
2211 Campus Drive, Evanston, IL 60208
alirezat@kellogg.northwestern.edu

Arun G. Chandrasekhar
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305
and NBER
arungc@stanford.edu

1. INTRODUCTION

One of the most prevalent assumptions in the network economics literature is that decision makers have correct, complete, and common knowledge about the structure of the networks they interact on. For instance, various models of social learning over networks (Banerjee, 1992; Smith and Sørensen, 2000; Mossel, Sly, and Tamuz, 2015), the literature on network games of strategic complementarities (Ballester, Calvó-Armengol, and Zenou, 2006; Calvó-Armengol, Patacchini, and Zenou, 2009; Candogan, Bimpikis, and Ozdaglar, 2012), as well as the growing literature on the econometrics of network formation (Sheng, 2016; Leung, 2015a; Menzel, 2017; de Paula, Richards-Shubik, and Tamer, 2018) implicitly or explicitly assume that agents have full knowledge about the underlying network structure.

In this paper, we assess the empirical content and theoretical implications of this assumption. Using observational data, we document substantial departures from full network knowledge, underscoring that the assumption of complete network information may serve as a poor approximation to the real world. We then investigate the theoretical implications of relaxing this assumption in each of the three different domains mentioned above — social learning, network games, and estimation of models of network formation. Our results illustrate that, while a sensible first step, this assumption is not innocuous: violations of full network knowledge may have qualitatively and quantitatively important implications for the predictions of each of these models.

Our empirical investigation uses data collected in 75 villages in Karnataka, India, where we have previously collected detailed network data across all 16,500+ households (Banerjee et al., 2016). Our network data consists of information on whether or not a link exists across a number of informational, financial, and social dimensions for over 98% of pairs of households in each village. Against this backdrop, we returned to the 75 villages and asked 4,554 respondents about existence of various forms of linkages between five pairs of individuals from their village. In particular, we first asked the respondents whether they are unable to offer a response — either because they are not certain enough or because they do not know at least one of the individuals in the corresponding pair. If the respondents were able to offer an answer, we then solicited their guess about the link’s existence.

As our main empirical finding, we document a substantial lack of knowledge about the network structure, both in terms of respondents’ uncertainty as well as the extent to which they correctly identified the presence or absence of links in the underlying network.

Consistent with a basic lack of network knowledge, we find that almost half (46%) of respondents reported that they “don’t know” the link status between a given pair (j, k) of individuals in their village. We then regressed the dummy variable capturing respondent’s uncertainty regarding the link on the respondent’s distance to j and k , controlling for the respondent’s eigenvector centrality, the average centrality of j and k , and a number of demographic controls (such as caste, amenities, geography). We find that respondents are

more likely to express uncertainty about the existence of a link between j and k the further they are in the network: a one-step increase in the average distance of j and k to the respondent is associated with a 12.2 percentage point (pp) increase ($p=0$) in the probability of expressing uncertainty. This effect is sizable, specially relative to the mean of 46%.

We then investigate the quality of the respondents' information by regressing whether the respondent correctly identifies the link status of a given pair (j, k) on network distance, while controlling for centralities and demographic covariates. We find that the quality of the respondent's information also declines steeply in social distance: a one-step increase in the average distance of j and k to the respondent is associated with a 9.3pp decline ($p=0$) in the respondent's probability of correctly identifying that j and k are linked in the underlying network. This also captures a sizable effect, especially given that the mean rate of correctly identifying a link's status is only 20%.¹ Even conditional on expressing certainty, we find that being one step further is associated with a 1.1pp to 2.1pp decline ($p=0.005$) in the probability of correctly identifying that j and k are linked.

The bulk of the aforementioned results are not explained by other demographic covariates. We include geographic distance between respondents and jk pairs, dummies for whether the respondent is of the same caste or subcaste, has the same household amenities (e.g., electrification, roof type, ownership) as j and k , interactions of all of these covariates with actual jk link status, as well as village fixed effects and even respondent fixed effects. Across specifications, our conclusions are predominantly unchanged both qualitatively and quantitatively. We then conduct a post-LASSO exercise to determine a sparse set of covariates that predict our key outcomes, namely, the respondent's correct assessment of link status as well as her claim whether to know or not know the status of a given link. Both in terms of correct assessments and degree of certainty, we find that it is the average network distance to j and k that is selected (along with respondent centrality and the average centrality of j and k). In contrast, except for caste and subcaste, all other demographic variables — including geographic distance between the respondent and jk — are not selected. This finding is consistent with a story where agents learn about the structure of the network through interactions on and through the network, where knowledge tends to be localized.

We next turn to the theoretical implications of our empirical findings by showing that relaxing the widespread assumption that the network structure is commonly known can have first-order implications for a range of qualitative and quantitative results in various contexts. As a first application, we focus on a simple model of Bayesian learning over social networks. We show that partial network knowledge may lead to identification problems that can impede learning, even though information would have been efficiently aggregated had individuals faced no uncertainty about the network structure. This rests on the observation that knowledge about the intricate details of the network structure provides the agents with

¹This indicates systematic misspecification: respondents are more likely than not to think a link exists. A back-of-the-envelope calculation shows that they must be overestimating the average degree by at least 38%.

valuable information on how to discern the correlations and redundancies in their neighbors' estimates. Our result thus demonstrates that predictions about learning dynamics may be sensitive to the extent and nature of agents' knowledge about the social network structure.

As a second case study, we illustrate how assuming full network knowledge may lead to biased structural estimates if in fact decision makers face some uncertainty about the underlying network. We focus on the canonical network interaction game of [Ballester et al. \(2006\)](#) with strategic complementarities. This game, which serves as one of the workhorse models for studying strategic network interactions in the literature, allows for simple structural estimates, counterfactuals, and policy prescriptions. We show that uncertainty about the extent of strategic complementarities outside one's neighborhood results in a systematic shift in equilibrium actions (relative to the complete information benchmark). Hence, ignoring the possibility of incomplete network knowledge may lead to biased structural estimates for the degree of complementarities, mis-specified cost-benefit analyses, and potentially counterproductive intervention policies (such as identifying "key players" to target for interventions).

Finally, we shift our attention to investigating the implications of network knowledge for parameter identification in network formation games. We study a fairly standard econometric model of network formation, where the literature has grappled with partial identification of parameters when agents have complete information about the realization of pairwise preference shocks. However, we show that the equilibrium network in the game with incomplete information provides the econometrician with sufficiently rich observations to point-identify all structural parameters. This result suggests that introducing the (more realistic) assumption of incomplete information may transform a model that is fundamentally hard to identify to a straightforwardly estimable model.

Our work is not the first paper to introduce incomplete information into the above mentioned contexts. For instance, in the context of social learning, [Acemoglu et al. \(2011\)](#) and [Lobel and Sadler \(2015\)](#) allow for randomly generated networks of various kinds. Similarly, papers such as [Leung \(2015b\)](#) and more recently [Ridder and Sheng \(2017\)](#) study network formation with incomplete information. Our goal here is primarily pedagogical, strongly motivated by the data. Specifically, our results are meant to clarify the extent to which a potentially unrealistic assumption plays a critical role in a wide variety of applications and how incorporating partial or lack of network knowledge into fairly standard models can lead to vastly different conclusions.

In addition to the network theory and econometrics literatures discussed above, our paper is related to works in sociology that explore limits to knowledge about others' traits as well as perceptions of friendships in networks ([Friedkin, 1983](#); [Krackhardt, 1987, 2014](#)). For instance, using network data among several departments at the University of Chicago and Columbia University, [Friedkin \(1983\)](#) showed that a respondent i was less likely to know about another faculty member j 's current research if the respondent was further from j in

the network. This is in line with findings of [Alatas et al. \(2016\)](#), who show that subjects are much less likely to know about the wealth status of individuals in their village who are further away in the network. Relative to these findings, our paper emphasizes that knowledge about the network itself is similarly limited and localized through the network.²

The rest of the paper is organized as follows. Section 2 describes the setting, data, and sample statistics. Section 3 contains our main empirical findings, where we document how network knowledge relates to the average distance between the respondents and the pair of individuals they are being inquired about. In Section 4, we discuss how our findings regarding limited network knowledge can be relevant for theoretical, applied, and econometric work. Section 5 is a conclusion. All proofs are provided in Appendix A. An online appendix contains some additional empirical results.

2. SETTING AND DATA

2.1. Network data collection. We collected data in 75 villages in Karnataka, India, where we have previously worked. The villages span 5 districts around Bangalore where we have collected network data in the past: Wave I in 2006 ([Banerjee et al., 2013](#)) and Wave II in 2012 ([Banerjee et al., 2016](#)). In 2012, we collected network data from 89% of the over 16,500 households across the 75 villages. These links spanned financial, social, and informational relationships across 12 dimensions.³ Because such a high share of households were surveyed and every household could name links to any other household in the census in each village, we have 98.8% of all links in the resulting undirected, unweighted network. It is against this backdrop that we conducted a subsequent survey to explore network knowledge.

2.2. Knowledge data collection. We conducted a network knowledge questionnaire in 2015 by asking 4,554 randomly selected individuals across the 75 villages about network relationships between various other individuals in their community. We asked every respondent i about network relationships among five distinct pairs j and k . This was stratified in the following way. We ensured that for each i , one jk pair of each distance 1–1.5, distance 2–2.5, distance 3–3.5 distance 4–4.5, and distance 5+ (including unreachables) was selected in the survey, using our pre-existing network data.⁴ For each jk pair we asked about the existence of specific types of network relationships: informational, financial, and social. We

²[Krackhardt \(1987, 2014\)](#) was the first to advocate for collecting data not only about network interactions, but also about others’ perceptions of such interactions. [Krackhardt \(1987\)](#) called these cognitive social structures (CSS). Our data could be interpreted as sampling from CSSs in 75 distinct networks.

³(1) whose house the respondent visits, (2) who visits her house, (3) kin, (4) whom they socialize with, (5) who gives information when there is a medical need, (6) whom they gives advice to, (7) whom they get advice from, (8) whom they lend material goods to, (9) whom they borrow material goods from, (10) whom they borrow money from, (11) whom they lend money to, (12) with whom they go to pray (e.g., at a temple or mosque).

⁴We also asked i about several demographic traits of j and k . These include if they have children, household size, main household occupation, monthly income, television ownership, religious details, political disposition, and land ownership.

analyze responses to informational, financial, and social relationships, denoted by index r .⁵ The resulting object is a graph \mathbf{g}_r for each r .

For every pair jk and relationship type r , we asked i whether $g_{jk,r} = 1$ or $g_{jk,r} = 0$. The respondent could tell us their estimate or tell us that they “Don’t know.” A response of “Don’t know” could arise for two reasons: either the respondent does not feel certain enough to offer a guess or because the respondent does not know who either j or k are. The respondent i ’s response is therefore $h_{i,jk,r} \in \{1, 0, \text{Don’t Know}\}$.⁶

Our outcomes of interest are whether the respondent correctly identified link status, i.e.,

$$y_{i,jk,r} = 1\{h_{i,jk,r} = g_{jk,r}\},$$

and whether they were sure enough to offer a guess to being with, i.e.,

$$\text{DK}_{i,jk,r} = 1\{h_{i,jk,r} = \text{Don’t Know}\}.$$

We are chiefly interested in how $y_{i,jk,r}$ and $\text{DK}_{i,jk,r}$ depend on the network distance between the respondent i and the pair jk , conditional on the centralities of the nodes involved and demographic covariates. While we consider the respondents’ outcomes $y_{i,jk,r} = 1\{h_{i,jk,r} = g_{jk,r}\}$ for each type of relationship r separately in our regressions, we follow previous work and measure distances and centralities in the union network $\mathbf{g} = \cup_r \mathbf{g}_r$, according to which two agents are assumed to be linked if either agent reports having any of the 12 relationship dimensions with the other.⁷

2.3. Sample statistics. Respondent characteristics are summarized in Table 1, Panel A. 59% of the respondents were female and the average age in the sample was 42.55. Males were typically born in the village (51%) whereas most females married in (only 11% were born in the village). The average degree is 19.5 which indicates a sparse network (the average number of households in the village is 196). Panel B describe pair characteristics. 99% of the pairs belonged to the same connected component as the respondent (i.e., the pair was ‘reachable’), and in our stratified sample, 34% of the time they were linked. On average, the distance from i to jk was 2.13, and 59% of the time they were of the same caste category. Panel C presents our knowledge survey outcomes. 46% of the time the respondents expressed that they did not know the linking status of the pairs in question. When willing to offer a guess, the respondents correctly identified the linking status of the pair only 20% of the time. Conditional on a link existing, a correct guess was offered 57%

⁵Specifically, we collapsed some of the Banerjee et al. (2013) survey, leaving us with 3 dimensions: (1) social interaction; (2) borrowing/lending of money; (3) giving/receiving advice before an important decision;

⁶Note that even though the subjects report whether or not they believe a given pair of individuals are linked, we do not observe the respondents’ subjective beliefs (i.e., the probability that i assigns to the existence of a link j and k).

⁷This choice is made to generate the regressors as independent variables because we believe that this generates the most natural definition of social distance for our setting. Individuals might learn about a financial relationships existing in other parts of the network, for example, through financial ties but also through information and social ties.

of the time. This immediately illustrates that respondents systematically overestimated the existence of a link.⁸

3. NETWORK KNOWLEDGE IN THE NETWORK

3.1. Knowledge and distance. Our summary statistics in Table 1 already indicate poor average knowledge about the network. We next explore how the relative position of nodes influences i 's ability to know whether j and k are linked.

To this end, we run core regressions of the following form:

$$\begin{aligned}
 (1) \quad Y_{i,jk,r} = & \alpha + \theta g_{jk,r} + \lambda_r + \beta^{\text{not}} \cdot \left(\frac{\text{dist}(i,j) + \text{dist}(i,k)}{2} \right) \cdot (1 - g_{jk,r}) \\
 & + \beta^{\text{link}} \cdot \left(\frac{\text{dist}(i,j) + \text{dist}(i,k)}{2} \right) \cdot g_{jk,r} \\
 & + \delta^{\text{not}} \cdot \text{Avg. Centrality}_{jk} \cdot (1 - g_{jk,r}) + \delta^{\text{link}} \cdot \text{Avg. Centrality}_{jk} \cdot g_{jk,r} \\
 & + \gamma^{\text{not}} \cdot \text{Centrality}_i \cdot (1 - g_{jk,r}) + \gamma^{\text{link}} \cdot \text{Centrality}_i \cdot g_{jk,r} \\
 & + X'_{i,jk} \eta + \epsilon_{i,jk,r}.
 \end{aligned}$$

Here, $Y_{i,jk,r}$ will either be $y_{i,jk,r}$, which is a dummy that measures whether i correctly identified the link status between j and k , or $\text{DK}_{i,jk}$, which is whether i reported not knowing the status.

Taking for now the case where $y_{i,jk,r}$ is the outcome variable of interest, coefficient θ measures the difference in the probability of being correct if the link $g_{jk,r}$ exists, whereas coefficients λ_r are relationship type-fixed effects. Coefficient β^{not} measures the marginal change in i 's probability of correctly identifying that j and k are not linked when the average distance increases by 1, while β^{link} records how knowledge changes with distance when $g_{jk,r} = 1$.⁹ Similarly, γ^{not} and γ^{link} measure the marginal effect of being one standard deviation more central on the probability of being correct when the link does not exist, and that probability when the link does exist. δ^{not} and δ^{link} play the same role but now looking at j and k 's centralities. Finally, $X_{i,jk}$ will include a vector of covariates: average geographic distance of i from j and k , their geographic centralities, and dummies for each of whether i is of the same caste, subcaste, has the same electrification status, has the same roof type, and has the same ownership status for each of j and k , as well as all these variables' interactions with linked status $g_{jk,r}$.

⁸Assuming that each respondent simply offers a response about the existence of a link uniformly at random, a back-of-the-envelope calculation shows that the implied density of the network by our respondents is at least 37.9% higher than that in the data.

⁹Note that even if respondents have biased guesses about whether links exist when they do not truly know (meaning that their guesses may differ on average from the true rate) this does not affect the test of whether β^{not} or β^{link} is non-zero.

The main results can be seen in the raw data in Panels A–C of Figure 1. Panel A depicts the share of correct guesses by average network distance between i and the (j, k) pair. It is immediate to see that the fraction of correct guesses steadily declines in distance down to the unreachable pairs. Panel B illustrates that the share of respondents who do not know jk 's link status increases across network distance. In Panel C, we condition on having a view on the link status and then look at whether the guess is correct or not. We see in the raw data that there is a negative correlation between the share of correct guesses and the network distance, meaning that even beyond having a view on the link, there is residual information which is more accurate for pairs that are on average socially closer to the respondents.

We can also consider how network knowledge varies with the centralities of j and k . In Panels D–F of Figure 1, we plot the same outcomes of interest, varying average centralities. Here, we see that when j and k are more central, i is more likely to offer a correct guess, is less likely to say that she “Doesn’t know” the link status, and has better information about the existence of a link conditional on having a view. Note that more central j and k are mechanically closer to any arbitrary i , so the relationship between centrality and knowledge is likely a direct consequence of the relationships in Panels A–C.¹⁰

3.1.1. *How the likelihood of guessing correctly varies by distance.* Table 2 contains our main network distance results. Panel A looks at how whether i knows jk 's linking status, $y_{i,jk,r}$ depends on the relative network distances between parties and their centralities. Column 1 presents the regression just with the network position variables, column 2 adds the aforementioned vector of demographic controls contained in $X_{i,jk}$, column 3 adds village fixed effects, while column 4 also includes respondent fixed effects. Note that the specification with respondent fixed effects holds the demographic and network characteristics of the respondent fixed, comparing close versus far jk , within respondent.¹¹

We focus on column 1 for exposition. We see that having j and k be on the same connected component as i leads to a 6.5pp increase ($p=0.001$) in the probability of guessing the status correctly when there is no link. Further, this corresponds to a 54.7pp increase ($p=0$) in the probability of guessing the status correctly when there is a link.

Next, we study how respondents' knowledge depends on their network distance to the pair in question. Note that to understand the effect of distance, we need to condition on reachability through the network, as unreachable individuals have infinite distance to the respondent. We find that being one step further than j and k on average leads to a 1.8pp decline ($p=0$) in the probability of guessing correctly when there is no link, and more importantly a 10.9pp decline ($p=0$) in the probability of guessing correctly when there is a link. This is a very large effect relative to the mean (19.9%). Furthermore, columns 2–4

¹⁰This finding may also help explain why respondents can identify the more central individuals fairly accurately (Banerjee et al., 2016).

¹¹Online Appendix B presents the same tables with all network covariates.

illustrate that this pattern is robust to a wide range of specifications, clearly indicating that the quality of respondent’s information is steeply declining in her network distance to the pair and is not explained by similarities in demographics such as wealth, caste, electrification, geography, and so on.¹²

3.1.2. *How the likelihood of being certain varies by distance.* Note that the correct guess outcome in Panel A combines both having a view on a given jk relationship and guessing correctly about the existence (or not) of that relationship. Panel B of Table 2 focuses on the first component. The specifications and control sets are the same as in Panel A, but now our outcome variable is $DK_{i,jk,r}$, a dummy for whether the respondent declared that they “Don’t know” whether j and k are linked.

We find that respondents are considerably more likely to have a view about the network structure local to them and are much more uncertain at larger distances. Once again, we focus on column 1 for exposition, though the results are largely robust to controls and fixed effects. Relative to a mean of 46%, we observe that being reachable leads to a 32.1pp decline ($p=0$) in the probability of declaring “Don’t know” when $g_{jk,r} = 0$ and a 46.6pp decline ($p=0$) when $g_{jk,r} = 1$. These large magnitudes suggests that i has extremely limited knowledge when j and k are unreachable. Then, conditional on reachability, every extra step leads to a 10.9pp (or 14.8pp) ($p=0$) increase in the probability of declaring “Don’t know” when there is no link (when there is a link) which again is a very large effect relative to the mean.

Next, we further decompose the “Don’t know” response into its two components — not knowing that either j or k exist and knowing of j and k , but not having a view on their relationship. We explore this decomposition in Table 3. Columns 1 and 2 consider whether respondent i knows of both j and k as a function of distance and reachability. Note that this variable is defined at the (i, j, k) level and does not vary across relationship types. We therefore collapse the data in columns 1 and 2 to that level. First note that i knows of both j and k only 64% of the time. This means that approximately three-fourths of the “Don’t know” responses come from not knowing anything about one or both of the nodes in question. Unsurprisingly, the patterns of knowledge as a function of reachability and distance look very similar to those in Table 2, Panel B. In columns 3 and 4, we ask whether conditional on knowing that j and k exist, there is a residual relationship between $DK_{i,jk,r}$ and the distance between i and jk . We find that all of the qualitative patterns survive in this conditional regression: reachability is correlated with a lower likelihood of “Don’t know”, while larger social distances make a “Don’t know” response more likely. For example, when

¹²We note that adding controls does reduce the magnitude of the coefficient on reachability when there is a link. The other three coefficients are stable across specifications.

$g_{jk,r} = 1$, moving jk one step further from i is associated with a 4.5pp increase ($p=0$) in $DK_{i,jk,r}$.¹³

3.1.3. Is there residual information conditional on a degree of certainty? Having established that the likelihood that a guess is correct and the likelihood of having a view both decline in distance, we now study whether these results are solely driven by respondents' uncertainty ("Don't know") or whether within a certainty bin, i 's responses are more accurate when jk are socially closer. Table 2, Panel C presents the results of regressions that condition on having a view about the relationship. The control sets across columns are the same as those used in Panels A and B.

We first note that these regressions condition on individuals i , pairs jk , and relationship dimensions r for which respondent i claims to know about the relationship status. Given that we just showed that having a view is also a function of distance, this creates a censoring problem that could lead to bias in the estimated relationship of interest. We believe that the most logical form of bias should make it harder to detect a negative relationship between social distance and bias. This would be the case if i were more likely to have an opinion about far jk pairs who were nonetheless "closer" on unobservables and if those unobservables also corresponded to a higher likelihood of a correct guess.

Keeping that caveat in mind, we do find evidence that distance is correlated with making a correct guess, even conditioning on $DK_{i,jk,r} = 0$. We again find that being 1 step higher in terms of average distance to j and k leads to a decrease of between 1.1 and 2.1pp ($p=0.005$) in the probability of guessing correctly when there is a link between jk . However, there is no detectable pattern in distance when $g_{jk,r} = 0$.

That network distance still predicts accuracy, even conditional on having a view, is consistent with the idea that simply knowing something about j and k does not drive the entire relationship identified in Table 2, Panel A. Instead, a systematic bias develops more strongly outside the local radius.

3.2. What best predicts network knowledge? Our findings thus far illustrate that network knowledge is systematically predicted by and related to the social distance between respondent i and the pair jk in question. We also demonstrated that this correlation is not affected by the inclusion of numerous controls including amenities, caste, geography, village fixed effects, and respondent fixed effects.

We now ask something stronger: out of all the variables we have at our disposal, what best predicts network knowledge? For example, it may be geography, because people observe those who live nearby interacting with others, even if they are not themselves linked. Or it may be the network itself, because people interact on and through the network, so one

¹³We note that the regressions in columns 3 and 4 control on an outcome that we have already shown to be correlated with distance. See Section 3.1.3 below for a discussion of the likely bias.

is more likely to know about relationships among one’s friends or friends of friends rather than someone who is an effective stranger.

To this end, we conduct a post-LASSO procedure (Belloni and Chernozhukov (2013) and Belloni, Chernozhukov, and Hansen (2014b,a)) by regressing the outcome, whether i knows link jk exists or not (or how certain i is), on the network variables and all the demographic variables as before. Since this procedure penalizes putting in many parameters with the aim of selecting a sparse subset of covariates, it may wind up picking a minimal model that excludes network variables altogether. Alternatively, it may include network-based variables, viewing them as more predictive than any of our demographic covariates. Note that since LASSO is a shrinkage estimator, the post-LASSO procedure ensures that consistent estimates are being used.

Table 4 presents the results of this exercise. Column 1 contains the estimates when the outcome is whether the respondent’s guess is correct ($y_{i,jk,r}$). Three core network variables are selected: the distance to the pair when there is no link, centrality of i when jk are linked, and the average centrality of the pair jk when they are linked. As we have mentioned before, the centralities of j and k are highly (negatively) correlated with average distances with an arbitrary i . Additionally (but omitted from the display), the following demographics are selected: dummies for whether i and j , i and k , and j and k are of the same caste, subcaste, have the same occupation, house ownership status, roof type, and electrification status.¹⁴ In column 2 we predict the responses $DK_{i,jk,r}$ of the agent. We find that the distance is selected both when $g_{jk,r} = 0$ and $g_{jk,r} = 1$. Being one step further corresponds to a 8.99pp (or a 19.6%) ($p=0$) increase in the probability of declaring “Don’t know”. Two other network measures are selected here: the centrality of the respondent when $g_{jk,r} = 1$ and the average centralities of jk when $g_{jk,r} = 0$. In this case, a smaller set of demographics is selected: average geographic distance between j and k , and all variables relating to the castes and subcastes of i , j , and k .¹⁵ Note that in either case geographic distance from respondent to respondees or geographic centralities are not selected.

This suggests that an important set of predictive variables in terms of network knowledge or certainty about this knowledge consists of distance in the network from the respondent, how central the respondent is, and how central those in consideration are. These variables are selected by a post-LASSO procedure in a large horseshoe against a number of alternatives, including geographic variables. This is consistent with a story where people learn about their social structure by exploring/interacting through their social structure.

¹⁴23 village fixed effects and 1 dimension fixed effect are also selected.

¹⁵In addition, one occupation variable is selected along with 44 village fixed effects.

4. IMPORTANCE OF ASSUMPTIONS ON NETWORK KNOWLEDGE

In this section, we discuss how our findings regarding limited network knowledge can be relevant for theoretical, applied, and econometric work. Assuming that agents have complete information about the underlying network is an oft-used assumption in the literature (with applications ranging from social learning and peer effects to network formation games). In what follows, we demonstrate how such an assumption (or lack thereof) can have first-order implications for a range of qualitative and quantitative results in various contexts.

4.1. Social learning. As our starting point, we explore the implications of network uncertainty in the context of learning over social networks. Many models of Bayesian social learning assume that agents have complete knowledge about the social network structure.¹⁶ In what follows we use a simple framework to illustrate that relaxing the assumption of full network knowledge may lead to identification problems that can serve as impediments to learning, even when agents have access to enough information to uncover the state. Crucially, we also show that these identification problems do not arise when individuals face no uncertainty about the network.

Consider a collection of $n + 1$ individuals, denoted by $\{0, 1, \dots, n\}$, who wish to estimate an unknown state of the world $\theta \in \mathbb{R}$. Each agent i receives a noisy private signal $s_i = \theta + \epsilon_i$ about the underlying state, where the error terms $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ are drawn independently across agents. Prior to observing their private signals, all agents share an (improper) uniform prior belief about the state.

In addition to her private signals, each agent observes the point estimates of a subset of other individuals, whom we refer to as her neighbors. More specifically, we assume that agents are located on a directed, acyclic network that determines the patterns of observations: agent i can observe agent j 's point estimate ω_j about the state if and only if there is a directed link from agent j to agent i .¹⁷

We represent the potential lack of knowledge about the social network structure by assuming that agent 0 is ex ante uncertain about the patterns of connections among other individuals. More specifically, we assume that the underlying network is drawn randomly from the set $G = \{\mathbf{g}^1, \mathbf{g}^2\}$ with probabilities p^1 and $p^2 = 1 - p^1$, respectively, where

$$\mathbf{g}^i = \{ij : j \neq 1, 2\} \cup \{j0 : j \neq 0\}$$

¹⁶This includes both sequential learning models such as Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), and Smith and Sørensen (2000) as well as models of repeated network interactions such as Mossel, Sly, and Tamuz (2015) and Mossel, Olsman, and Tamuz (2016).

¹⁷A directed network is said to be acyclic if it contains no directed cycles. The assumption that the network is acyclic ensures that information flows are unidirectional. This assumption is also equivalent to imposing an exogenous sequence of timing for agents' observations and assuming that each agent i can only observe the point estimates of a subset of her predecessors, as in Banerjee (1992) and Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), among others.

and jk denotes the directed edge from agent j to agent k . Thus, whereas in both networks agent 0 can observe the estimates of all other individuals, agents labeled $j = 3, \dots, n$ can only observe the estimate of agent $i \in \{1, 2\}$ in network \mathbf{g}^i . The two networks are depicted in Figure 2. We have the following result.

PROPOSITION 1. *Suppose the social network is drawn randomly prior to the realization of the signals. As $n \rightarrow \infty$,*

- (a) *if agent 0 observes the realized network, then she learns the state almost surely;*
- (b) *if agent 0 has misspecified beliefs about the network, she mislearns the state almost surely;*
- (c) *if agent 0 does not observe the realized network, she remains uncertain about the state almost surely.*

Statement (a) of the above result establishes that regardless of what the realized network and state are, agent 0 uncovers the true state with arbitrarily high confidence as $n \rightarrow \infty$ as long as she has full knowledge about the underlying network. Note that even though 0's neighbors (partially) rely on a common source of information, agent 0 can use her knowledge about the network structure to account for any redundancies in her neighbors' estimates. More specifically, when the realized network is \mathbf{g}^i , the point estimate of each agent $j \in \{3, \dots, n\}$ is equal to $\omega_j = (s_j + s_i)/2$, while $\omega_1 = s_1$ and $\omega_2 = s_2$. Hence, with access to the point estimates of agents 1 and 2, agent 0 can simply back out agent j 's private signal s_j by computing $2\omega_j - \omega_i$. Such a calculation enables agent 0 to discern any correlation in her neighbors' estimates that is due to their common neighbor i . Consequently, agent 0's point estimate of the true state is given by

$$\omega_0 = \frac{1}{n+1} \left(s_0 + \omega_1 + \omega_2 + \sum_{j=3}^n (2\omega_j - \omega_i) \right) = \frac{1}{n+1} \sum_{j=0}^n s_j,$$

which by the law of large numbers, converges to the true state with probability one as $n \rightarrow \infty$.

Contrasting the above observation with statements (b) and (c) of Proposition 1 underscores the crucial role of *correct* and *complete* network knowledge for successful learning in part (a): whereas agent 0 learns the state with probability one when she observes the realized network, she fails to learn the state if she has either incorrect or incomplete knowledge about the realized network.

To see the intuition for this contrast, first consider the case in which agent 0 has misspecified beliefs about the social network (statement (b)). More specifically, suppose without loss of generality that the true realized network is \mathbf{g}^1 , but agent 0 believes (mistakenly) that the underlying network is \mathbf{g}^2 . Such a misspecification means that even though the point estimate of agent $j \in \{3, \dots, n\}$ is given by $\omega_j = (s_j + s_1)/2$, agent 0 assumes that the correlation in her neighbors' estimates are due to agent 2's private signal. As a result,

agent 0’s point estimate is given by

$$\omega_0 = \frac{1}{n+1} \left(s_0 + \omega_1 + \omega_2 + \sum_{j=3}^n (2\omega_j - \omega_2) \right) = \frac{1}{n+1} \left((n-2)(s_1 - s_2) + \sum_{j=0}^n s_j \right),$$

which converges to $\theta + s_1 - s_2$ with probability one, whenever the true state is θ . In other words, agent 0’s misspecified belief about the network structure manifests itself as a bias in her estimate about the state.

Finally, part (c) of Proposition 1 illustrates that uncertainty about the network structure can create an identification problem that would serve as an impediment to learning. In the context of the above example, no matter how many neighbors agent 0 has, she can only learn the state by properly accounting for the redundancies and patterns of correlations in her neighbors’ estimates. Yet, uncertainty about the identity of her neighbors’ neighbors means that agent 0 cannot identify the source of such redundancies. More specifically, when the true state and the realized network are θ and \mathbf{g}^1 , respectively, agent 0 has no way of distinguishing whether the underlying state-network pair is (θ, \mathbf{g}^1) or $(\theta + s_1 - s_2, \mathbf{g}^2)$, even as $n \rightarrow \infty$. This is despite the fact that she can hold and update beliefs not only about the state but also about the underlying network structure.¹⁸

Taken together, Proposition 1 illustrates that incorrect or incomplete knowledge about the social network can undermine information aggregation, even in fairly simple environments that would have otherwise led to learning. This observation, alongside our empirical findings in Section 3, suggests that Bayesian models that allow for uncertainty about the network structure (such as Lobel and Sadler (2015)) as well as behavioral models (such as DeGroot (1974), Eyster and Rabin (2014), and Li and Tan (2017)) that assume agents cannot fully account for how the network structure shapes their observations may be useful and important starting points for applied work to focus on.

4.2. Network interactions and peer effects. We next focus on network peer effect games and illustrate how assuming full network knowledge may lead to biased structural estimates when in fact agents face uncertainty about the underlying network.

To this end, we focus on the canonical network interaction model of Ballester, Calvó-Armengol, and Zenou (2006). Besides their widespread use, models that build off this linear-quadratic framework lend themselves to readily obtainable structural estimates, as well as counterfactual and policy analyses (such as which “key” players to target with an intervention). Applications include peer effects in education (Calvó-Armengol, Patacchini, and Zenou, 2009), local consumption externalities (Candogan, Bimpikis, and Ozdaglar,

¹⁸In this sense, our result on the role of network uncertainty is similar to Acemoglu, Chernozhukov, and Yildiz (2016), who argue that uncertainty about the signal generating process can result in an identification problem for a single Bayesian agent in isolation, as the same long-run frequency of signals may be consistent with multiple states.

2012), research collaboration among firms (Goyal and Moraga-González, 2001), among numerous others.

The basic setup is an n -player game in which agents' payoffs are given by

$$u_i(x_1, \dots, x_n) = x_i \theta_i - \frac{1}{2} x_i^2 + \alpha \sum_{j \neq i} g_{ij} x_i x_j,$$

where x_i denotes the action (say, effort) of agent i , g_{ij} regulates the strength of strategic complementarities between i and j , and parameters $\alpha \in [0, 1)$ and $\theta_i > 0$ are common knowledge among all agents.

As is common in the literature, pairwise interactions can be represented by a directed and weighted network $\mathbf{g} = [g_{ij}]$, with the convention that $g_{ii} = 0$ for all i . In a departure from the rest of the literature, however, we assume that this interaction network is drawn randomly prior to the start of the game. More specifically, we assume that for any pair of agents $i \neq j$, the weight g_{ij} is drawn independently from the rest of the weights according to a non-degenerate probability distribution with density f_{ij} . Throughout, we assume that $\sum_{j \neq i} \max \text{supp}(f_{ij}) \leq 1$ for all agents i .¹⁹

We consider two variants of the above game. In what we call the complete information benchmark, we assume that the realized network \mathbf{g} is observed by all agents prior to making their decisions. Therefore, ex post, our complete information benchmark coincides with the canonical model of Ballester, Calvo-Armengol, and Zenou (2006). On the other hand, under the incomplete information variant of the game, we assume that each agent i can only observe the set of weights $\{g_{ij} : j \neq i\}$ in her neighborhood. Hence, in contrast to the complete information benchmark, agents choose their actions with no knowledge about the strength of strategic complementarities among other agents. We have the following result:

PROPOSITION 2. *The vectors of equilibrium actions under the complete and incomplete information variants of the game are given by*

$$(2) \quad x^{\text{com}} = (I - \alpha \mathbf{g})^{-1} \theta$$

$$(3) \quad x^{\text{inc}} = \left(I + \alpha \mathbf{g} (I - \alpha \mathbb{E}[\mathbf{g}])^{-1} \right) \theta,$$

respectively. Furthermore, $\mathbb{E}[x_i^{\text{inc}}] < \mathbb{E}[x_i^{\text{com}}]$ for all agents i .

The above result establishes that equilibrium actions are sensitive to whether agents have full knowledge about the patterns of interactions in the network. More importantly however, Proposition 2 also illustrates that uncertainty about the network structure results in a systematic shift in equilibrium actions: each agent's action in the incomplete information game is in expectation less than her action in the complete information benchmark.

What this means, practically, is that estimates of θ and α obtained under the assumption of complete information may be biased if agents are in fact uncertain about the network

¹⁹This restriction on the support of the distributions is meant to guarantee that the matrix $I - \alpha \mathbf{g}$ is invertible with probability one.

structure. Such biases can in turn substantially affect not only counterfactuals but also policy prescriptions by potentially impacting (i) cost-benefit assessments, (ii) the predicted extent of externalities, and (iii) the relative network centralities of various agents; a statistic that is central to policy prescriptions in this literature.

4.3. Identification in network formation games. We conclude this section by investigating the implications of incomplete network knowledge for identification of structural parameters in network formation games. As our main result, we show that introducing the (arguably more realistic) assumption of incomplete information may transform a model that is fundamentally hard to identify to a straightforwardly estimable model.

Consider a simultaneous-move game of network formation with n players, indexed $\{1, \dots, n\}$. Each agent i is endowed with some attribute $z_i \in \{a, b\}$. These attributes, which are common knowledge among all agents and are observable to the econometrician, may represent agents' gender, caste, race, or other characteristics. Throughout, we use λ to denote the fraction of agents with attribute a .

Agents draw utilities by forming direct and indirect links with other individuals. More specifically, the utility function of agent i is given by

$$(4) \quad u_i(\mathbf{g}, z, \epsilon_i) = \sum_{j \neq i} g_{ij}(\alpha - \epsilon_{ij}) + \beta \sum_{j \neq i} g_{ij} \mathbf{1}\{z_j = z_i\} + \gamma \left(\frac{1}{n-1} \sum_{j \neq i} \sum_{k \neq i, j} g_{ij} g_{jk} \right),$$

where \mathbf{g} denotes the undirected, unweighted network of linkages (with $g_{ij} = 1$ if i and j are linked and $g_{ij} = 0$ otherwise), $z = (z_1, \dots, z_n)$ denotes the vector of population attributes, and ϵ_{ij} is a preference shock that determines agent i 's idiosyncratic cost of establishing a link to agent j . Throughout, we assume that preference shocks are drawn independently from a common probability distribution $F(\cdot)$ with a continuous density and full support over \mathbb{R} .

The first two terms on the right-hand side of (4) capture the (net) *direct* utility of forming linkages: agent i obtains a net marginal utility of $\alpha - \epsilon_{ij}$ of forming a direct link to agent j , with an additional marginal benefit of β if the two agents have identical attributes. On the other hand, the third term on the right-hand side of (4) allows for the possibility that *indirect* connections can also be payoff relevant. For example, when $\gamma > 0$, agent i draws a benefit if any of her neighbors form links with other individuals.²⁰

Players simultaneously announce the set of agents they wish to link to, with links formed by mutual consent. Formally, each player i chooses the action $x_{ij} \in \{0, 1\}$ indicating

²⁰As in Sheng (2016) and Mele (2017), the specification in (4) assumes that the indirect benefits of forming links are proportional to the number of connections of i 's neighbors. An alternative specification is to assume that indirect benefits only depend on the number of *distinct* connections of i 's neighbors, as in de Paula, Richards-Shubik, and Tamer (2018). Our choice of the specification in (4) is made to simplify the analysis. However, our main result on the relationship between network knowledge and the identifiability of structural parameters holds for a broader class of specifications.

whether she wishes to form a link with agent j . The undirected link (i, j) is then formed if both parties agree, i.e., $g_{ij} = x_{ij}x_{ji}$.

The econometrician is interested in estimating the structural parameters α , β , and γ in (4) while only observing the equilibrium network \mathbf{g} . The literature typically assumes that the realizations of preference shocks ϵ_{ij} are common knowledge among the agents and that agents are in a pairwise stable network (de Paula, Richards-Shubik, and Tamer, 2018; Leung, 2015a; Menzel, 2017; Sheng, 2016). This observability assumption means that agents' link formation decisions $x_{ij} \in \{0, 1\}$ can be contingent on the realization of all pairwise preference shocks.

The fact that network formation games of complete information tend to pose a challenge for identification has been at the forefront of the literature. To overcome this issue, the literature has resorted to imposing subsequent stronger assumptions on agents' payoffs, the number of preference shocks, or the network formation process (see, e.g., de Paula et al. (2018); Leung (2015a); Menzel (2017); Sheng (2016)).

In a departure from the literature, we assume that preference shocks are only locally observable: the realization of agent i 's preference shocks $\epsilon_i = (\epsilon_{i1}, \dots, \epsilon_{in})$ are i 's private information and hence are unobservable to all other agents. Given the incomplete information nature of this specification, we focus on (symmetric monotone) Bayes-Nash equilibrium as our solution concept, according to which each agent i chooses the action $x_{ij} = 1$ if the realized preference shock ϵ_{ij} falls below some given threshold that only depends on i and j 's attributes.

Before presenting our result, we remark that the notion of incomplete information in this context is slightly different from those in Subsections 4.1 and 4.2. Whereas our earlier examples focused on the observability (or lack thereof) of the network structure itself, the incomplete information variant of the network formation game above requires the agents to be uncertain about the realizations of the underlying pairwise preference shocks.

PROPOSITION 3. *Suppose the econometrician observes the realized network but not the realization of the preference shocks and $\lambda \neq 1/2$. If preference shocks are only locally observable, structural parameters α , β , and γ are point-identified.*

The proposition illustrates that relaxing the complete information assumption leads to a diametrically opposite result of that in the literature: observing the pattern of linkages provides the analyst with sufficiently rich information to point-identify all structural parameters. This contrast is driven by the fact that when players have limited information, they are unable to condition their decisions on the fine details of the entire set of realized shocks $\{\epsilon_{ij}\}$, leading to coarser linking strategies and hence facilitating identification. This illustrates that imposing the simple and — in view of our empirical findings in Section 3 — more realistic assumption of incomplete information transforms a model that is fundamentally hard to identify to a realistic and straightforwardly estimable model, without changing

the payoffs, imposing specific link formation dynamics, or restricting the number or types of links that can be formed.

5. CONCLUSION

We present a new set of stylized facts about the extent to which individuals are informed about network relationships in their own communities. We collected survey data in 75 communities where we had already documented the full network structure across a number of dimensions. We find that network knowledge is low and localized: respondents often do not know linking statuses among others in their village and that network knowledge decays steeply in network distance to the pair. These patterns are robust to the inclusion of numerous covariates (e.g., amenities, caste, geography) and are consistent with individuals learning about the social network through the network itself.

Our findings are at odds with many commonly-used modeling techniques in network economics across a range of questions and applications. Moving from a complete information framework to an incomplete information framework can lead to changes in the behavior of agents in the models and, thus different predictions or equilibria, thereby affecting qualitative conclusions and structural estimates. Further, we show that typical econometric models of network formation have the features that moving to the more realistic assumption of incomplete information may make parameters identifiable whereas under the less realistic assumption of complete information, parameters may not be identifiable or only partially identifiable. We are certainly not the first to explore incomplete information in any of these cases. However, we view our examples playing a pedagogical role, pointing out how this ubiquitous assumption in the literature, inconsistent with the data, may affect widespread conclusions.

Taken together, our empirical and theoretical results suggest that the complete information benchmark is not an accurate view of the world for the setting we study, and moreover, is not always a benign assumption. Rather, incorporating incomplete network information may yield theoretical predictions that better match the behaviors of agents in real-world networks. Incomplete information may also play a simplifying role in the theory, both yielding a smaller set of equilibria in games of cooperation and also leading to smaller identified sets in network econometric models.

REFERENCES

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2016): “Fragility of asymptotic agreement under Bayesian learning,” *Theoretical Economics*, 11, 187–225. 18
- ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): “Bayesian learning in social networks,” *The Review of Economic Studies*, 78, 1201–1236. 1, 17
- ALATAS, V., A. BANERJEE, A. G. CHANDRASEKHAR, R. HANNA, AND B. A. OLKEN (2016): “Network structure and the aggregation of information: Theory and evidence from Indonesia,” *American Economic Review*, 106, 1663–1704. 1
- BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): “Who’s who in networks. Wanted: The key player,” *Econometrica*, 74, 1403–1417. 1, 4.2
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): “The diffusion of microfinance,” *Science*, 341, 1236498. 2.1, 5
- (2016): “Gossip: Identifying Central Individuals in a Social Network,” NBER Working Paper No. 20422. 1, 2.1, 10
- BANERJEE, A. V. (1992): “A simple model of herd behavior,” *The Quarterly Journal of Economics*, 107, 797–817. 1, 16, 17
- BELLONI, A. AND V. CHERNOZHUKOV (2013): “Least squares after model selection in high-dimensional sparse models,” *Bernoulli*, 19, 521–547. 3.2
- BELLONI, A., V. CHERNOZHUKOV, AND C. HANSEN (2014a): “High-dimensional methods and inference on structural and treatment effects,” *The Journal of Economic Perspectives*, 29–50. 3.2
- (2014b): “Inference on treatment effects after selection among high-dimensional controls,” *The Review of Economic Studies*, 81, 608–650. 3.2
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A theory of fads, fashion, custom, and cultural change as information cascades,” *Journal of Political Economy*, 100, 992–1026. 16
- CALVÓ-ARMENGOL, A., E. PATACCINI, AND Y. ZENOU (2009): “Peer effects and social networks in education,” *Review of Economic Studies*, 76, 1239–1267. 1, 4.2
- CANDOGAN, O., K. BIMPIKIS, AND A. OZDAGLAR (2012): “Optimal pricing in networks with externalities,” *Operations Research*, 60, 883–905. 1, 4.2
- DE PAULA, A., S. RICHARDS-SHUBIK, AND E. TAMER (2018): “Identifying preferences in networks with bounded degree,” *Econometrica*, 86, 263–288. 1, 20, 4.3
- DEGROOT, M. H. (1974): “Reaching a consensus,” *Journal of the American Statistical Association*, 69, 118–121. 4.1
- EYSTER, E. AND M. RABIN (2014): “Extensive imitation is irrational and harmful,” *The Quarterly Journal of Economics*, 129, 1861–1898. 4.1
- FRIEDKIN, N. E. (1983): “Horizons of observability and limits of informal control in organizations,” *Social Forces*, 62, 54–77. 1

- GOYAL, S. AND J. L. MORAGA-GONZÁLEZ (2001): “R&D networks,” *The RAND Journal of Economics*, 32, 686–707. 4.2
- KRACKHARDT, D. (1987): “Cognitive social structures,” *Social networks*, 9, 109–134. 1, 2
- (2014): “A preliminary look at accuracy in egonets,” in *Contemporary Perspectives on Organizational Social Networks*, Emerald Group Publishing Limited, 277–293. 1, 2
- LEUNG, M. P. (2015a): “A random-field approach to inference in large models of network formation,” *Available at SSRN*. 1, 4.3
- (2015b): “Two-step estimation of network-formation models with incomplete information,” *Journal of Econometrics*, 188, 182–195. 1
- LI, W. AND X. TAN (2017): “Learning in local networks,” Working paper. 4.1
- LOBEL, I. AND E. SADLER (2015): “Information diffusion in networks through social learning,” *Theoretical Economics*, 10, 807–851. 1, 4.1
- MELE, A. (2017): “A structural model of dense network formation,” *Econometrica*, 85, 825–850. 20
- MENZEL, K. (2017): “Strategic network formation with many agents,” Working paper. 1, 4.3
- MOSSEL, E., N. OLSMAN, AND O. TAMUZ (2016): “Efficient Bayesian learning in social networks with Gaussian estimators,” in *Annual Allerton Conference on Communication, Control, and Computing*, Allerton House, Illinois, 425–432. 16
- MOSSEL, E., A. SLY, AND O. TAMUZ (2015): “Strategic learning and the topology of social networks,” *Econometrica*, 83, 1755–1794. 1, 16
- RIDDER, G. AND S. SHENG (2017): “Estimation of large network formation games,” Working paper. 1
- SHENG, S. (2016): “A structural econometric analysis of network formation games,” Working paper. 1, 20, 4.3
- SMITH, L. AND P. N. SØRENSEN (2000): “Pathological outcomes of observational learning,” *Econometrica*, 68, 371–398. 1, 16

FIGURES

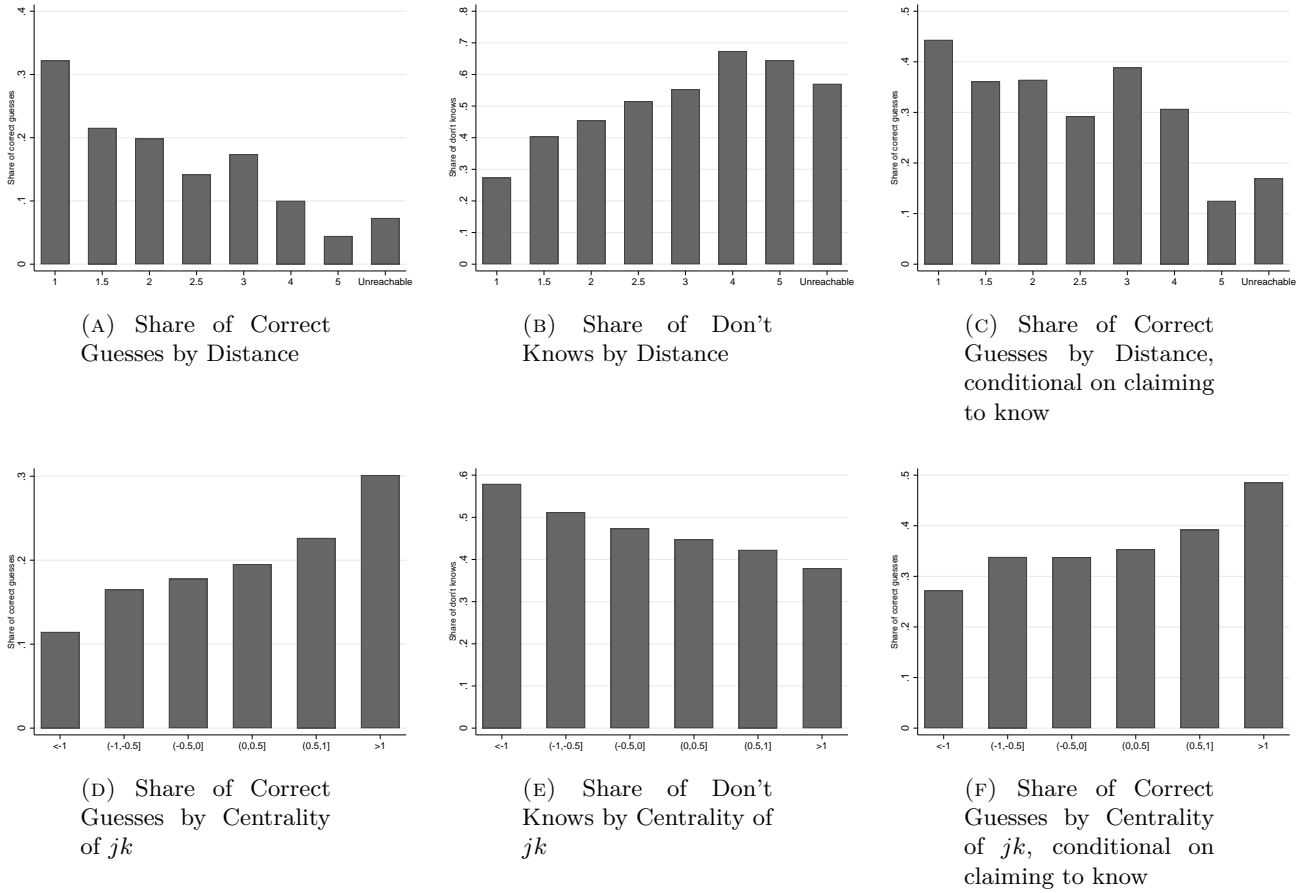


FIGURE 1. Panels A–C depict raw shares of correct guesses and declaring that the respondent doesn't know by distance. Panels D–F show the same by average centrality of j and k (standardized).




FIGURE 2. Networks g^1 (left) and g^2 (right). Arrows indicate the direction over which information can flow from one agent to another.

TABLES

TABLE 1. Sample Statistics

<i>Panel A: Respondent Characteristics</i>		
	mean	obs
Female	0.59	4838
Age	42.55	4838
General or OBC	0.65	4838
Born in the village	0.27	4838
Born in the village (female)	0.11	2854
Born in the village (male)	0.51	1984
Degree	19.51	4838
Eigenvector centrality	0.07	4838
<i>Panel B: Pair Characteristics</i>		
	mean	obs
j,k Reachable	0.99	21861
j,k Linked	0.34	21861
Avg. distance to j,k	2.13	21861
Avg. degree centrality of j,k	18.68	21861
Avg. eigenvector centrality of j,k	0.06	21861
Same caste category j,k	0.59	21834
Same subcaste j,k	0.44	21818
<i>Panel C: Knowledge</i>		
	mean	obs
Correct guess	0.20	66810
Correct guess (j,k not linked)	0.07	43884
Correct guess (j,k linked)	0.57	22926
Don't know	0.46	66810

TABLE 2. Network Knowledge as a function of Distance and Centralities

<i>Panel A: Correct guesses (including don't knows as wrong guesses)</i>				
VARIABLES	(1)	(2)	(3)	(4)
	Correct guess	Correct guess	Correct guess	Correct guess
1{Reachable j,k, Not linked j,k}	0.0647 (0.0196) [0.001]	0.0639 (0.0192) [0.001]	0.0549 (0.0194) [0.006]	0.0420 (0.0224) [0.065]
1{Reachable j,k, Linked j,k}	0.547 (0.0227) [0.000]	0.147 (0.112) [0.193]	0.122 (0.108) [0.261]	0.174 (0.142) [0.224]
Avg. distance to j,k, Not linked j,k	-0.0177 (0.00398) [0.000]	-0.0153 (0.00389) [0.000]	-0.0115 (0.00400) [0.005]	-0.0133 (0.00437) [0.003]
Avg. distance to j,k, Linked j,k	-0.109 (0.00729) [0.000]	-0.0943 (0.00656) [0.000]	-0.0920 (0.00607) [0.000]	-0.0925 (0.00673) [0.000]
Observations	66,810	66,810	66,810	66,810
Depvar mean	0.199	0.199	0.199	0.199
<i>Panel B: Don't knows</i>				
VARIABLES	(1)	(2)	(3)	(4)
	Don't know	Don't know	Don't know	Don't know
1{Reachable j,k, Not linked j,k}	-0.321 (0.0377) [0.000]	-0.271 (0.0362) [0.000]	-0.246 (0.0288) [0.000]	-0.253 (0.0335) [0.000]
1{Reachable j,k, Linked j,k}	-0.466 (0.0381) [0.000]	-0.355 (0.153) [0.023]	-0.244 (0.163) [0.139]	-0.228 (0.216) [0.293]
Avg. distance to j,k, Not linked j,k	0.109 (0.00866) [0.000]	0.0933 (0.00939) [0.000]	0.0854 (0.00653) [0.000]	0.0811 (0.00686) [0.000]
Avg. distance to j,k, Linked j,k	0.148 (0.00873) [0.000]	0.123 (0.00788) [0.000]	0.118 (0.00784) [0.000]	0.122 (0.00887) [0.000]
Observations	66,810	66,810	66,810	66,810
Depvar mean	0.463	0.463	0.463	0.463
<i>Panel C: Correct guesses (conditional on sure guesses only)</i>				
VARIABLES	(1)	(2)	(3)	(4)
	Correct guess	Correct guess	Correct guess	Correct guess
1{Reachable j,k, Not linked j,k}	0.00979 (0.0432) [0.821]	0.0348 (0.0424) [0.414]	0.0327 (0.0368) [0.378]	-0.00996 (0.0484) [0.837]
1{Reachable j,k, Linked j,k}	0.534 (0.0423) [0.000]	-0.158 (0.108) [0.148]	-0.101 (0.127) [0.427]	0.108 (0.359) [0.764]
Avg. distance to j,k, Not linked j,k	0.00668 (0.00699) [0.342]	0.00541 (0.00770) [0.484]	0.00857 (0.00684) [0.214]	0.000630 (0.00720) [0.930]
Avg. distance to j,k, Linked j,k	-0.0189 (0.00697) [0.008]	-0.0208 (0.00711) [0.005]	-0.0177 (0.00727) [0.017]	-0.0110 (0.00924) [0.239]
Observations	35,909	35,909	35,909	35,909
Depvar mean	0.371	0.371	0.371	0.371

Notes: Standard errors (clustered at village level) are reported in parentheses. p -values are reported in brackets. All columns control for eigenvector centrality of i and avg. eigenvector centrality of j,k , when g_{jk} is 1 and 0. All columns also include dimension fixed-effects. Columns (2), (3) and (4) control for demographic covariates. Demographic controls include average geographic distance of i from j and k and dummies for each of whether i is of the same caste, subcaste, has the same electrification status, has the same roof type, has the same number of rooms, and has the same house ownership status for each of j and k , when g_{jk} is 1 and 0. They also include the above-mentioned dummies for when j and k have the same demographic traits. Columns (3) and (4) include village and respondent fixed-effects, respectively.

TABLE 3. Network Knowledge, conditional on knowing j,k , as a function of Distance and Centralities

VARIABLES	(1) Know pair	(2) Know pair	(3) Don't know	(4) Don't know
1{Reachable j,k , Not linked j,k }	0.337 (0.0393) [0.000]	0.289 (0.0389) [0.000]	-0.0671 (0.0355) [0.063]	-0.0490 (0.0337) [0.150]
1{Reachable j,k , Linked j,k }	0.446 (0.0395) [0.000]	0.349 (0.174) [0.049]	-0.122 (0.0338) [0.001]	-0.0548 (0.0672) [0.418]
Avg. distance to j,k , Not linked j,k	-0.109 (0.00923) [0.000]	-0.0959 (0.00975) [0.000]	0.0323 (0.00736) [0.000]	0.0229 (0.00728) [0.002]
Avg. distance to j,k , Linked j,k	-0.137 (0.00948) [0.000]	-0.116 (0.00834) [0.000]	0.0447 (0.00641) [0.000]	0.0349 (0.00738) [0.000]
Observations	22,270	22,270	42,750	42,750
Know pair only			✓	✓
Depvar mean	0.640	0.640	0.160	0.160
Number of (i, jk) s			14250	14250

Notes: Standard errors (clustered at village level) are reported in parentheses. p -values are reported in brackets. All columns control for eigenvector centrality of i and avg. eigenvector centrality of j,k , when g_{jk} is 1 and 0. Columns (3) and (4) include dimension fixed-effects. Columns (2) and (4) control for demographic covariates. Demographic controls include average geographic distance of i from j and k and dummies for each of whether i is of the same caste, subcaste, has the same electrification status, has the same roof type, has the same number of rooms, and has the same house ownership status for each of j and k , when g_{jk} is 1 and 0. They also include the above-mentioned dummies for when j and k have the same demographic traits.

TABLE 4. Post-LASSO Estimates

VARIABLES	(1) Correct guess	(2) Don't know
Avg. distance to j,k, Not linked j,k	-0.0269 (0.00367) [0.000]	0.105 (0.00603) [0.000]
Avg. distance to j,k, Linked j,k		0.0899 (0.00554) [0.000]
Centrality (std.) of i, Linked j,k	0.0134 (0.00582) [0.024]	-0.0125 (0.00560) [0.029]
Avg. centrality (std.) of j,k, Not linked j,k		-0.00211 (0.00692) [0.761]
Avg. centrality (std.) of j,k, Linked j,k	0.0121 (0.00520) [0.022]	
Observations	57,597	63,558
Reachable pairs only	✓	✓
LASSO-selected Demographic controls	✓	✓
LASSO-selected Village FE	✓	✓
Depvar mean	0.204	0.459

Notes: Standard errors (clustered at village level) are reported in parentheses. p -values are reported in brackets. Demographic controls include average geographic distance of i from j and k and dummies for each of whether i is of the same caste, subcaste, has the same electrification status, has the same roof type, has the same number of rooms, and has the same house ownership status for each of j and k , when g_{jk} is 1 and 0. They also include the above-mentioned dummies for when j and k have the same demographic traits. Post-LASSO procedure used with optimal penalty parameter to select variables.

APPENDIX A. PROOFS

Proof of Proposition 1.

Proof of part (a). Throughout, let $\omega^n = (\omega_1, \dots, \omega_n)$, where ω_i denotes the point estimates of agent i . Without loss of generality, suppose the true underlying state and the realized network are θ and \mathbf{g}^2 , respectively. By assumption, the realized network is observable to agent 0. Consequently, agent 0's posterior likelihood ratio for an arbitrary pair of states $\hat{\theta}$ and $\tilde{\theta}$ given her observations is equal to

$$(5) \quad \frac{f(\hat{\theta}|s_0, \omega^n, \mathbf{g}^2)}{f(\tilde{\theta}|s_0, \omega^n, \mathbf{g}^2)} = \frac{f(s_0, \omega^n|\hat{\theta}, \mathbf{g}^2)}{f(s_0, \omega^n|\tilde{\theta}, \mathbf{g}^2)} = \frac{f(s_0|\hat{\theta})f(\omega^n|\hat{\theta}, \mathbf{g}^2)}{f(s_0|\tilde{\theta})f(\omega^n|\tilde{\theta}, \mathbf{g}^2)},$$

where we are using the assumption that agents have a common uniform prior beliefs over the real line and the fact that signal s_0 is conditionally independent from the point estimates of all other agents. On the other hand, since agents $j \in \{3, \dots, n\}$ can observe agent 2's point estimate $\omega_2 = s_2$, it is immediate that $\omega_j = (s_j + s_2)/2$. Therefore, for any arbitrary state $\tilde{\theta} \in \mathbb{R}$,

$$f(\omega^n|\tilde{\theta}, \mathbf{g}^2) = \frac{1}{(\sqrt{2\pi}\sigma)^n} \exp\left(-\frac{1}{2\sigma^2} \left((\tilde{\theta} - \omega_1)^2 + (\tilde{\theta} - \omega_2)^2 + \sum_{i=3}^n (\tilde{\theta} - 2\omega_i + \omega_2)^2 \right)\right).$$

The juxtaposition of the above equation with equation (5) therefore implies that agent 0's posterior log-likelihood ratio for any given pair of states $\hat{\theta}$ and $\tilde{\theta}$ is equal to

$$(6) \quad \log \frac{f(\hat{\theta}|s_0, \omega^n, \mathbf{g}^2)}{f(\tilde{\theta}|s_0, \omega^n, \mathbf{g}^2)} = \frac{1}{2\sigma^2} (\tilde{\theta} - \hat{\theta}) \left((n+1)(\hat{\theta} + \tilde{\theta}) - 2s_0 - 2\omega_1 - 2\omega_2 - 2 \sum_{i=3}^n (2\omega_i - \omega_2) \right).$$

Set $\hat{\theta} = \theta$, where θ is the underlying state of the world. Multiplying both sides of the above equation by $1/n$ and taking the limit as $n \rightarrow \infty$ implies that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \frac{f(\theta|s_0, \omega^n, \mathbf{g}^2)}{f(\tilde{\theta}|s_0, \omega^n, \mathbf{g}^2)} = \frac{1}{2\sigma^2} (\tilde{\theta} - \theta)^2,$$

where we are using the fact that by the law of large numbers, $\lim_{n \rightarrow \infty} (1/n) \sum_{i=3}^n \omega_i = (\omega_2 + \theta)/2$ with probability 1. Therefore, for any $\tilde{\theta} \neq \theta$,

$$\lim_{n \rightarrow \infty} \frac{f(\theta|s_0, \omega^n, \mathbf{g}^2)}{f(\tilde{\theta}|s_0, \omega^n, \mathbf{g}^2)} = \infty$$

with probability one. Consequently, agent 0 assigns an asymptotic belief of one on the true state θ as $n \rightarrow \infty$ almost surely. \square

Proof of part (b). Without loss of generality let θ and \mathbf{g}^1 denote the realized state and network, respectively. However, suppose that agent 0 has misspecified beliefs about the network by mistakenly believing that the underlying network is \mathbf{g}^2 . Following an argument

identical to the proof of part (a), the log-likelihood ratio of agent 0's posterior beliefs is given by (6). Consequently, the log-likelihood ratio of agent 0's posterior beliefs as $n \rightarrow \infty$ satisfies

$$(7) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \log \frac{f(\hat{\theta}|s_0, \omega^n, \mathbf{g}^2)}{f(\tilde{\theta}|s_0, \omega^n, \mathbf{g}^2)} = \frac{1}{2\sigma^2} (\tilde{\theta} - \hat{\theta})(\hat{\theta} + \tilde{\theta} - 2(\theta + \omega_1 - \omega_2))$$

almost surely. In the above expression, we are using the fact that the true underlying network is \mathbf{g}^1 , and hence, $\lim_{n \rightarrow \infty} (1/n) \sum_{i=3}^n \omega_i = (\omega_1 + \theta)/2$ with probability one. Setting $\hat{\theta} = \theta + \omega_1 - \omega_2$ therefore implies that the right-hand side of (7) is strictly positive for all $\tilde{\theta} \neq \theta + \omega_1 - \omega_2$. Consequently,

$$(8) \quad \lim_{n \rightarrow \infty} \frac{f(\hat{\theta}|s_0, \omega^n, \mathbf{g}^2)}{f(\tilde{\theta}|s_0, \omega^n, \mathbf{g}^2)} = \infty$$

almost surely for all $\tilde{\theta} \neq \theta + \omega_1 - \omega_2$. In other words, as $n \rightarrow \infty$, agent 0 becomes certain that the realized state is $\theta + \omega_1 - \omega_2$; a state that is distinct from the true underlying state θ with probability one. \square

Proof of part (c). Without loss of generality let θ and \mathbf{g}^1 denote the realized state and network, respectively, and let $\hat{\theta} \neq \theta$. Agent 0's posterior likelihood ratios are given by

$$\frac{f(\theta|s_0, \omega^n)}{f(\hat{\theta}|s_0, \omega^n)} = \frac{\mathbb{P}(\mathbf{g}^1)f(s_0, \omega^n|\theta, \mathbf{g}^1) + \mathbb{P}(\mathbf{g}^2)f(s_0, \omega^n|\theta, \mathbf{g}^2)}{\mathbb{P}(\mathbf{g}^1)f(s_0, \omega^n|\hat{\theta}, \mathbf{g}^1) + \mathbb{P}(\mathbf{g}^2)f(s_0, \omega^n|\hat{\theta}, \mathbf{g}^2)},$$

where we are using Bayes' rule and the assumption that agents have improper uniform prior beliefs. Consequently,

$$\frac{f(\theta|s_0, \omega^n)}{f(\hat{\theta}|s_0, \omega^n)} \leq \frac{\mathbb{P}(\mathbf{g}^1)f(s_0, \omega^n|\theta, \mathbf{g}^1)}{\mathbb{P}(\mathbf{g}^2)f(s_0, \omega^n|\hat{\theta}, \mathbf{g}^2)} + \frac{f(s_0, \omega^n|\theta, \mathbf{g}^2)}{f(s_0, \omega^n|\hat{\theta}, \mathbf{g}^2)} = \frac{\mathbb{P}(\mathbf{g}^1)f(s_0|\theta)f(\omega^n|\theta, \mathbf{g}^1)}{\mathbb{P}(\mathbf{g}^2)f(s_0|\hat{\theta})f(\omega^n|\hat{\theta}, \mathbf{g}^2)} + \frac{f(\theta|s_0, \omega^n, \mathbf{g}^2)}{f(\hat{\theta}|s_0, \omega^n, \mathbf{g}^2)}.$$

Let $\hat{\theta} = \theta + \omega_1 - \omega_2$. Equation (8) in the proof of part (b) implies that the last term on the right-hand side of the above equation converges to zero as $n \rightarrow \infty$ almost surely whenever the true state-network pair is (θ, \mathbf{g}^1) . Therefore,

$$(9) \quad \limsup_{n \rightarrow \infty} \frac{f(\theta|s_0, \omega^n)}{f(\hat{\theta}|s_0, \omega^n)} \leq \frac{\mathbb{P}(\mathbf{g}^1)}{\mathbb{P}(\mathbf{g}^2)} \frac{f(s_0|\theta)}{f(s_0|\hat{\theta})} \limsup_{n \rightarrow \infty} \frac{f(\omega^n|\theta, \mathbf{g}^1)}{f(\omega^n|\hat{\theta}, \mathbf{g}^2)}.$$

On the other hand, note that

$$\begin{aligned} \log \frac{f(\omega^n|\theta, \mathbf{g}^1)}{f(\omega^n|\hat{\theta}, \mathbf{g}^2)} &= \frac{1}{2\sigma^2} \left((\hat{\theta} - \omega_1)^2 - (\theta - \omega_2)^2 \right) + \frac{1}{2\sigma^2} \left((\hat{\theta} - \omega_2)^2 - (\theta - \omega_1)^2 \right) \\ &\quad + \frac{1}{2\sigma^2} \sum_{i=3}^n \left((\hat{\theta} - 2\omega_i + \omega_2)^2 - (\theta - 2\omega_i + \omega_1)^2 \right). \end{aligned}$$

Since $\hat{\theta} = \theta + \omega_1 - \omega_2$, the first and third terms on the right-hand side of the above equation is equal to zero. As a result,

$$\log \frac{f(\omega^n | \theta, \mathbf{g}^1)}{f(\omega^n | \hat{\theta}, \mathbf{g}^2)} = 2(\omega_1 - \omega_2)(\theta - \omega_2) / \sigma^2$$

which is a finite constant that does not depend on n . Hence, the inequality in (9) reduces to

$$\limsup_{n \rightarrow \infty} \frac{f(\theta | s_0, \omega^n)}{f(\hat{\theta} | s_0, \omega^n)} < \infty$$

almost surely. Hence, the belief that agent 0 assigns to the true state θ remains bounded away from 1. This observation, coupled with the fact that a Bayesian agent never rules out the true state as impossible implies that agent 0 remains uncertain forever with probability one. \square

Proof of Proposition 2.

We start by analyzing the complete information benchmark. Under complete information about the realized network, the first-order condition of agent i is given by

$$x_i = \theta_i + \alpha \sum_{j \neq i} g_{ij} x_j.$$

Therefore, the vector of equilibrium actions in the complete information variant of the game is equal to $x^{\text{com}} = (I - \alpha \mathbf{g})^{-1} \theta$, thus establishing (2). Note that the assumptions that $\sum_{j \neq i} \max \text{supp}(f_{ij}) \leq 1$ and $\alpha < 1$ guarantee that $(I - \alpha \mathbf{g})^{-1}$ always exists and is element-wise non-negative.

Next, consider the incomplete information game, in which agent i only observes the realization of the weights in the set $\{g_{ij} : j \neq i\}$, with no information about g_{jk} for $j \neq i$. We establish (3) by verifying that the vector $x = (I + \alpha \mathbf{g} L) \theta$ satisfies the best-response equations of all agents simultaneously, where $L = (I - \alpha \mathbb{E}[\mathbf{g}])^{-1}$. To this end, recall that the first-order condition of agent i is given by

$$(10) \quad x_i = \theta_i + \alpha \sum_{j \neq i} g_{ij} \mathbb{E}_i[x_j],$$

where $\mathbb{E}_i[x_j]$ is i 's expectation of j 's action conditional on observing her own local neighborhood weights g_{ij} . Suppose that the action of each agent $j \neq i$ is given by the conjectured actions $x_j = \theta_j + \alpha \sum_{k \neq j} \sum_{r=1}^n g_{jk} \ell_{kr} \theta_r$. Plugging this expression into the left-hand side of agent i 's first-order condition in (10) implies that

$$x_i = \theta_i + \alpha \sum_{j \neq i} g_{ij} \theta_j + \alpha^2 \sum_{j \neq i} \sum_{k \neq j} \sum_{r=1}^n g_{ij} \mathbb{E}[g_{jk}] \ell_{kr} \theta_r,$$

where we are using the fact that $\mathbb{E}_i[g_{jk}] = \mathbb{E}[g_{jk}]$. This is a consequence of the assumption that the information set of agent i provides no information about g_{jk} when $j \neq i$. Furthermore, the fact that $L = \sum_{m=0}^{\infty} \alpha^m \mathbb{E}^m[\mathbf{g}]$ implies that $\alpha \mathbb{E}[\mathbf{g}]L = L - I$. Therefore, it is immediate that

$$x_i = \theta_i + \alpha \sum_{j \neq i} \sum_{r=1}^n g_{ij} \ell_{jr} \theta_r.$$

Consequently, the strategy profile $x = (I + \alpha \mathbf{g}L)\theta$ satisfies the first-order conditions of all agents simultaneously, establishing (3).

We complete the proof of Proposition 2 by establishing that $\mathbb{E}[x_i^{\text{inc}}] < \mathbb{E}[x_i^{\text{com}}]$ for all agents i . As a first observation, note that equation (3) implies the ex ante vector of equilibrium actions in the incomplete information game is given by $\mathbb{E}[x^{\text{inc}}] = \mathbb{E}[I + \alpha \mathbf{g}(I - \alpha \mathbb{E}[\mathbf{g}])^{-1}]\theta$. Consequently,

$$(11) \quad \mathbb{E}[x^{\text{inc}}] = \sum_{k=0}^{\infty} \alpha^k \mathbb{E}^k[\mathbf{g}]\theta.$$

On the other hand, equation (2) implies that the ex ante vector of equilibrium actions of actions in the complete information benchmark is given by

$$(12) \quad \mathbb{E}[x^{\text{com}}] = \mathbb{E}[(I - \alpha \mathbf{g})^{-1}]\theta = \sum_{k=0}^{\infty} \alpha^k \mathbb{E}[\mathbf{g}^k]\theta.$$

The juxtaposition of (11) and (12) implies that the proof is complete once we show that $\mathbb{E}[\mathbf{g}^k] \geq \mathbb{E}^k[\mathbf{g}]$ for all non-negative integers k , with at least one inequality being strict. We show this by establishing that $\mathbb{E}[g_{j_1, j_2} g_{j_2, j_3} \cdots g_{j_{k-1}, j_k}] \geq \mathbb{E}[g_{j_1, j_2}] \cdots \mathbb{E}[g_{j_{k-1}, j_k}]$ for any sequence of agents j_1, j_2, \dots, j_k and all k .

To prove this claim, note that if $j_m = j_{m+1}$, then both sides of the above inequality are equal to zero. Therefore, suppose that the sequence of agents j_1, j_2, \dots, j_k is such that $j_m \neq j_{m+1}$ throughout. We have,

$$\mathbb{E}[g_{j_1, j_2} g_{j_2, j_3} \cdots g_{j_{k-1}, j_k}] = \mathbb{E} \left[\prod_{i=1}^n \prod_{l \neq i} \prod_{r \in B(i, l)} g_{j_r, j_{r+1}} \right] = \prod_{i=1}^n \prod_{l \neq i} \mathbb{E} \left[\prod_{r \in B(i, l)} g_{j_r, j_{r+1}} \right]$$

where $B(i, l) = \{r : j_r = i, j_{r+1} = l\}$ and the second equality is a consequence of the assumption that all elements of \mathbf{g} are drawn independently from one another. Let $b_{il} = |B(i, l)|$. We have,

$$\mathbb{E}[g_{j_1, j_2} g_{j_2, j_3} \cdots g_{j_{k-1}, j_k}] = \prod_{i=1}^n \prod_{l \neq i} \mathbb{E}[g_{il}^{b_{il}}] \geq \prod_{i=1}^n \prod_{l \neq i} \mathbb{E}[g_{il}]^{b_{il}},$$

where the second implication is a consequence of Jensen's inequality. Furthermore, note that the inequality is strict whenever there exists pair of agents i and l in the sequence such

that $b_{il} \geq 2$. As a result,

$$\mathbb{E}[g_{j_1, j_2} g_{j_2, j_3} \cdots g_{j_{k-1}, j_k}] \geq \prod_{i=1}^n \prod_{l \neq i} \prod_{r \in B(i, l)} \mathbb{E}[g_{j_r, j_{r+1}}] = \mathbb{E}[g_{j_1, j_2}] \mathbb{E}[g_{j_2, j_3}] \cdots \mathbb{E}[g_{j_{k-1}, j_k}].$$

The above inequality guarantees $\mathbb{E}[\mathbf{g}^k] \geq \mathbb{E}^k[\mathbf{g}]$ for all integer $k \geq 0$ and hence completes the proof. \square

Proof of Proposition 3. Consider a monotone symmetric equilibrium of the incomplete information game, according to which each agent i with attribute $z_i \in \{a, b\}$ chooses to link to an agent j with attribute z_j if and only if ϵ_{ij} falls below the threshold $\tau_{z_i z_j}$. Hence, any monotone symmetric equilibrium is characterized by four thresholds τ_{aa} , τ_{ab} , τ_{ba} , and τ_{bb} .

As a starting point, we determine the moments of these thresholds that are identifiable from the observed network \mathbf{g} . Recall that the formation of each link requires the mutual consent of both parties. This means that a link between agents i and j with attributes $z_i = z_j = a$ exists if and only if $\epsilon_{ij}, \epsilon_{ji} \leq \tau_{aa}$. Consequently,

$$(13) \quad p_{aa} = F^2(\tau_{aa}),$$

where p_{aa} is the fraction of linkages within type- a agents and $F(\cdot)$ denotes the common distribution of preference shocks. A similar argument implies that

$$(14) \quad p_{bb} = F^2(\tau_{bb})$$

$$(15) \quad p_{ab} = F(\tau_{ab})F(\tau_{ba}),$$

where p_{bb} and p_{ab} denote the fraction of linkages within type- b agents and across agents of types a and b , respectively.

Next we use the equilibrium conditions to relate the model's structural parameters α , β , and γ to the observables. More specifically, note that for τ_{aa} to correspond to an equilibrium threshold, agent i with attribute $z_i = a$ must be indifferent between linking to an agent j with the same attribute whenever $\epsilon_{ij} = \tau_{aa}$. In other words,

$$\tau_{aa} = \alpha + \beta + \frac{\gamma}{n-1} \sum_{k \neq i, j} \mathbb{E}[g_{jk}].$$

Furthermore, note that $\mathbb{E}[g_{jk}] = F^2(\tau_{aa})$ if $z_j = z_k = a$, whereas $\mathbb{E}[g_{jk}] = F(\tau_{ab})F(\tau_{ba})$ if $z_j = a$ and $z_k = b$. Replacing for these terms from equations (13) and (15) therefore implies that

$$(16) \quad F^{-1}(\sqrt{p_{aa}}) = \alpha + \beta + \frac{\gamma}{n-1} \left((n\lambda - 2)p_{aa} + n(1-\lambda)p_{ab} \right),$$

where λ denotes the fraction of agents with attribute a . A similar argument applied to linkages within type- b agents implies that

$$(17) \quad F^{-1}(\sqrt{p_{bb}}) = \alpha + \beta + \frac{\gamma}{n-1} \left(n\lambda p_{ab} + (n(1-\lambda) - 2)p_{bb} \right).$$

Equations (16) and (17) provide the econometrician with two equations that relate network observables to the structural parameters α , β , and γ and hence can be used for identification. Specifically, as long as

$$(18) \quad \lambda \neq \frac{(n-2)(p_{bb} - p_{ab}) + 2(p_{aa} - p_{ab})}{n(p_{aa} + p_{bb} - 2p_{ab})},$$

the two equations are linearly independent and hence γ and $\alpha + \beta$ are point-identified.

With the above in hand, the next step is to determine conditions under which α and β are separably identifiable. To this end, note that τ_{ab} and τ_{ba} correspond to equilibrium thresholds if and only if

$$\begin{aligned} \tau_{ab} &= \alpha + \frac{\gamma}{n-1} \left((n\lambda - 1)p_{ab} + (n(1-\lambda) - 1)p_{bb} \right) \\ \tau_{ba} &= \alpha + \frac{\gamma}{n-1} \left((n\lambda - 1)p_{aa} + (n(1-\lambda) - 1)p_{ab} \right), \end{aligned}$$

where once again we are using equations (13)–(15) to replace for the thresholds in terms of the observables fractions p_{aa} , p_{ab} , and p_{bb} . Consequently,

$$\begin{aligned} \tau_{ab} &= F^{-1}(\sqrt{p_{bb}}) - \beta + \gamma(p_{bb} - p_{ab})/(n-1) \\ \tau_{ba} &= F^{-1}(\sqrt{p_{aa}}) - \beta + \gamma(p_{aa} - p_{ab})/(n-1). \end{aligned}$$

Using equation (15) one more time implies that $G(\beta, \gamma) = p_{ab}$, where

$$G(\beta, \gamma) = F \left(F^{-1}(\sqrt{p_{aa}}) - \beta + \gamma \frac{p_{aa} - p_{ab}}{n-1} \right) \cdot F \left(F^{-1}(\sqrt{p_{bb}}) - \beta + \gamma \frac{p_{bb} - p_{ab}}{n-1} \right).$$

Recall that, by assumption, the preference shocks have full support over \mathbb{R} , which implies that F is strictly increasing. This means that $G(\beta, \gamma)$ is strictly decreasing over its domain. This observation, alongside the fact that $\lim_{\beta \rightarrow \infty} G(\beta, \gamma) = 0$ and $\lim_{\beta \rightarrow -\infty} G(\beta, \gamma) = 1$, implies that for every γ , there exists a unique β such that $G(\beta, \gamma) = p_{ab}$. This implying that β is also point-identified. Since $\alpha + \beta$ was already identified, this means that all structural parameters are point-identified for all values of n .

The proof is complete once we verify that, as $n \rightarrow \infty$, the identification requirement on λ expressed in (18) is always satisfied as long as $\lambda \neq 1/2$. To this end, suppose that (18) is violated, that is,

$$\lambda = \frac{(n-2)(p_{bb} - p_{ab}) + 2(p_{aa} - p_{ab})}{n(p_{aa} + p_{bb} - 2p_{ab})},$$

which reduces to

$$(19) \quad \lambda = \frac{p_{bb} - p_{ab}}{p_{aa} + p_{bb} - 2p_{ab}}$$

as $n \rightarrow \infty$. Plugging back this expression into equations (16) and (17) implies that (18) is violated asymptotically only if $p_{aa} = p_{bb}$, which by (19), is equivalent to $\lambda = 1/2$. Thus, as long as $\lambda \neq 1/2$, the identification requirement (18) on λ is always satisfied as $n \rightarrow \infty$. \square

**ONLINE APPENDIX
NOT FOR PUBLICATION**

APPENDIX B. MAIN TABLES WITH CENTRALITIES DISPLAYED

TABLE B.1. Network Knowledge as a function of Distance and Centralities

VARIABLES	(1) Correct guess	(2) Correct guess	(3) Correct guess	(4) Correct guess
1{Reachable j,k, Not linked j,k}	0.0647 (0.0196) [0.001]	0.0639 (0.0192) [0.001]	0.0549 (0.0194) [0.006]	0.0420 (0.0224) [0.065]
1{Reachable j,k, Linked j,k}	0.547 (0.0227) [0.000]	0.147 (0.112) [0.193]	0.122 (0.108) [0.261]	0.174 (0.142) [0.224]
Avg. distance to j,k, Not linked j,k	-0.0177 (0.00398) [0.000]	-0.0153 (0.00389) [0.000]	-0.0115 (0.00400) [0.005]	-0.0133 (0.00437) [0.003]
Avg. distance to j,k, Linked j,k	-0.109 (0.00729) [0.000]	-0.0943 (0.00656) [0.000]	-0.0920 (0.00607) [0.000]	-0.0925 (0.00673) [0.000]
Centrality (std.) of i, Not linked j,k	0.00265 (0.00285) [0.356]	0.00224 (0.00286) [0.437]	0.00325 (0.00280) [0.250]	
Centrality (std.) of i, Linked j,k	5.61e-05 (0.00521) [0.991]	0.00152 (0.00520) [0.772]	0.00247 (0.00510) [0.629]	-6.32e-06 (0.00588) [0.999]
Avg. centrality (std.) of j,k, Not linked j,k	-0.000971 (0.00343) [0.778]	-0.000282 (0.00328) [0.932]	0.00263 (0.00319) [0.412]	-0.000425 (0.00349) [0.903]
Avg. centrality (std.) of j,k, Linked j,k	-0.00820 (0.00513) [0.114]	-0.00782 (0.00521) [0.138]	-0.00704 (0.00505) [0.167]	-0.00619 (0.00524) [0.241]
Observations	66,810	66,810	66,810	66,810
Dimension FE	✓	✓	✓	✓
Village FE			✓	
Respondent FE				✓
Demographic controls		✓	✓	✓
Depvar mean	0.199	0.199	0.199	0.199

Notes: Standard errors (clustered at village level) are reported in parentheses. p -values are reported in brackets. Demographic controls include average geographic distance of i from j and k and dummies for each of whether i is of the same caste, subcaste, has the same electrification status, has the same roof type, has the same number of rooms, and has the same house ownership status for each of j and k , when g_{jk} is 1 and 0. They also include the above-mentioned dummies for when j and k have the same demographic traits.

TABLE B.2. Don't Know as a function of Distance and Centralities

VARIABLES	(1) Don't know	(2) Don't know	(3) Don't know	(4) Don't know
1{Reachable j,k, Not linked j,k}	-0.321 (0.0377) [0.000]	-0.271 (0.0362) [0.000]	-0.246 (0.0288) [0.000]	-0.253 (0.0335) [0.000]
1{Reachable j,k, Linked j,k}	-0.466 (0.0381) [0.000]	-0.355 (0.153) [0.023]	-0.244 (0.163) [0.139]	-0.228 (0.216) [0.293]
Avg. distance to j,k, Not linked j,k	0.109 (0.00866) [0.000]	0.0933 (0.00939) [0.000]	0.0854 (0.00653) [0.000]	0.0811 (0.00686) [0.000]
Avg. distance to j,k, Linked j,k	0.148 (0.00873) [0.000]	0.123 (0.00788) [0.000]	0.118 (0.00784) [0.000]	0.122 (0.00887) [0.000]
Centrality (std.) of i, Not linked j,k	0.00201 (0.00608) [0.742]	0.00139 (0.00617) [0.822]	-0.00383 (0.00531) [0.472]	
Centrality (std.) of i, Linked j,k	-0.00335 (0.00595) [0.574]	-0.00565 (0.00587) [0.339]	-0.00979 (0.00571) [0.091]	-0.00396 (0.00649) [0.543]
Avg. centrality (std.) of j,k, Not linked j,k	-0.00535 (0.00761) [0.484]	-0.00438 (0.00740) [0.556]	-0.00927 (0.00675) [0.174]	-0.00413 (0.00641) [0.522]
Avg. centrality (std.) of j,k, Linked j,k	0.00241 (0.00596) [0.687]	0.00109 (0.00614) [0.859]	0.000736 (0.00604) [0.903]	0.00668 (0.00650) [0.307]
Observations	66,810	66,810	66,810	66,810
Dimension FE	✓	✓	✓	✓
Village FE			✓	
Respondent FE				✓
Demographic controls		✓	✓	✓
Depvar mean	0.463	0.463	0.463	0.463

Notes: Standard errors (clustered at village level) are reported in parentheses. p -values are reported in brackets. Demographic controls include average geographic distance of i from j and k and dummies for each of whether i is of the same caste, subcaste, has the same electrification status, has the same roof type, has the same number of rooms, and has the same house ownership status for each of j and k , when g_{jk} is 1 and 0. They also include the above-mentioned dummies for when j and k have the same demographic traits.

TABLE B.3. Correct Guess, conditional on sure guesses only, as a function of distance and centralities

VARIABLES	(1) Correct guess	(2) Correct guess	(3) Correct guess	(4) Correct guess
1{Reachable j,k, Not linked j,k}	0.00979 (0.0432) [0.821]	0.0348 (0.0424) [0.414]	0.0327 (0.0368) [0.378]	-0.00996 (0.0484) [0.837]
1{Reachable j,k, Linked j,k}	0.534 (0.0423) [0.000]	-0.158 (0.108) [0.148]	-0.101 (0.127) [0.427]	0.108 (0.359) [0.764]
Avg. distance to j,k, Not linked j,k	0.00668 (0.00699) [0.342]	0.00541 (0.00770) [0.484]	0.00857 (0.00684) [0.214]	0.000630 (0.00720) [0.930]
Avg. distance to j,k, Linked j,k	-0.0189 (0.00697) [0.008]	-0.0208 (0.00711) [0.005]	-0.0177 (0.00727) [0.017]	-0.0110 (0.00924) [0.239]
Centrality (std.) of i, Not linked j,k	0.00605 (0.00530) [0.257]	0.00496 (0.00533) [0.355]	0.00213 (0.00526) [0.686]	
Centrality (std.) of i, Linked j,k	-0.00265 (0.00510) [0.605]	-0.00169 (0.00523) [0.747]	-0.00366 (0.00520) [0.484]	0.00105 (0.00851) [0.902]
Avg. centrality (std.) of j,k, Not linked j,k	-0.00440 (0.00605) [0.470]	-0.00280 (0.00608) [0.647]	-0.00276 (0.00544) [0.613]	-0.00474 (0.00666) [0.479]
Avg. centrality (std.) of j,k, Linked j,k	-0.00999 (0.00587) [0.093]	-0.0111 (0.00603) [0.070]	-0.0103 (0.00631) [0.105]	-0.00679 (0.00768) [0.379]
Observations	35,909	35,909	35,909	35,909
Sure guess only	✓	✓	✓	✓
Dimension FE	✓	✓	✓	✓
Village FE			✓	
Respondent FE				✓
Demographic controls		✓	✓	✓
Depvar mean	0.371	0.371	0.371	0.371

Notes: Standard errors (clustered at village level) are reported in parentheses. p -values are reported in brackets. Demographic controls include average geographic distance of i from j and k and dummies for each of whether i is of the same caste, subcaste, has the same electrification status, has the same roof type, has the same number of rooms, and has the same house ownership status for each of j and k , when g_{jk} is 1 and 0. They also include the above-mentioned dummies for when j and k have the same demographic traits.

TABLE B.4. Network Knowledge, conditional on knowing j,k , as a function of Distance and Centralities

VARIABLES	(1) Know pair	(2) Know pair	(3) Don't know	(4) Don't know
1{Reachable j,k , Not linked j,k }	0.337 (0.0393) [0.000]	0.289 (0.0389) [0.000]	-0.0671 (0.0355) [0.063]	-0.0490 (0.0337) [0.150]
1{Reachable j,k , Linked j,k }	0.446 (0.0395) [0.000]	0.349 (0.174) [0.049]	-0.122 (0.0338) [0.001]	-0.0548 (0.0672) [0.418]
Avg. distance to j,k , Not linked j,k	-0.109 (0.00923) [0.000]	-0.0959 (0.00975) [0.000]	0.0323 (0.00736) [0.000]	0.0229 (0.00728) [0.002]
Avg. distance to j,k , Linked j,k	-0.137 (0.00948) [0.000]	-0.116 (0.00834) [0.000]	0.0447 (0.00641) [0.000]	0.0349 (0.00738) [0.000]
Centrality (std.) of i , Not linked j,k	0.000393 (0.00531) [0.941]	0.000918 (0.00544) [0.866]	0.00344 (0.00490) [0.485]	0.00270 (0.00482) [0.577]
Centrality (std.) of i , Linked j,k	0.00944 (0.00577) [0.106]	0.0102 (0.00565) [0.074]	0.00644 (0.00521) [0.220]	0.00424 (0.00523) [0.420]
Avg. centrality (std.) of j,k , Not linked j,k	0.00927 (0.00784) [0.241]	0.00697 (0.00758) [0.361]	0.00272 (0.00653) [0.679]	0.000500 (0.00663) [0.940]
Avg. centrality (std.) of j,k , Linked j,k	-0.00227 (0.00602) [0.707]	-0.00296 (0.00602) [0.625]	-0.000628 (0.00369) [0.865]	-0.00357 (0.00404) [0.380]
Observations	22,270	22,270	42,750	42,750
Know pair only			✓	✓
Dimension FE	N/A	N/A	✓	✓
Demographic controls		✓		✓
Depvar mean	0.640	0.640	0.160	0.160

Notes: Standard errors (clustered at village level) are reported in parentheses. p -values are reported in brackets. Demographic controls include average geographic distance of i from j and k and dummies for each of whether i is of the same caste, subcaste, has the same electrification status, has the same roof type, has the same number of rooms, and has the same house ownership status for each of j and k , when g_{jk} is 1 and 0. They also include the above-mentioned dummies for when j and k have the same demographic traits.