Multiplex Network Diffusion

Valentina A. Assenova*[1] and Nidhi Sridhar[2]

[1] *Management Department, The Wharton School, University of Pennsylvania*
[2] *Department of Computer and Information Science, University of Pennsylvania*

January 29, 2019

Abstract

The diffusion of information through social networks presents an important area of research in economics, with applications to a variety of processes including the spread of new technologies. The majority of existing research has focused on modeling diffusion in simplex (one-layer) networks, yet many networks are multiplex (multi-layered), spanning many kinds of ties. We consider the role that multiplexity plays in diffusion, specifically through information bridging across structurally similar (vs. dissimilar) layers. Our simulations reveal that these layers exhibit distinct properties affecting diffusion, such as connectedness, clustering, and density, and are associated with different diffusion curves over time.

Keywords: Networks; Multiplexity; Diffusion
JEL Codes: O3; D8; Z1

*Correspondence: vaa@wharton.upenn.edu. Address: 3620 Locust Walk, Philadelphia, PA 19104. This research was generously supported by the Mack Institute for Innovation Managment at the Wharton School of the University of Pennsylvania. The authors gatefully acknowledge helpful comments from Greta Hsu, Rahul Kapoor, and participants at the 2018 meeting of the International Network for Analytic Sociologists (INAS) at the Stanford Institute for Economic Policy Research (SIEPR). The authors are solely responsible for any errors.

1 Introduction

The role that networks play in the transfer and spread of information has generated considerable interest in recent research (Chandrasekhar et al., 2018; Banerjee et al., 2018; Alatas et al., 2015). Network connections have been shown to play a vital role in many diffusion processes involving information transfer and information aggregation through ties (Banerjee et al., 2018; Alatas et al., 2015), such as the spread of ideas in online and virtual communities (Goel et al., 2016; Salganik et al., 2006), the adoption of new technologies (Angst et al., 2010), the implementation of new policies (Curran, 2015), and the search for knowledge in organizations (Rosenkopf and Almeida, 2003; Powell et al., 1996; Singh, 2005).

Much of the existing research has examined network ties from the point of view of networks comprising a single type of tie, for example friendship (Rossman, 2014; Centola, 2015). Yet, recent research suggests that networks comprising multiple kinds of ties – multiplex networks – are important for diffusion and generate novel dynamical behavior that differs from behavior in single-tie (monoplex) networks (Yagan and Gligor, 2012; Goel et al., 2016; Wang et al., 2017). In light of these dynamics, recent work by Myers and Leskovec (2012) and Melnik et al. (2013) has called for understanding network complexity and its consequences.

Our aim in this article is to examine this question by using computer simulations on empirically-derived multiplex networks, and in doing so, to contribute to the emerging literature on network complexity and its consequences. Methodologically, we draw on data from a prior study in development economics, collected by Banerjee et al.(2013) in 43 villages in Karnataka, India. These data comprise multiple networks and information about the diffusion of microfinance in each of these networks. To study the role of multiplexity in these dynamics, we first identify distinct groupings of ties associated with different information bridging mechanisms: redundant bridging through ties spanning structurally similar subgraphs ("homogeneous multiplexity") and non-redundant bridging through ties spanning structurally dissimilar subgraphs ("heterogeneous multiplexity"). We focus on structural redundancy as an extension of prior results that network redundancy facilitates information search but limits information novelty (Reagans and Zuckerman, 2008; Aral and Alstyne, 2011). We therefore examine the diffusion consequences of these distinct types of multiplexity by conducting network simulations and testing for differences in structural properties.

Our findings have both theoretical and empirical implications. Theoretically, we argue that multiplex networks can give rise to multiple information cascades with different diffusion curves depending on whether information is bridged through ties that span structurally homogeneous (vs. heterogeneous) subgraphs. Empirically, we demonstrate that homogeneous and heterogeneous subgraphs differ in important structural properties affecting the speed and breadth of diffusion, such as clustering, connectedness, and density, and produce different adoption curves. These results suggest that a community's receptivity to new technologies (or ideas, practices, etc.) depends on the extent to which different types of ties connecting individuals represent structurally distinct pathways for obtaining information and social validation. We discuss several reasons why subgraph heterogeneity promotes diffusion, and conclude with implications for future studies.

2 Multiplex Networks and Diffusion

Prior research distinguishes among three mechanisms behind diffusion: (i) social contagion, (ii) social influence, and (iii) social learning (Rogers, 1983; Young, 2009, 2011; Wejnert, 2002). Social contagion involves a logic of mimicry – one adopts because one's contacts have adopted (Strang and Soule, 1998). An example of this is buying a new gadget because all of one's close friends have one. Social influence, by contrast, involves a logic of norm-setting: one adopts when there is a critical mass of prior adopters (Rogers, 1983). An example of this is buying a new gadget because most people in one's school or workplace have bought one. The last mechanism – social learning – involves information-gathering from others by talking to them and observing their behavior (Young, 2009; Chandrasekhar et al., 2018; Banerjee et al., 2018). As the term implies, this last mechanism involves information processing and judgment and decision-making before one adopts.

Evidence and theory concur that these mechanisms produce different diffusion curves (Young, 2009; Wejnert, 2002; Coleman et al., 1966). For instance, both social influence and social contagion produce rapid initial adoption and subsequent tapering to a steady-state adoption level, resembling a log function. Social learning, by contrast, produces a distinct S-shaped curve characterized by slow initial adoption, rapid spike, and gradual saturation (Young, 2009, 2011). In many cases, learning happens within a social context, that is, within a network of contacts from whom one obtains information (Rossman, 2015; Chandrasekhar et al., 2018; Aral et al., 2009).

2.1 Why does multiplexity matter?

Under conditions of imperfect information, agents are motivated to search for information broadly by consulting with those connected to them. Other people's opinions, communications, and observed behaviors can potentially reveal novel information about the value of a new technology (or idea, practice, etc.) (Assenova, 2018; Chandrasekhar et al., 2018; Banerjee et al., 2018). As more people adopt, the information uncertainty around the potential value is reduced and more people are induced to adopt.

Although prior work has identified important drivers of diffusion through learning in networks, such as social similarity (Aral et al., 2009), network consolidation (Centola, 2015), network positions ("centrality," "structural equivalence") (Burt, 1987; Banerjee et al., 2018, 2013a), and "wide" ties (Centola and Macy, 2007), the existing literature has largely overlooked variation in how multiplex ties are composed. By tie composition, we mean the constitutive types of ties that connect different agents in a multiplex network. For instance, an agent can be connected to others by friendship, co-authorship, advisory ties, and so forth. These ties can overlap in myriads ways. One can be friends with a co-author, for instance, or a former advisee. Having many types of connections presents more opportunities for an agent to learn, but only if those connections are not redundant, i.e., if one's co-author, advisor, and friend are not the same person. Thus, the ability to learn – and diffuse – depends on structural dissimilarity in the connections through which an agent is connected to others in a multiplex network.

At a network level, multiplexity means that ties span different sub-networks, for example friendship and co-authorship (Gould, 1991; Shipilov, 2012). On a dyadic level, multiplexity increases the "width" of a tie connecting two actors, meaning that these actors are now linked through multiple information pathways (Smith and Papachristos, 2016; Centola and Macy, 2007). Multiplex ties are therefore thought to represent stronger forms of connectivity and to involve higher levels of trust and cooperation (Ferriani et al., 2013; Shipilov, 2012; Uzzi, 1996). For these reasons, many prior studies have viewed multiplexity as beneficial for learning and diffusion (Centola and Macy, 2007; Gould, 1991; Centola, 2015).

Yet, multiplexity can also involve similar combinations of ties that are redundant in their overall topologies (Ferriani et al., 2013; Shipilov, 2012; Smith and Papachristos, 2016). Specifically, at a network level, multiplex ties can span structurally similar subgraphs – for instance friendship and advice connections – or structurally dissimilar subgraphs – for instance family and co-working connections. Our aim is to distinguish between these types of multiplex ties and propose ways in which these types affect information redundancy and learning within multiplex networks over time.

2.2 Multiplexity and Information Redundancy

Many of the benefits of network connections for information transfer depend on non-redundancy – the notion that different network ties provide novel and unique sources of information (Burt, 1992). Brokers, for instance, are in structural positions of advantage insofar as they are situated as a "bridge" between two unique sources of information (Burt, 2012; Ryall and Sorenson, 2007). This notion of information non-redundancy, we argue, is at the heart of understanding the value of multiplex ties for information bridging. When the same sets of actors are connected through the same sets of edges across multiple types of ties, these ties do not provide novel sources of information through either the ties or the actors. The opposite is true when different actors are connected through structurally dissimilar sets of edges. Drawing on these distinctions, we define homogeneous and heterogeneous multiplex ties as follows:

Homogeneous Multiplex Ties.– Homogeneous multiplex ties span multiple, but structurally similar, subgraphs of a multiplex network. Structural similarity means that the overall tie topologies – the patterns of connections – among actors involve similar sets of nodes and similar configurations of edges among these nodes. When graphed as networks, these sub-graphs appear the same and exhibit similar configurations of nodes and edges. For example, within a community (a closed social group), the network of friendship connections (who is friends with whom) might resemble the network of social connections (who interacts socially with whom) because people that one spends time with are more likely to become friends, and people who are friends are more likely to spend time together. In cases such as these, the friendship and social networks will look structurally similar, with the same sets of individuals being connected to each other through both social and friendship ties. In networks that are highly similar, even though the ties represent different types of relations, because these relations are formed among the same sets of individuals across both networks, the emergent multiplex ties are homogeneous, meaning that they contribute information from the same sources.

Heterogeneous Multiplex Ties.– Heterogeneous multiplex ties span multiple but structurally dissimilar subgraphs of a multiplex network. Structural dissimilarity means that the overall tie topologies among actors involve different sets of nodes and different configurations of edges among these nodes. When graphed as networks, these subgraphs look different and exhibit distinct configurations of nodes and edges in a graph. For instance, if one examines a research network within a university, and subsequently examines patterns of friendship formation among colleagues at this university, one might discover that that the co-author and friendship networks overlap in who is connected to whom somewhat, but that in many ways these networks are structurally distinct.

2.3 Information Bridging

The concept of network redundancy, which has generally been applied to single-tie networks, pertains to the notion of having multiple people from whom one can access information, but who all know the same things. Prior research has argued that redundancy diminishes the value of social connections for obtaining diverse and relevant information through shared social connections (Burt, 1992; Aral and Alstyne, 2011; Reagans and Zuckerman, 2008). Thus, even though an actor may be central in a network and have many social pathways to accessing information (i.e. have power in the network), this actor may nevertheless lack knowledge in the sense of having a large diversity of information (Reagans and Zuckerman, 2008). Thus, networks can be optimized for either informational diversity (and low redundancy) or high bandwidth (and high redundancy) (Aral and Alstyne, 2011). Yet prior work has not linked this notion of redundancy to the concept of multiplexity when actors simultaneously interact across multiple, interconnected networks.

This concept of redundant information and sources of social validation can be extended to multiplex networks by examining overlap and redundancy in ties across different layers (subgraphs) of a network. In multiplex networks, redundancy can take two forms: having similar patterns of connections across network subgraphs, and having many shared connections (triadic closure) within a single subgraph. Thus, the presence of multiplex ties alone is not guaranteed to increase either the diversity of one's contacts, or the diversity of information from these contacts.

Homogeneous multiplex ties, although they bridge multiple subgraphs of a network, fail to bridge information "expansively," that is, in ways that include nodes in the network that would not be otherwise reached through other subgraphs. When ties in one subgraph are isomorphic to ties in another subgraph, even though these ties "span" multiple subgraphs, they are in essence not providing any additional value: they are structurally redundant. This redundancy reduces the value of spanning multiple subgraphs, and hence the benefits of multiplex ties for promoting broader and faster information diffusion within a network.

From an information bridging perspective, heterogeneous multiplex ties represent "bridges" across non-redundant sources of information that are otherwise disconnected. When subgraphs within a multiplex network are not perfectly isomorphic, ties that span these subgraphs have the capacity to spread information transmitted in one subgraph into another subgraph, producing multi-stage information cascades. The presence of these ties means that information can travel faster and reach more nodes that are otherwise not reached through connections present within a single subgraph of the network.

2.4 Structural Properties

Prior research has shown that there is a close relationship between the types of connections linking actors and the appearance of the overall topology of the network emerging from these ties. Romantic networks, for instance, resemble chains of dyads, which generally remain fairly sparse (Bearman et al., 2004), whereas communication networks are dense (Cardillo et al., 2013). Multiplex networks combine multiple types of ties that themselves can be organized as different topological structures when examined as subgraphs of the larger network. These subgraphs can thus differ in their structural properties depending on the type of network ties that comprise them (Ferriani et al., 2013; Kim and Goh, 2013). Networks that are highly consolidated through many shared actors can actually create highly fragmented communities where information does not flow freely across group boundaries (Centola, 2015; Vedres and Stark, 2010). Yet, the lack of consolidation through some overlap in actors or ties can similarly reduce communication and cause integration problems within networks, which undermine diffusion (Centola, 2015; Wang et al., 2017).

Topologically, homogeneous and heterogeneous multiplex ties could exhibit different network properties owing to their different functions in information search. Homogeneous multiplex ties, for instance, arise when agents deepen their ties with existing contacts to obtain better information (i.e. search along the "intensive margin"), whereas heterogeneous ties arise when agents seek novel information outside of existing contacts (i.e. search along the "extensive margin") (Ferriani et al., 2013; Smith and Papachristos, 2016; Gómez-Gardeñes et al., 2012). The topology – and function – of different subgraphs within a multiplex network should therefore depend on the type of information bridging that these ties provide. Multiplex networks comprising mostly homogeneous ties should be more likely to involve the same sets of agents, effectively “collapsing” into less connected, lower density structures when redundancy is accounted for. By contrast, networks comprising mostly heterogeneous multiplex ties should be likely to exhibit high connectedness and low clustering, owing to non-redundancy.

3 Data

The network data we use for our simulations come from an intervention conducted in 2006 by Banerjee et al. (2013) designed to promote the spread of microfinance (MF). Data collection and materials and methods are described in detail in Banerjee et al. (2013). The intervention that Banerjee et al. (2013) conducted was to “seed” the social networks of each village with information among people from the villages who could act as “leaders” (influencers) to promote MF. The timing of the intervention was designed to first raise awareness about MF and then leverage network connections in the villages to promote adoption (enrollment).

Individuals in the sampling frame were asked to name others with whom they engaged in eight distinct types of interactions: (1) money borrowing/lending (Money), (2) goods exchange of kerosene and rice (Goods), (3) home visits (Visits), (4) advice giving/receiving (Advice), (5) kinship relations (Kinship), (6) assistance with medical emergencies (Medical), (7) attendance of social events such as marriages and festivals in the village (Social), and (8) praying together at temple, church, or mosque (Religious). The data were intended to capture multiple dimensions of social interactions in the villages. Surveys were administered face-to-face and the names of contacts were recorded separately and subsequently coded with unique identifiers. Data from these surveys were then compiled into adjacency matrices for each village and each type of interaction, for a total of 344 graphs. In our analyses, we use data for the 43 villages for which social network data were collected. Of these, we were unable to obtain results for three villages (IDs:4, 15 and 23) because eigenvector centralities of the nodes in these villages were of the order e^{-07} and our model converged during initialization itself, as described below.[1]

4 Methods

4.1 Subgraph Generation

The data comprised eight types of relations connecting individuals in each village: Money, Advice, Goods, Visit, Medical, Social, Religious and Kinship. We used these data to generate a new "super" graph for each village with eight edge properties for each of these relations. The property could take on binary values: 0 or 1, 1 indicating that the edge in this relationship was present and 0 otherwise. For example, consider two villagers, X and Y. Suppose that X and Y had an edge (or were connected) in the Money, Medical, Advice, and Religious graphs. In our new super graph, the edge XY will have 8 edge properties whose values are as follows:

$$Money : 1 \quad Advice : 1 \quad Goods : 0 \quad Visit : 0 \quad Medical : 1 \quad Social : 0 \quad Religious : 1 \quad Kinship : 0$$

We generated the subgraphs $G_{exchange}, G_{communal}, G_{heterogenous}$ from each super graph by taking the pairwise union of the ties that belonged to each of these sets. These sets were defined based on hamming distance analyses of structural similarity, as described below. We performed this procedure for each village and stored these subgraphs as .rda files for our model

[1] The data and code for our analyses and results will be made publicly available upon publication through the Harvard Dataverse (`doi:10.7910/DVN/V3DZCU`), curated by the Institute for Quantitative Social Science.

simulations.

4.2 Structural Similarity

Using these super graphs, we conducted analyses based on the hamming distance across layers of the multiplex networks to detect distinct layers based on structural similarity. The hamming distance between the elements of two graphs g_1 and g_2 with adjacency matrices $\mathbf{A^{(1)}}$ and $\mathbf{A^{(2)}}$ was computed as:

$$d_H(g_1, g_2) = \sum_{i \neq j}^{N} \left[A_{ij}^{(1)} \neq A_{ij}^{(2)} \right]$$

This measure approaches zero as the structural similarity between two graphs increases. We used the `sna` package in `R` to compute the hamming distances across all pairwise combinations of subgraphs in each multiplex network (Butts, 2008). After computing the hamming distances, we then performed cluster analysis in `Stata 15 MP` using the `cluster` command to identify distinct groupings of structurally dissimilar subgraphs. These analyses revealed two primary clusters – the "communal" and "exchange" clusters – which comprised graphs that were structurally similar. The exchange cluster comprised {Money, Visit, Goods, Advice} ties. The communal cluster comprised {Social, Medical, Religious, Kinship} ties. Fig1 below illustrates these clusters.

5 Model Definition

We consider a village graph $G = (V, E)$ where V denotes the set of people in the village and E denotes the set of relations between these people. We then define subgraphs for each village graph, $G_{exchange}, G_{communal}$ and $G_{heterogenous}$ where each subgraph consists of nodes and relations that share a relation in at-least two types of Exchange, Communal or Heterogeneous relationships respectively. We compute this by creating subgraphs of villages

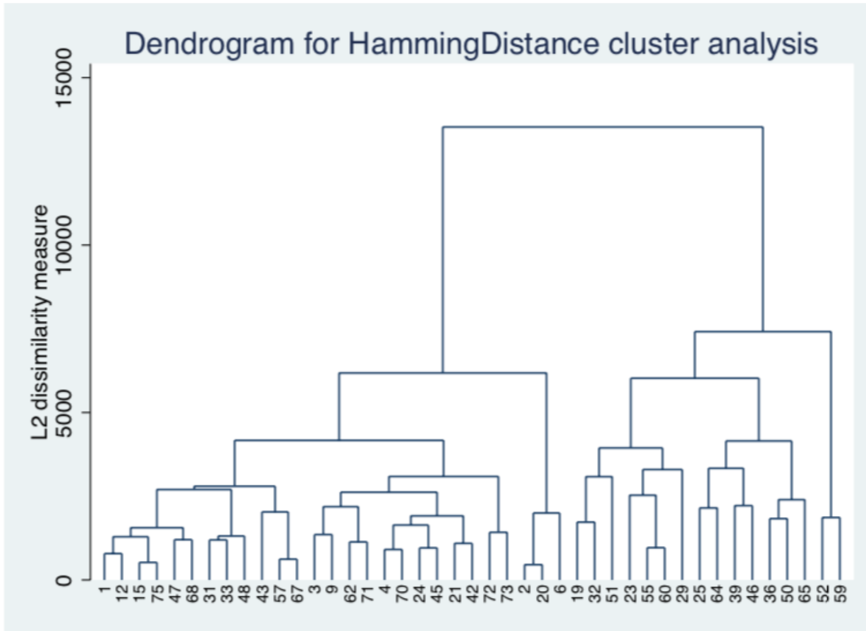

Figure 1: Hamming Distance Cluster Analysis

in each of the eight relations: Money, Advice, Goods, Visit, Medical, Social, Religious and Kinship. The $G_{exchange}$ is computed by taking the pairwise union of edges that exist in any of two Exchange type subgraphs- {Money, Visit, Goods, Advice}. Similarly, the $G_{Communal}$ is computed by the same idea with edges in the Communal type subgraphs - {Social, Medical, Religious, Kinship}. The heterogeneous subgraph is the subgraph from the village graph G that consists of edges not present in both $G_{exchange}$ and $G_{communal}$.

Now that the graphs have been defined, we move on to defining the model for adoption. For each villager v $\in$ V, we introduce a concept of initial probability $\theta_{v,t=1}$ and adoption threshold ϕ_v and an indicator variable to determine the adoption status of the villager V, A_v. If the villager has adopted, $A_v = 1$ else $A_v = 0$. The eigenvector centrality for each villager E_v as well as the set of neighbors(out degree of the villager) N_v is also computed.

In our model, we define that a user adopts when the probability to adopt at the given time instant t, $\theta_{v,t}$ is greater than the adoption threshold ϕ_v. To compute this probability

$\theta_{v,t}$, we define it as follows

$$\theta_{v,t} = \theta_{v,t-1} + \frac{\sum_{j \text{ in } N_v} E_j * A_j}{\text{count}(N_v)}$$

and

$$A_v = \begin{cases} 1, & \text{if } \theta_{v,t} \geq \phi_v \\ 0, & \text{otherwise} \end{cases}$$

6 Simulation

The model defined was simulated for each of the 43 villages. The adoption threshold for each villager was ϕ_v was drawn from a uniform distribution. While the initial probability of adoption at time t = 1, $\theta_{v,t=1}$ was drawn from a distribution such that the fraction of initial seed of adopters in both the simulated as well as the actual case were the same. For example, if Village 1 had 13% of initial adopters as observed. it was assured that 13 % of the villagers in the simulated case were initial adopters as well. However, these were picked randomly over 100 iterations. This step was done to ensure that the initial conditions before the simulation were similar. The simulation was performed in `R` using the packages `network, dplyr, sna` and `iGraph`.

For each village, the graphs showing adoption, at every iteration were stored as .gexf files and then loaded into Gephi for visualization. Graph properties for each village and subgraph were also computed and stored for further analysis.

7 Results

The fraction of adoption over time for a sample network (Village 2) is shown in Figure 2(a). The results for other villages were qualitatively similar and are omitted for brevity of presentation. The final states of adoption in each of the subgraphs of Village 2 are shown in Figure 2(b)-(c). The labeled nodes indicate that those individuals adopted ($A_v = 1$), while

Algorithm 1 Basic Algorithm

for *each village* V **do**
  Compute subgraphs $G_{exchange}, G_{communal}, G_{heterogenous}$
Repeat 100 iterations:
  for *each subgraph* g **do**
    Initialize A_1
    Compute E
    Initialize $\theta_{t=1}$, ϕ
    for *time* $t \in (2, T)$ **do**
      for *each villager* $v \in V$ **do**
        $\theta_{v,t} = \theta_{v,t-1} + \frac{\sum_{j \text{ in } N_v} E_j * A_j}{\text{count}(N_v)}$
        if $\theta_{v,t} \geq \phi_v$ **then**
          $A_v = 1$
    Store Fraction Adopted at Iteration
  Compute Average Fraction Adopted

the unlabeled black dots are the nodes that did not adopt by the end of the simulation. As this figure shows, the fraction of adoption in heterogeneous networks is much higher than that observed in the homogeneous (Exchange, Communal) subgraphs. We obtained the same result across all villages. Further, the adoption fraction was always higher in the Exchange subgraphs than in the Communal subgraphs.

Fig.3 illustrates the diffusion over time (at t = 1, 2, 3) in the heterogeneous and homogeneous (Exchange) subgraphs of the network of village 2 (the results for other villages were qualitatively similar). The left-hand-side panel (Fig.3(a),(c),(e)) shows diffusion within the subgraph of homogeneous (Exchange) multiplex ties, while the right-hand-side panel (Fig.3(b),(d),(f)) shows diffusion within the subgraph of heterogeneous multiplex ties. Colored nodes denote adopters ($A_v = 1$), while black nodes denote non-adopters.

As observed, the fraction of adopters in the heterogeneous subgraph was much higher over time than the fraction of adopters in the homogeneous subgraph. We observe the same results for all villages in our sample. We therefore tested for differences in means for the properties and adoption rates in these subgraphs by conducting paired two-sample t-tests with an equality of means null hypothesis. We report the results of these analyses in Table 1. The results revealed that, indeed, the average rate of adoption across these two types

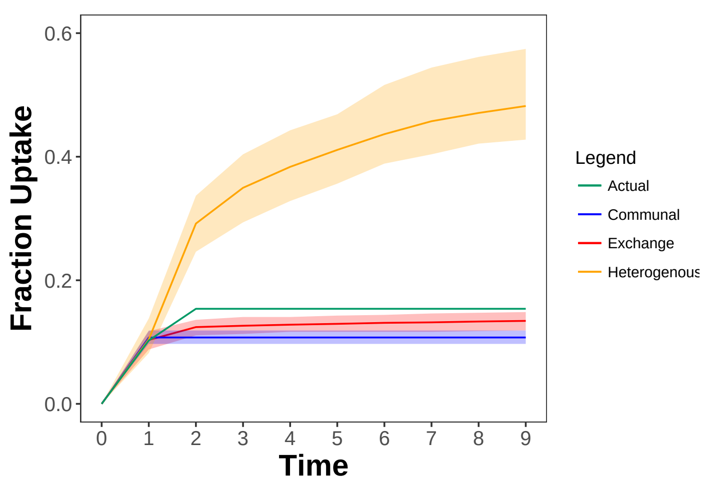

(a) Adoption over time

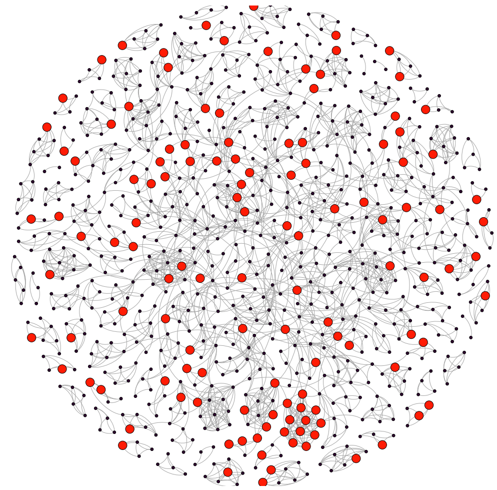

(b) Exchange Subgraph

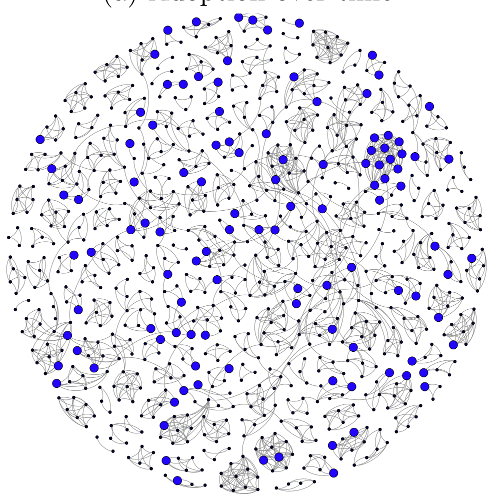

(c) Communal Subgraph

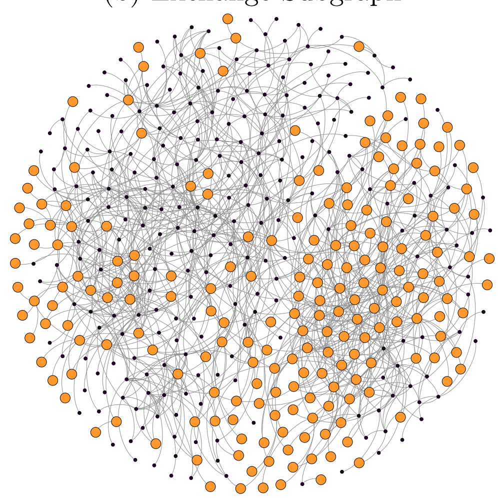

(d) Heterogeneous Subgraph

Figure 2: Adoption S-Curves in Subgraphs: Village 2

Table 1: Subgraph Properties and Diffusion

Subgraph Type	homogeneous		heterogeneous		difference		paired t-test (two-tailed)
variable	m1	se	m2	se	m1-m2	se	t-statistic
connectedness	0.51	0.01	0.93	0.01	-0.42	0.01	48.26
clustering	0.78	0.01	0.10	0.00	0.67	0.01	76.47
density	0.01	0.00	0.01	0.00	0.00	0.00	45.81
adoption rate	0.34	0.01	0.47	0.01	-0.13	0.01	22.38

of subgraphs across all villages in the sample was different: adoption in the homogeneous subgraphs averaged 34 percent, compared with 47 percent in the heterogeneous subgraphs (Table 1, $|t|=22.3$, $p<0.001$).

An interesting result is illustrated in Figure 4 where we plot the fraction of adoption

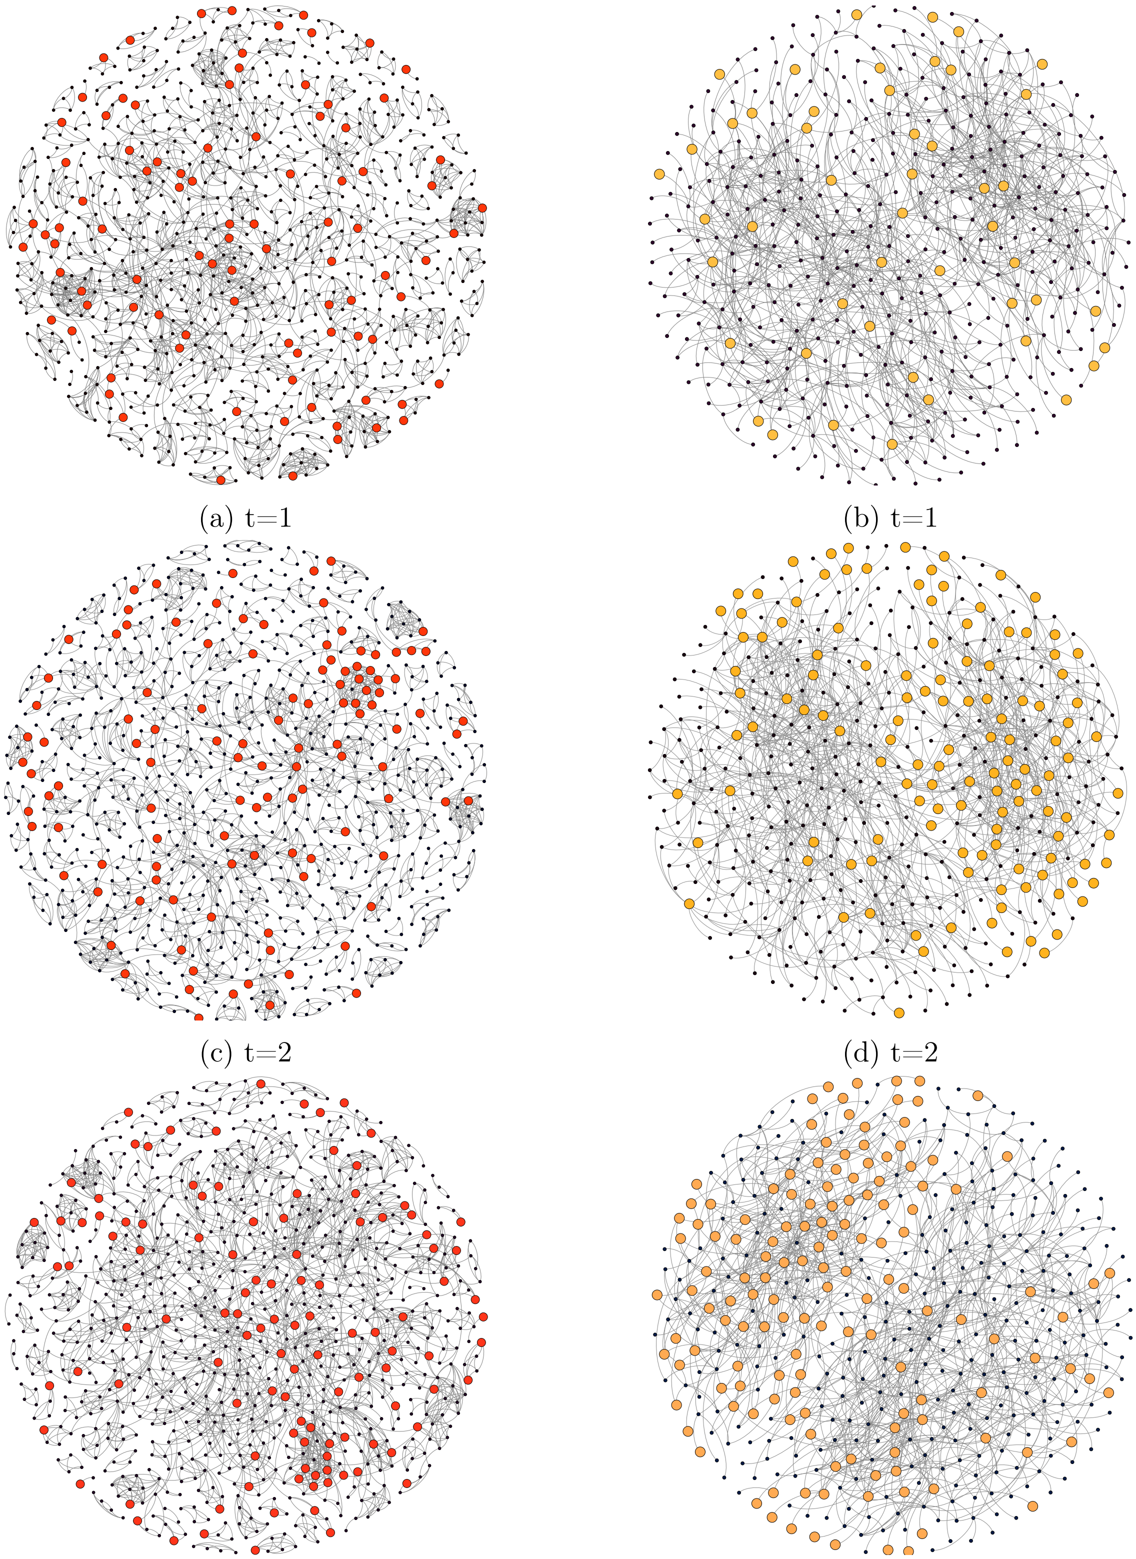

(a) t=1 (b) t=1

(c) t=2 (d) t=2

(e) t=3 (f) t=3

Figure 3: Temporal Dynamics Heterogeneous and Homogeneous Subgraphs: Village 2

against network properties of the subgraphs across all villages. The yellow squares, which represent the villages in Heterogeneous graphs, have observable differences from the properties of the Exchange and Communal subgraphs of these villages. Specifically, they have high connectedness and edge density while having extremely low clustering coefficients compared to those of the Exchange and Communal subgraphs. The higher adoption in Exchange subgraphs can also be explained in terms of these plots: these subgraphs tended to be more connected with higher edge densities and lower clustering coefficients compared to the Communal subgraphs.

Turning first to the subgraph properties, we posited that the connectedness, clustering, and density of the homogeneous and heterogeneous subgraphs would be different. Indeed, the homogeneous subgraphs had about half the average connectedness of the heterogeneous subgraphs ($|t|=48.2$, $p<0.001$), about seven times higher clustering ($|t|=76.4$, $p<0.001$), and about half the density of the heterogeneous subgraphs ($|t|=45.8$, $p<0.001$). These results suggest that there were multiple S-curves and information cascades unfolding within the villages, through different layers of the social structure. These layers also appear to have had very different structural properties associated with diffusion, such as clustering, connectedness, and density.

8 Conclusions

In this article, we examined diffusion dynamics on multiplex (multi-layered) networks by conducting simulations and structural analyses on the networks collected as a part of a prior study of the diffusion of microfinance in India (Banerjee et al., 2013b,a). Our findings revealed that there were multiple information cascades and S-curves of adoption unfolding on these networks. We also found that dynamics of adoption within different "layers" of the multiplex networks in each village were different. Specifically, heterogeneous layers – comprising ties that spanned both economic and social domains of interaction – were associ-

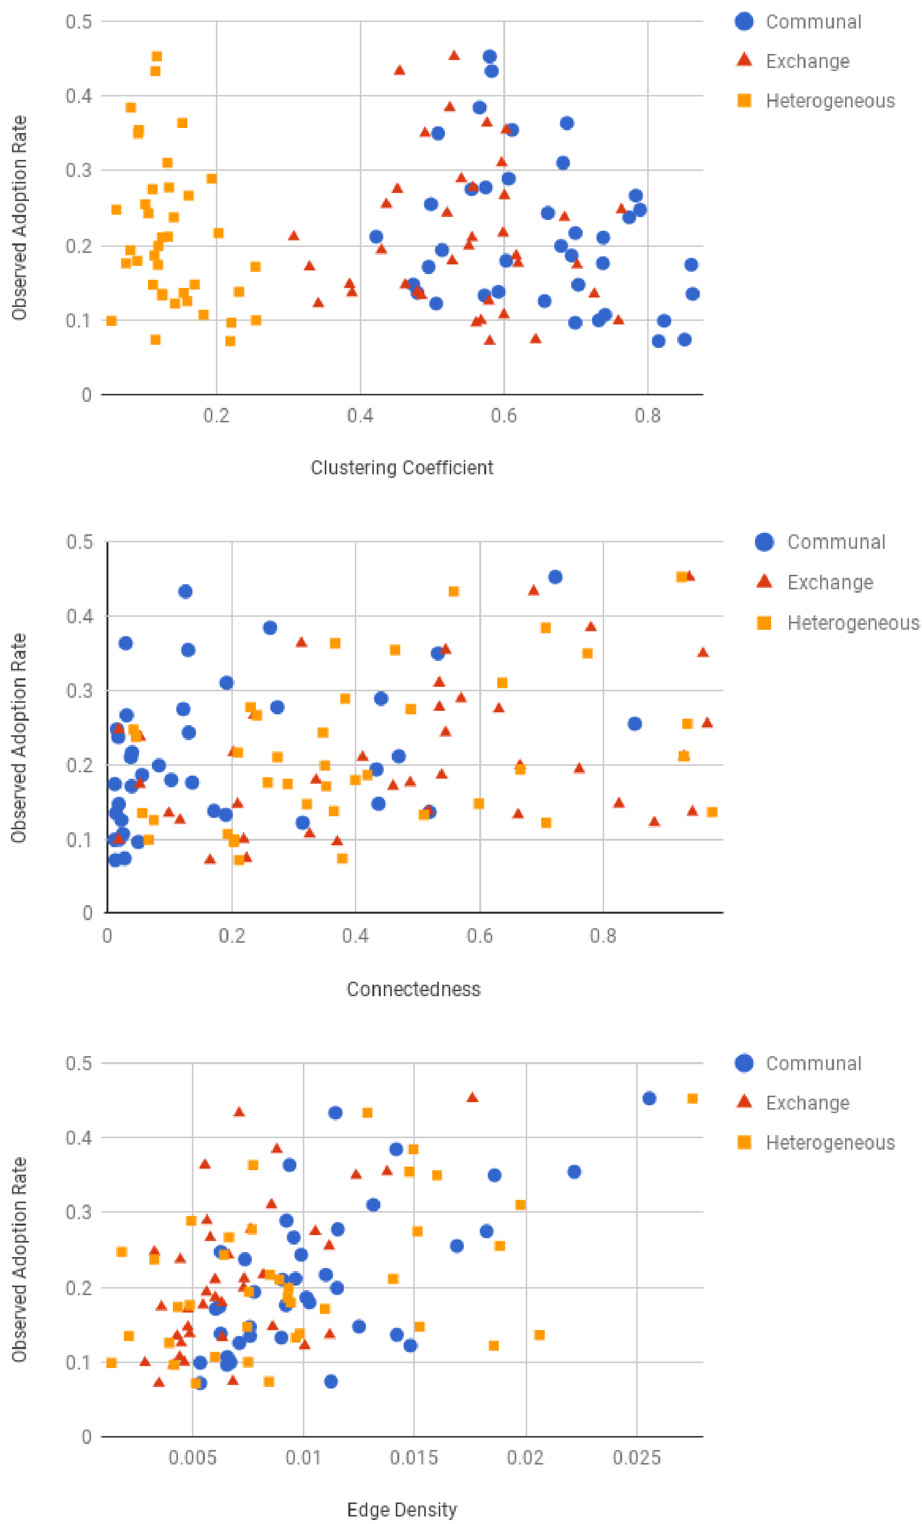

Figure 4: Observed Adoption and Subgraph Properties

ated with faster and broader diffusion than homogeneous layers. Further, these component layers exhibited different structural properties associated with diffusion, such as connectedness, clustering, and density. This evidence suggests that homogeneous and heterogeneous layers of the network were associated with different information bridging mechanisms and informational non-redundancy.

Our findings contribute to three distinct areas of research: network multiplexity (Brummitt et al., 2012; Kim and Goh, 2013; Smith and Papachristos, 2016), social diffusion processes on networks (Centola, 2015; Rossman, 2015; Banerjee et al., 2013a, 2018) and information transmission in networks (Burt, 1987; Aral et al., 2009; Brummitt et al., 2012; Alatas et al., 2015). First, in relation to studies of network multiplexity, we demonstrate that the structural homogeneity of multiplex networks, indicative of trust and social cohesion (Shipilov, 2012; Ferriani et al., 2013), is associated with lower connectedness and higher clustering than structural heterogeneity. Indeed, the two distinct types of layers present in multiplex graphs appear to have very different structural properties. Second, in relation to social diffusion processes, the results show that homogeneity in network layers appears to suppress – rather than promote – diffusion over time, owing to structural redundancy. Third, in relation to prior studies of the network properties conducive to information transmission and information cascades, our findings demonstrate that broad information diffusion is more likely to occur when multiplex networks comprise structurally heterogeneous (vs. homogeneous) layers. We believe that these results are attributable to how tie bridging across subgraphs of the network affect properties associated with faster and broader information spread, such as connectedness, clustering, and density.

These results demonstrate that different "layers" of multiplex networks indeed differ in their ability to induce virality and widespread contagion (Goel et al., 2016) by the types of information bridging that they provide. Specifically, heterogeneous layers bridge agents that do not otherwise come into contact and create linkages within the overall network that promote rapid diffusion (Banerjee et al., 2018), whereas homogeneous layers reinforce existing ties among agents through multi-layer closure. Thus, homogeneous and heterogeneous layers serve different functions, the former enabling trust and cohesion, and the latter enabling information diversity.

Our research suggests that many of the results previously associated with information bridging within networks (Brummitt et al., 2012; Aral and Alstyne, 2011) may not generalize to multiplex networks owing to these countervailing mechanisms. For instance, agents that are "central" in one layer but disconnected from another layer should be less effective than agents that are more sparsely – but more broadly – connected. Our results align with recent research examining information cascades on these networks, which shows that dynamical behavior in complex networks differs from that of simple networks (Myers and Leskovec, 2012; Melnik et al., 2013). Prior results that wide – or multiplex – ties promote diffusion (Centola and Macy, 2007; Centola, 2015) may therefore hold only for processes where multiple cascades are unlikely to arise.

Our model currently works on the assumption that information diffusion unfolds in fixed networks, which are not growing in the number of new actors or in the number of new ties added over time. Future research could extend this model to dynamic networks, in which new actors enter the network and ties are both formed and severed, to study how these dynamics affect information redundancy and diffusion. Empirical studies could also extend and apply these insights to a variety of processes that rely on information transmission through network ties, such as social sanctioning, gossip, and knowledge production. Finally, examinations of multiplexity as it relates to processes such as cooperation and competition in information acquisition present fruitful directions for future work.

References

Alatas, V., A. V. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. Olken (2015). Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia.

Angst, C. M., R. Agarwal, R. H. Smith, V. Sambamurthy, and K. Kelley (2010). Social Contagion and Information Technology Diffusion: The Adoption of Electronic Medical Records in U.S. Hospitals. *Management Science 56*(8), 1219–1241.

Aral, S. and M. V. Alstyne (2011). Networks, Information and Brokerage: The Diversity-Bandwidth Tradeoff. *American Journal of Sociology 117*(1), 90–171.

Aral, S., L. Muchnik, and A. Sundararajan (2009). Distinguishing Influence-Based Contagion from Homophily-Driven Diffusion in Dynamic Networks. *Proceedings of the National Academy of Sciences of the United States of America 106*(51), 21544–21549.

Assenova, V. A. (2018). Modeling the diffusion of complex innovations as a process of opinion formation through social networks. *PLoS ONE 13*(5), e0196699.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013a). Replication data for: "The Diffusion of Microfinance" https://hdl.handle.net/1902.1/21538, Havard Dataverse V9.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013b). The Diffusion of Microfinance. *Science 341*(6144), 1236498.

Banerjee, A. V., E. Breza, A. G. Chandrasekhar, and M. Mobius (2018). Naive Learning with Uninformed Agents.

Banerjee, A. V., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2018). Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials.

Bearman, P. S., J. Moody, and K. Stovel (2004). Chains of Affection: The Structure of Adolescent Romantic and Sexual Networks. *American Journal of Sociology 110*(1), 44–91.

Brummitt, C., K.-M. Lee, and K.-I. Goh (2012). Multiplexity-Facilitated Cascades in Networks. *Physical Review E 85*(4), 045102(R).

Burt, R. S. (1987). Social Contagion and Innovation: Cohesion versus Structural Equivalence. *American Journal of Sociology 92*(6), 1287.

Burt, R. S. (1992). *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Burt, R. S. (2012). Network-Related Personality and the Agency Question: Multirole Evidence from a Virtual World. *American Journal of Sociology 118*(3), 543–591.

Butts, C. T. (2008). Social Network Analysis with sna. *Journal of Statistical Software 24*(6), 1–51.

Cardillo, A., J. Gómez-Gardeñes, M. Zanin, M. Romance, D. Papo, F. del Pozo, and S. Boccaletti (2013). Emergence of Network Features from Multiplexity. *Scientific Reports 3*, 1344.

Centola, D. (2015). The Social Origins of Networks and Diffusion. *American Journal of Sociology 120*(5), 1295–1338.

Centola, D. and M. Macy (2007). Complex Contagions and the Weakness of Long Ties. *American Journal of Sociology 113*(3), 702–734.

Chandrasekhar, A. G., H. Larreguy, and J. P. Xandri (2018). Testing Models of Social Learning on Networks: Evidence from Two Experiments.

Coleman, J. S., E. Katz, and H. Menzel (1966). *Medical Innovation: A Diffusion Study*, Volume 9. Indianapolis: Bobbs-Merrill Co.

Curran, F. C. (2015). Expanding Downward: Innovation, Diffusion, and State Policy Adoptions of Universal Preschool. *Education Policy Analysis Archives 23*(36), 1–26.

Ferriani, S., F. Fonti, and R. Corrado (2013). The Social and Economic Bases of Network Multiplexity: Exploring the Emergence of Multiplex Ties. *Strategic Organization 11*(1), 7–34.

Goel, S., A. Anderson, J. Hofman, and D. J. Watts (2016). The Structural Virality of Online Diffusion. *Management Science 62*(1), 180 – 196.

Gómez-Gardeñes, J., I. Reinares, A. Arenas, and L. M. Floría (2012). Evolution of Cooperation in Multiplex Networks. *Scientific Reports 2*, 1–6.

Gould, R. V. (1991). Multiple Networks and Mobilization in the Paris Commune, 1871. *American Sociological Review 56*(6), 716–729.

Kim, J. Y. and K.-I. Goh (2013). Coevolution and Correlated Multiplexity in Multiplex Networks. *Physical Review Letters 111*(5), 058702.

Melnik, S., J. A. Ward, J. P. Gleeson, and M. A. Porter (2013). Multi-stage complex contagions. *Chaos 23*(1).

Myers, S. A. and J. Leskovec (2012). Clash of the contagions: Cooperation and competition in information diffusion. In *Proceedings - IEEE International Conference on Data Mining, ICDM*, pp. 539–548.

Powell, W. W., K. W. Koput, and L. Smith-Doerr (1996). Interorganizational Collaboration and the Locus of Innovation: Networks of Learning in Biotechnology. *Administrative Science Quarterly 41*(1), 116–145.

Reagans, R. E. and E. W. Zuckerman (2008). Why knowledge does not equal power: The network redundancy trade-off. *Industrial and Corporate Change 17*(5), 903–944.

Rogers, E. M. (1983). *Diffusion of Innovations* (3rd ed. ed.). New York: The Free Press.

Rosenkopf, L. and P. Almeida (2003). Overcoming Local Search Through Alliances and Mobility. *Management Science 49*(6), 751–766.

Rossman, G. (2014). The Diffusion of the Legitimate and the Diffusion of Legitimacy. *Sociological Science 1*, 49–69.

Rossman, G. (2015). *Climbing the Charts: What Radio Airplay Tells Us about the Diffusion of Innovation*. Princeton, N.J.: Princeton University Press.

Ryall, M. D. and O. Sorenson (2007). Brokers and Competitive Advantage. *Management Science 53*(4), 566–583.

Salganik, M. J., P. S. Dodds, and D. J. Watts (2006). Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science 311*(5762), 854–856.

Shipilov, A. (2012). Strategic Multiplexity. *Strategic Organization 10*(3), 215–222.

Singh, J. (2005). Collaborative Networks as Determinants of Knowledge Diffusion Patterns. *Management Science 51*(5), 756–770.

Smith, C. M. and A. V. Papachristos (2016). Trust Thy Crooked Neighbor: Multiplexity in Chicago Organized Crime Networks. *American Sociological Review 81*(4), 644–667.

Strang, D. and S. A. Soule (1998). Diffusion in Organizations and Social Movements: From Hybrid Corn to Poison Pills. *Annual Review of Sociology 24*, 265–290.

Uzzi, B. (1996). The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect. *American Sociological Review 61*(4), 674–698.

Vedres, B. and D. Stark (2010). Structural Folds: Generative Disruption in Overlapping Groups. *American Journal of Sociology 115*(4), 1150–1190.

Wang, P., L. J. Zhang, X. J. Xu, and G. Xiao (2017). Heuristic strategies for persuader selection in contagions on complex networks. *PLoS ONE 12*(1).

Wejnert, B. (2002). Integrating Models of Diffusion of Innovations: A Conceptual Framework. *Annual Review of Sociology 28*, 297–326.

Yagan, O. and V. Gligor (2012). Analysis of complex contagions in random multiplex networks. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics 86*(3).

Young, H. P. (2009). Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning. *American Economic Review 99*(5), 1899–1924.

Young, H. P. (2011). The Dynamics of Social Innovation. *Proceedings of the National Academy of Sciences of the United States of America 108 Suppl*(SUPPL. 4), 21285–91.