Friends with Costs and Benefits: Community Formation with Myopic, Boundedly-Rational Actors

Naina Balepur[*]
nainab2@illinois.edu
University of Illinois
Urbana-Champaign
Urbana, IL, USA

Andy Lee[*]
andy2@illinois.edu
University of Illinois
Urbana-Champaign
Urbana, IL, USA

Hari Sundaram
hs1@illinois.edu
University of Illinois
Urbana-Champaign
Urbana, IL, USA

ABSTRACT

In this paper we address how complex social communities emerge from local decisions by individuals with limited attention and knowledge. This problem is critical; if we understand community formation mechanisms, it may be possible to intervene to improve social welfare. We propose an interpretable, novel model for attributed community formation driven by resource-bounded individuals' strategic, selfish behavior. In our stylized model, attributed individuals act strategically in two dimensions: attribute and network structure. Agents are endowed with limited attention, and communication costs limit the number of active connections. In each time step, each agent proposes a new friendship. Agents then accept proposals, decline proposals, or remove friends, consistent with their strategy to maximize payoff. We identify criteria (number of stable triads) for convergence to some community structure and prove that our community formation model converges to a stable network. Ablations justify the ecological validity of our model and show that each aspect of the model is essential. Our empirical results on a physical world microfinance community demonstrate excellent model fits compared to baseline models.

1 INTRODUCTION

How do complex social communities emerge from behaviors of attributed, selfish, resource-bounded individuals (*i.e.*, with limited attention and who lack global network knowledge)? Developing an ecologically valid community formation mechanism is a challenging problem. If we can do so, we can improve social welfare by intervening via network-specific policy recommendations. For example, improving social welfare is salient in communities where misinformation spreads [4]. While we understand mechanisms by which global network structure [8] and decentralized search [31] emerge from individual actions, we lack similar community formation mechanisms. Community *detection* has been extensively studied [22, 34, 35, 38, 41, 49], but there has been limited work explaining how communities form through individual action.

While non-strategic models [20, 1, 48, 28] may generate networks with similar structural properties to social networks, they are not explanatory models. Individual choice plays a crucial role in community formation. Individuals have preferences over the structure of these communities; some may prefer dense groups with high trust, others may prefer to connect to a variety of people, maximizing social capital [12, 27, 11, 10]. People may also prefer to connect to those who are similar or different from themselves.

We propose a model for *attributed* community formation, where communities emerge through individuals' strategic, selfish behavior. In our stylized model, attributed individuals strategize in two dimensions: who they want to be friends with (homophily vs. heterophily) and what local network structures they desire (highly embedded vs. social capital). We know from Schelling's seminal work [44, 45] that an individual's preference for homophilous interaction with neighbors can lead to highly segregated communities. Conversely, friendship diversity can provide structural advantages [27]. We also know that some individuals prefer highly embedded communities that foster trust [26], yet others seek social capital for the structural advantages [11]. Inspired by the work of Dunbar [19], individuals in our model can support only a limited number of active communications. By triadic closure [32], individuals largely rely on their friends to introduce them to potential connections with some connections resulting from random encounters.

In this work, attributed community formation proceeds as follows. First, each agent examines their neighbors' connections as candidates for friendship. One additional node is revealed at random as a candidate, which mimics friendships arising from serendipitous encounters. After every agent proposes to the utility-maximizing candidate, all agents examine incoming proposals for friendship. They then *myopically* choose to do one of three things: 1) accept a proposal, 2) remove an existing friendship, 3) do nothing — whichever maximizes their payoff. Since all agents have a constant endowment of resources, and communication with friends is costly, dropping a low utility connection may prove strategic. In this work, we consider a population-level mixing of strategies; thus, we can view this model as an evolutionary game. We then identify criteria for convergence to a community structure. Our empirical results on a large microfinance dataset [6] demonstrate excellent fits to the model. Our contributions are as follows:

Evolutionarily Stable Communities: We propose a model for community formation that is evolutionarily stable. We define a notion of stability over triads; in contrast prior work considers pairwise stability [37]. We show that for any finite network the number of stable triads is non-decreasing, and so our model must converge to a non-unique equilibrium state.

Identification of Essential Agent Properties: In this work, we identify and test necessary behavioral properties of agents. Our agents are resource-constrained, restricted to local network knowledge only, and have preferences over both network structure and neighbor attributes. Prior work does not consider these properties in conjunction. We demonstrate through ablation tests that these three ecologically valid behavioral assumptions are necessary

---
[*]Authors contributed equally to this research.

to explain the formation of complex community structures. We elaborate on these model properties in Section 4.

Interpretability: Our model parameters characterize the nature of the agents in the network in a way that is easy to understand. Prior work on community formation [16, 37] uses decision functions whose parameters are difficult to interpret and rely upon vast amounts of data. We identify mixing proportions over strategies which help us explain and predict the network outcome.

2 RELATED WORK

There are many social network formation models that focus on producing graphs with some subset of observed properties of real networks. Within this broader class of network formation models are models which focus on community formation.

Non-strategic Network Formation: Classic non-strategic network formation models include the Erdős–Rényi [20], Barabási–Albert [1], Watts-Strogatz [48], and the Stochastic Block Model [28]. The model in [43] merges cliques to form a network with strong community structure. While these models may generate networks which have similar structural properties to social networks, they are not plausibly explanatory models due to the lack of individual choice.

Strategic Network Formation: Network formation games [21, 3, 2, 30] have been extensively studied and efficiency bounds are often known. Utility functions in such games are usually interpretable and the behavior of agents understandable. Such network games include local formation games [21], geodisic utility games [30], and information passing games [30]. In the directed setting there is work in discrete choice models of social network formation [40]. However, agents in these models often have complete network information and no resource bounds. Agents also rarely consider both structure and attributes in their edge formation decisions.

Best Response Strategies: Agents taking best response strategies make optimal selfish decisions while making assumptions about others' strategies. Much of the work on best response for network formation [23, 18, 24] builds off the work of Bala and Goyal [5]. They present network formation as a non-cooperative game where agents' actions lead to equilibrium social networks with simple structures. We also present a non-cooperative game for network formation, but our agents do not make assumptions about others' strategies, and only observe the current state of the network.

Bounded Rationality: Bounded rationality, first detailed by Simon [47], has been studied in a variety of contexts. It is known that in general agents violate rationality axioms [39] and do not maximize expected utility [46]. Based on this prior work we consider agents who do not perfectly maximize expected utility. Our baseline models (detailed in [16] and [37]) also attempt to incorporate bounded rationality. In both, agents consider the formation of only one edge in each iteration; so while knowledge is indeed limited, the choice of this edge is not strategic.

Triangle Closing: Triadic closure is a commonly observed phenomena in social networks [27], and prior network formation models have captured this property. The clique merging model of [43] produces networks with large numbers of triangles. Random walk models such as the process described in [33] also generate networks with high triangle count, however this model is not attributed.

Attribute Assortativity: Many network formation models capture homophilic preferences, and this area has been surveyed [42]. Some [25, 9] follow empirical methods for estimating the homophilic preferences of agents in network formation. Others [17] follow an approach more similar to our own, focusing on thresholds of assortativity and the impact on the final network. However, we seek a process that also involves preferences over structure. In real world networks, some people may prefer tightly knit groups with high trust, whereas others may prefer to connect to a variety of people, maximizing social capital [12, 27, 11, 10].

3 PROBLEM STATEMENT

We model community formation as an evolutionary game amongst strategic, boundedly-rational agents. Consider a world with agent set V; each agent v is endowed with immutable attributes τ_v, and a bounded resource. In this game, each agent seeks to form friendships with other agents, which incur a communication cost. Thus, agents are constrained to only maintain at most κ connections. At any time t, there exists an undirected network $G_t = (V, E_t)$. Due to the resource constraint, each agent must strategically choose friendships to maximize their payoff.

Agents have latent preferences in d dimensions. We model this by a d dimensional strategy space S. As in ecological games, we assume that for each of the d dimensions, agents pick a pure strategy such as homophily or heterophily. We assume a population-level mixing of strategies π over the d dimensions. Thus, each agent v has a strategy $s_v \in S$ with each dimension sampled according to π.

Each agent v chooses actions to maximize their payoff. The payoff of each action is determined by v's strategy s_v and the state of the network G_t. After each agent simultaneously modifies their network by taking an action during time t, the resultant network is G_{t+1}. We ask: will a local edge-formation process with strategic, resource-bounded agents converge to a network with stable community properties at some time t^* ?

4 ESSENTIAL AGENT PROPERTIES

We determine that the following agent properties are essential to future work in ecologically-valid, strategic community formation:

Resource-Constrained Agents: Maintaining active connections is costly, and humans are cognitively constrained [19]. We introduce a limit on agents' social interactions by way of a degree constraint κ — agents can only maintain κ-many edges at a time. We implement this by considering each agent to have a social interaction budget of 1, where each edge has cost $c = 1/\kappa$. Thus, we require that the following inequality must be true for all agents $v \in V$: $c * \delta(v) \leq 1$, where $\delta(v)$ is the degree of agent v. Agents assign no inherent utility to any "leftover" budget.

Local Network Knowledge: In social networks, agents do not know global network properties (*e.g.,* the total number of agents, the distance between themselves and distant agents), so we ensure that they only have knowledge of local structure. An agent may propose to form an edge with another agent if there is a path of length two between them, or if they were introduced by chance. We adopt this restriction because it represents triadic closure [32], an agent's connection introducing them to their neighbors.

Core Strategy Space: People consider many factors when choosing friends. However, prior work suggests that two factors: homophily (befriending similar others) [13, 14, 36, 32] and local structure (social capital, embeddedness) [11, 10, 26] are *both* essential strategic considerations. Thus, when forming friendships, our agents consider both the attribute of the friend, and the local network structure induced by the friendship.

We show through ablations that the properties detailed above are essential to our model. The expressive and diverse communities our model produces are not possible without each of these elements.

5 COMMUNITY FORMATION MODEL

In this section we give a detailed model description. We begin with the strategy space and the strategies' utility functions. We give a high level description of the network formation process as well as a pseudo-code description. Finally we give the stopping criteria for our model and illustrate its validity. We introduce considerable notation in this section which is summarized in Table 1.

Table 1: Summary of Notation

Notation	Description
S	d dimensional strategy space
π	Population-level mixing distribution over S
$s_v \sim \pi$	Agent v's private strategy drawn using π
τ_v	Immutable attribute vector of agent v
$\delta(v)$	Degree of agent v
κ	Degree constraint for all agents
$N(v)$	Set of nodes in the neighborhood of v
$N^\tau(v)$	Set of nodes in $N(v)$ with the same type as v
U_v	Utility agent v derives from $N(v)$
Δ_v	Number of triangles v is part of
I_v	Number independent nodes in $N(v)$
$d_G(u,v)$	Distance between nodes u and v in graph G
P_v^R	Set of nodes who have proposed to v

We consider population-level strategy mixing. We assume all agents are present at $t = 0$ (no addition or deletion of nodes). For $|V|$-many agents, the network is parameterized by κ, π, Ω. We define κ as the maximum degree of any agent.

In this paper we consider two dimensions: attributes and local structure; and exactly two pure strategies in each dimension: homophily (befriend similar), heterophily (befriend dissimilar) for attributes and social capital (maximize the breadth of friends), embeddedness (maximize the number of mutual friends) for local structure. Since our strategy space has two dimensions ($d = 2$), the population mixing distribution has two parameters, $\pi = (\alpha, \beta)$.

We define α as the proportion of agents adopting homophily, $1-\alpha$ as heterophily, β as the proportion of agents adopting social capital, $1 - \beta$ as embeddedness, and Ω as the distribution over immutable attributes. These parameters α, β, Ω are independent. Each agent $v \in V$ is assigned a strategy s_v and type τ_v at $t = 0$ according to parameters α, β, and Ω. In subsequent iterations, each agent v uses this strategy s_v as well as the current state of the network G_t to make changes to their neighborhood. The type of each agent τ_v

Table 2: Utility functions U_v^a and U_v^s for each strategy

Strategy	Utility Function	Description
H_m	$|N^\tau(v)|/\kappa$	Neighbors with attribute τ_v
H_r	$|N(v) \setminus N^\tau(v)|/\kappa$	Neighbors without attribute τ_v
L_e	$\Delta_v/\binom{\kappa}{2}$	Triangles where v is a vertex
L_c	I_v/κ	Independent neighbors of v

is public, but the strategy s_v is private. For tractability we assume that there are only two categorical values of τ.

5.1 Strategy Space

We denote the set of attribute strategies H and structural strategies L. Each agent has strategy $s_v = (s_v^a, s_v^s)$ drawn from $H \times L$. The homophilic strategy is denoted H_m and the heterophilic strategy H_r. Homophilic agent v derives attribute utility from neighbors of the same type as v. Heterophilic agent v derives attribute utility from neighbors of a different type. Let $N^\tau(v) = \{u | u \in N(v), \tau_u = \tau_v\}$, the set of v's neighbors who have the same type as v. Then $N(v) \setminus N^\tau(v)$ gives the set of neighbors of v of a different type.

We denote the strategy of embedded agents L_e and that of social capital agents L_c. Agent v desiring embeddedness derives structural utility from a connection to u if the edge between v and u is part of one or more triangles in the network. An agent v desiring social capital gets structural utility from neighbors who are disconnected. Let $\Delta_v = |\{(u, w)|u, w \in N(v), (u, w) \in E\}|$, the number of triangles where v is a vertex. Additionally let $I_v = |\{u|u \in N(v), \nexists w \text{ s.t. } w \in N(v), (u, w) \in E\}|$, the number of components of size 1 in the subgraph induced by v's neighborhood $N(v)$ on G.

We give detailed utility functions in Table 2. Agent v's strategy s_v is drawn from $\{H_m, H_r\} \times \{L_e, L_c\}$. The aggregate utility of v is the sum of the attribute and structural utility, $U_v = U_v^a + U_v^s$. We normalize attribute utilities and social capital utility by κ, the degree constraint for all nodes, and the embedded utility by $\binom{\kappa}{2}$ to ensure they all remain in $[0, 1]$. Thus, $U_v \in [0, 2]$, for all agents v.

5.2 Model Description

Consider the network at iteration t, $G_t = (V, E_t)$, with agents V and edges E_t. At time G_0 we have a trivial network, $E_0 = \emptyset$. We drop the t subscript where the meaning is clear. Each iteration consists of two stages: first the proposal stage, and then the action stage. Agents within each stage move simultaneously. We give the algorithm details below, followed by a simplified algorithm. We visually depict the algorithm in Figure 1.

1) Proposal Stage: Let $d_G(u, v)$ be the distance between u and v in G. For each agent $v \in V$, an agent $u \in V \setminus \{v\}$ is revealed to v uniformly at random, representing a chance encounter. Then v may choose to propose to one agent in $\{u\} \cup \{w | d_G(v, w) = 2\}$. We will denote the degree of agent v as $\delta(v)$. Agents may only propose if their degree is less than the constraint; or $\delta(v) < \kappa$. Because v may propose to at most one agent, they will propose to the agent giving the maximum utility.
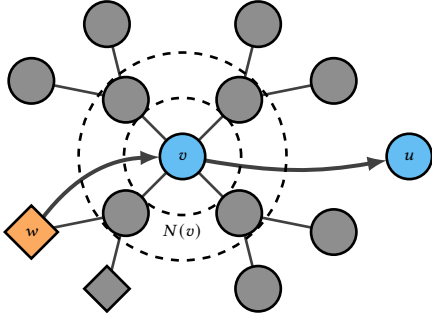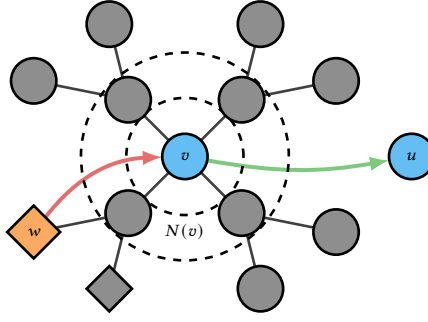
Figure 1a: Proposal stage at time t
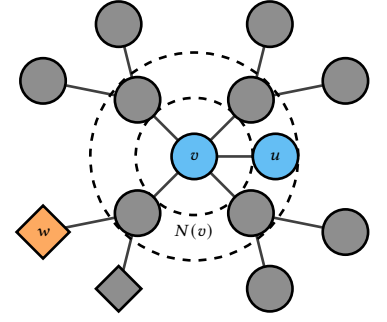
Figure 1b: Action stage at time t

Figure 1c: Network at time $t + 1$

Figure 1: Depiction of our proposal stage at time t, action stage at time t, and the network at time $t + 1$ following these stages. Node shape represents agent attribute τ_v; color is for attribute emphasis for relevant nodes. Each agent makes a proposal and responds with an action at time t; we omit most of these for clarity. At time t, agent v proposes to u, the random revelation. Agent w proposes to v. During the action stage, agent v does not add (w, v) while u agrees to add (u, v). Thus u is now in $N(v)$.

Algorithm 1: Proposal and Action Stages

Data: G the current network

Result: G' the network after actions

1 P an empty map from $V \rightarrow V$;

// Proposal Stage

2 **for** $v \in V$ **do**

3     $C \leftarrow$ the set of agents v may propose to;

4     $C_{\max} \leftarrow$ the agent in C that will maximize v's utility;

5     $U_{\max} \leftarrow$ v's utility change after addition of $(v, C_{\max})$;

6   **if** $U_{max} > 0$ and $\delta(v) < \kappa$ **then**

7        $P(v) \leftarrow C_{\max}$ // v will propose to $C_{\max}$

// Action Stage

8 $B_{\text{delete}}, B_{\text{add}} \leftarrow getBestActions(G, P)$ // get highest utility edge deletion and addition, assuming each proposal is accepted

9 **for** $v \in V$ **do**

10   **if** $\delta(v) < \kappa - 1$ and adding $(v, B_{add}(v))$ increases v's utility more than deleting $(v, B_{delete}(v))$ **then**

11   Add $(v, B_{\text{add}}(v))$ to E;

12   **else if** Deleting $(v, B_{delete}(v))$ increases v's utility **then**

13   Remove $(v, B_{\text{delete}}(v))$ from E;

---

2) Action Stage: The choice set for agent v comprises the nodes who proposed to v and current neighbors of v; potential actions are proposal acceptance, edge deletion, or no action. If an agent has degree $\delta(v) \geq \kappa - 1$ and has made a proposal, we will require they take a cost non-increasing action. Otherwise they will select whichever action maximizes their utility in this iteration, including doing nothing. If an agent has made a proposal, they assume their proposal has been accepted in the action stage. Each agent selects an action, and after every agent has selected, all agents execute their choices simultaneously. In a given iteration each agent may take one action.

5.3 Stopping Criteria

We define a triad of agents $u, v, w \in V$ as stable if the subgraph induced by u, v, w is of maximum utility for all agents; *i.e.*, no agent would choose to drop an edge. We stop our network formation simulations when the number of stable triads steadies.

Definition 5.1 (Stable triads). Let $G[u, v, w]$ be the subgraph of G induced by vertices $u, v, w \in V$. Agents u, v, w form a stable triad if and only if one of the following cases is true (see Figure 2):

(1) $s_u = s_v = s_w = (L_e, H_m)$ and $\tau_u = \tau_v = \tau_w$. The fully connected subgraph induced by u, v, w $G[u, v, w]$, is stable. A triple of agents with a preference for embeddedness and homophily of the same type achieve maximum utility from their connection.

(2) $s_u = s_v = s_w = (L_c, H_m)$ and $\tau_u = \tau_v = \tau_w$. Then $G[u, v, w]$ with $(u, v), (u, w) \in E, (v, w) \notin E$ is stable. A triple of homophilic agents with a preference for social capital who are on a line is also stable as v connecting to w would decrease the utility for all agents.

(3) $s_u = s_v = s_w = (L_c, H_r)$. Suppose $\tau_v = \tau_w$ and $\tau_u \neq \tau_v, \tau_w$. Then $G[u, v, w]$ with $(u, v), (u, w) \in E, (v, w) \notin E$ is stable. Similar to the prior case, a triple of heterophilic social capital agents on a line is also stable.
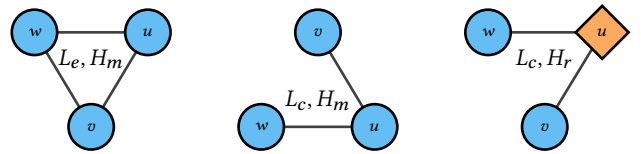

Figure 2: Three cases of stable triads— agents u, v, w.

THEOREM 5.2. *The number of stable triads in G is non-decreasing.*

PROOF. By definition, the subgraph induced by a stable triad contains nodes at maximum utility. Thus, alteration of the triad by edge addition or deletion would strictly harm all agents. Note that

4

for $s_u, s_v \in \{H_m\} \times L$ and $\tau_u = \tau_v$, if $(u,v) \in E$ then removing (u,v) decreases $|N^\tau(u)|$ by 1 for u and decreases $|N^\tau(v)|$ by 1 for v.

For the following cases we consider any stable triad u, v, w. When we consider alteration of one edge (*e.g.,* (u,v)), the proof extends to any equivalent edge (*e.g.,* (u,w)).

Case 1: $s_u = s_v = s_w = (L_e, H_m)$. Let us consider the deletion of (u,v). As $\tau_u = \tau_v$ this gives a reduction of $\frac{1}{\kappa}$ in U_u^a and U_v^a. Furthermore this must decrease Δ_u, Δ_v, and Δ_w.

Case 2: $s_u = s_v = s_w = (L_c, H_m)$ or $s_u = s_v = s_w = (L_c, H_r)$. The heterophily case is symmetric. Consider the deletion of (u,v). As $\tau_u = \tau_v$ this gives a $\frac{1}{\kappa}$ reduction in U_u^a and U_v^a.

Suppose there exists an agent x such that $(u,x),(v,x) \in E$. Deleting (u,v) increases the structural utility of u by at most $\frac{1}{\kappa}$ as x may now be an isolated neighbor of u. Because agents may make at most two connections in an iteration, if such a x were to exist then either (u,x) was added first, (v,x) was added first or (u,x) and (v,x) were added simultaneously. If WLOG (u,x) was added first then (v,x) would provide v with at most $\frac{1}{\kappa}$ attribute utility and reduce structural utility by the same. Then it must be that (u,x) and (v,x) are added simultaneously. However, because agents assume their proposals are accepted when evaluating candidates, there is no expected gain in U_u or U_v. Thus no such x may exist.

Next we consider the addition of (v,w). Because v and w are both connected to u, I_v and I_w would both be reduced by the addition of (v,w), reducing U_v^s, U_w^s, and U_u^s. □

As the number of stable triads in G is non-decreasing and G is finite, the number of stable triads must eventually converge. In simulation we consider the stable triad count to stabilize when standard deviation of the prior κ iterations is less than some small ε. We choose κ-many iterations because this would allow for each agent to completely change their neighborhood.

Though we do not prove this, we conjecture that convergence occurs in $O(\text{poly}(|V|, \kappa))$. Consider a network with countably infinite vertices. If $\kappa = 1$, it could not be the case that the network does not converge; in fact it would converge very quickly. Even for larger values of κ, maximum degree agents in stable triads are unlikely to continually switch between neighbors, as stable triad neighbors grant high payoff. By results from the coupon collector's problem, because of uniformly random revelations, agents will quickly find neighbors who grant them attribute utility. In practice our simulations do not require a large number of iterations.

6 SIMULATIONS

In this section we describe our first contribution: formation of evolutionarily stable communities. We explore the parameter space via simulation to demonstrate our model's ability to create networks with varied stable community structures.

In order to effectively explore the parameter space, we fix some parameters, while varying others. We use classical sociological results to fix values for $|V|$ and κ. From Dunbar [19], people can maintain active communities of ≈ 150; we fix $|V| = 150$ for our simulations. Dunbar also claims that a person's 5 strongest connections are their "support clique", or core group of emotional support. The closest 15 are the "sympathy group", who provide high cost support and time. We are interested in active relationships, though not necessarily restricted to only the closest connections. Thus

we fix $\kappa = 10$. We also fix Ω as a uniform distribution over two attributes for simplicity. We independently vary α and β over a 9×9 parameter space grid: $[0,1] \times [0,1]$. These parameters most directly control community structure and assortativity. To account for variation due to chance encounters in the proposal stage, we run each grid cell ten times and report the average metrics.

Triangle count depends strongly on chosen parameters. As expected, when a higher proportion of agents follow the embeddedness and homophily strategies, more triangles are produced (Figure 3a). Embedded agents' utility functions cause them to directly maximize triangle count. The correlation with homophilous preferences

(a) Triangle Count (hundreds)

(b) Assortativity Coefficient

(c) Stable Triad Count (hundreds)

(d) Average Utility of Agents

(e) Number of Louvain communities
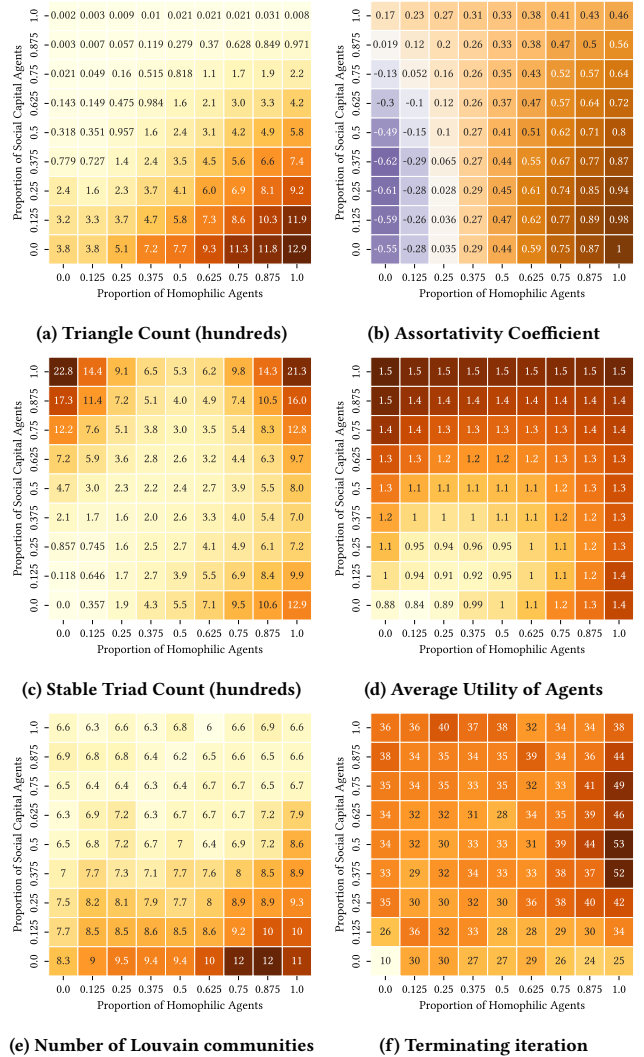
(f) Terminating iteration

Figure 3: Metrics over parameter space. Subfigure (b) shows that the proportion of social capital agents modulates the rate of change of the assortativity coefficient. Subfigures (c) and (d) show a bowl shape; stable communities form at the corners of our parameter space but not when embeddedness and heterophily are both high. The number of communities in (e) correlates strongly with the degree of embeddedness as agents desiring embeddedness form more neighborhoods that are closely knit.

Table 3: Ablation fits for two small villages with good standard model fit (top) and bad standard model fit (bottom).

Village 6 models	Best α	Best β	Loss
No resource constraint	0.2500	1.0000	0.1770
Global knowledge	0.3750	0.5000	0.0843
Ignore attribute	0.3750	0.5000	0.0658
Ignore structure	0.3750	0.5000	0.2689
Standard (no ablation)	0.1875	0.3750	**0.0419**

Village 10 models	Best α	Best β	Loss
No resource constraint	0.5000	1.0000	0.2643
Global knowledge	0.5000	0.5000	0.1863
Ignore attribute	0.2500	0.1250	**0.0866**
Ignore structure	0.3750	0.5000	0.2552
Standard (no ablation)	0.4375	0.2500	0.1228

is due to our use of a binary attribute — when two agents with heterophilous preferences connect on the basis of type, a third agent may give attribute utility to at most one of these agents.

Assortativity coefficient (Figure 3b) depends primarily on the level of homophily, which follows from the explicit optimization for assortativity and disassortativity by homophilous and heterophilous agents respectively. There is a weaker inverse correlation in the magnitude of the assortativity coefficient based on β. This may be due to agents preferring social capital deriving utility from any disconnected agent regardless of assortativity.

Due to the high variance in topology which results from different parameter values, our model can produce a wide variety of community structures. When agents prefer embeddedness the network forms cliques which are loosely connected. This roughly matches the idea of the strength of weak ties [27] wherein agents may form strong, tightly knit communities which are more loosely connected. Similarly, higher preference for homophily may create more tightly knit communities as cliques of uniform type may more easily form. This is reflected by the number of communities (Figure 3e) detected by the Louvain algorithm increasing with preference for embeddedness and homophily.
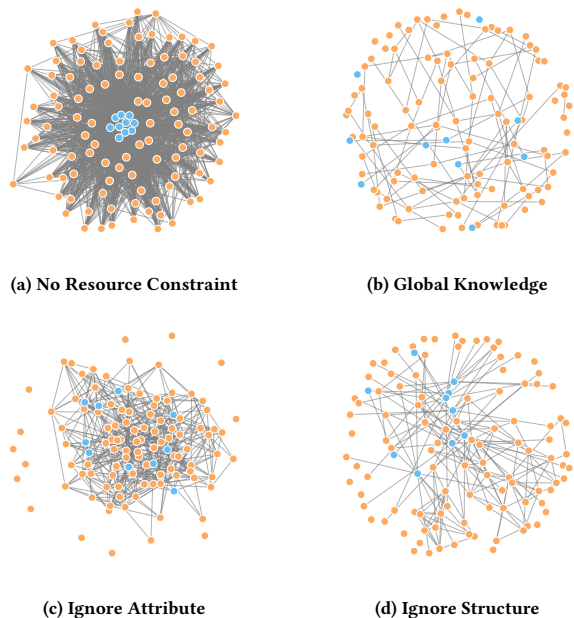
7 ABLATIONS

In this section we present results for our second contribution: identification of necessary agent behaviors. Our model relies upon boundedly-rational, resource constrained agents pursuing attribute and structural objectives to produce complex communities. Thus we utilize ablation tests to determine the necessity of each of these properties. Each ablation test removes only one property.

No resource constraint: Removal of resource constraint. Effectively, $\kappa = |V| - 1$; an agent can form as many edges as possible.

Global knowledge: Removal of local knowledge constraint. Agents are able to propose to anyone in the network, not just those within distance two and those who are revealed to them.

Ignore attribute: Removal of payoff from homophily and heterophily. Agents solely care about structural properties.

Ignore structure: Removal of payoff from social capital and embeddedness. Agents solely care about neighbors' attributes.

(a) No Resource Constraint   (b) Global Knowledge

(c) Ignore Attribute   (d) Ignore Structure

Figure 4: Ablation networks fit to Village 6. Color represents attribute. Notice the dense network in (a) due to lack of budget. Refer to Figure 5(a) and 5(b) for the actual network and our fit.

Our ablations reveal the strength of our ecologically valid model (see Table 3, ablation networks in Figure 4; original network and our model fit in Figure 5). The standard model generally outperforms the ablation models in both Village 6 (where our standard model achieves low loss) and Village 10 (where our standard model loss is higher). The standard model is outperformed only by the ablation where attribute is ignored for Village 10.

Here, the lower loss comes from the ablation model more accurately matching the distribution of assortativity coefficients. The better performance of the ablation is possibly due to the model's ability to produce agent neighborhoods with mixed attributes. Our standard model assigns each agent a pure strategy of homophily or heterophily, which makes achieving a mixed neighborhood difficult. Though agents in the global knowledge ablation have more information, it doesn't capture the fact that individuals in real networks are not omniscient, and likely make friends through mutual connections. Removing resource constraints produces a more dense network than real networks since individuals in real networks do not have the resources to befriend everyone. Similarly, ignoring structure does not produce sufficiently dense neighborhoods.

Further modeling subtleties may be introduced to increase our accuracy. A strict preference per agent for structure or attribute is likely too coarse; agents may exhibit different preferences over different subsets of the population. A vector of attributes would more accurately capture true attribute spaces with utility derived from vector similarity or dissimilarity. Agents could have homophilous or heterophilous preferences over different subsets of the attributes.

8 EXPERIMENTS

In this section we discuss our third contribution: interpretability. We evaluate our model on real network datasets and compare our

results to baselines. Unlike the baseline models, we are able to identify population-level mixing proportions over strategies which help us explain and predict the network outcome.

8.1 Dataset

We carefully chose a network dataset that best captured the intended use of our model. We require that a network be: attributed, undirected, and social, with active, voluntary edges. The village microfinance dataset collected by Banerjee et al. [7] meets our inclusion criteria. We use this dataset to demonstrate that our model can be fit to 75 unique social attributed networks.

The village network data comes from a survey of social networks in 75 Indian villages. The village networks are disjoint and the dataset contains demographic information for each household. We choose to observe the network at the level of the household rather than the individual as survey eligibility was determined by the presence of a head of household finances. Additionally this contraction avoids spurious homophily, heterophily, and triangles.

Banerjee et al. [6] provide twelve edge-sets for each of the 75 villages, corresponding to household survey questions (*e.g.,* who would you borrow money from, give advice to, etc.). Though the questions are directed, the authors provide undirected edges. We desire edges that encompass a wide range of social interactions, as our model is a general social network formation model. To ensure this, we consider the union over all edge-sets *minus* any edges that involve exchange of money or goods. We exclude monetary edges because these are not purely social interactions; they require some resources outside what is required for friendship.

We use the survey data to assign household attributes. We assume that number of rooms in a home is a proxy for wealth. We normalize this by dividing the number of rooms by the square root of number of beds + 1 [29]. We desire a binary attribute for wealth for each household, so we categorize households as having ≤ 2 normalized beds, or > 2. We choose this cutoff because it splits the households most evenly into two categories over all villages.

8.2 Baseline Models

Here we describe the baseline models; we attempted a best faith implementation of both. While Mele [37] provided code, neither it nor the required packages were documented enough to utilize.

Mele [37] This community formation model places attributed agents in unobserved communities at $t = 0$. In each iteration, a pair of agents is selected randomly with probability depending on network structure, community membership, and agents' attributes. Agents keep or create an edge if the sum of their utilities for the edge is positive. Some random variation is introduced to model unobserved strategies. The utility of an agent is dependent on the network structure, community membership, and agents' attributes. The paper provides six model variations; results show that one performs best, so we utilize this model for our analysis.

Christakis et. al. [16] This community formation model has agents that use a utility function parameterized by: a preference over attributes, for homophily, for the number of common connections, for distance in the network, and over the value of an observable characteristic. Agents have identical preferences.

One potential edge is introduced in each iteration; if the utility for the edge is positive for both agents then the edge is formed.

Both baselines, like our model, consider attribute and structure in formation of edges. In contrast to our own model, both baselines lack resource constraints and have hard to explain parameters.

8.3 Evaluation Metrics

We refer to an observed network as G^d, and use $\widehat{G}^d$ when referring to a simulated network which attempts to replicate G^d. Here we describe the metric used to evaluate this replication. For each vertex we calculate attribute assortativity a_i and local clustering coefficient c_i. The assortativity coefficient a_i indicates the level of homophily or heterophily exhibited by a vertex v, and is given by $a_i = \frac{1}{2}\left(\frac{2N^\tau(v) - \delta(v)}{\delta(v)} + 1\right)$. An assortativity coefficient of one indicates complete homophily and zero complete heterophily. The clustering coefficient of a vertex v is given by $\frac{\Delta_v}{\delta(v)(\delta(v)-1)}$. For both networks $\widehat{G}^d$ and G^d, we calculate a distribution of these pairs (a_i, c_i) and a weight p_i which corresponds to the relative frequency of (a_i, c_i) in each network. We utilize the 1-Wasserstein distance as a measure of similarity between distributions. The distance between points in the distribution is given by the L-2 (Euclidean) norm.

We also include a global component in our loss metric to ensure good fit on global assortativity and edge density. A network with many edges could be similar in distribution to one with few edges if they have similar local clustering coefficients. The global component alleviates this issue. Thus, we compute the difference between triangle count and assortativity coefficient by a symmetric mean absolute percentage error, and take the average of these errors. This makes up the global component. Note that very sparse networks have undefined clustering coefficients which we set to zero. We report the average of distributional and global components as the loss between the simulated $\widehat{G}^d$ and the observed network G^d.

8.4 Parameter Fitting

Given an attributed network dataset, some parameters are defined by the dataset. Our model has five free parameters: $|V|$, κ, Ω, α, and β. We look only at static datasets, so $|V|$ and Ω are fixed. The remaining three parameters need to be discovered.

We use grid search to discover the best fitting κ, α, and β for G^d. We consider 75 village networks, and we desire a shared κ over all villages, since κ represents a cognitive constraint on the ability of individuals to form edges. To achieve this, we run a coarse grid search to determine the best-fitting $\kappa \in \{5, 10, 15\}$. Then, using this κ, we find the best fitting α and β on a finer grid. The parameters in the baseline models [37, 16] are neither explanatory nor bounded. To fit them, we execute a random search over a plausible parameter space informed by the reported parameters in the papers and then a finer search over the discovered best-fitting space. We utilize our evaluation metrics (Section 8.3) to determine best fit.

8.5 Our Experimental Results

We now present the results of our model fitting, and compare to the baseline models. We executed all simulations five times, computed the losses from these networks, and then averaged the losses.

Table 4: Village Microfinance Fitting Results. We include the best-fitting (top half) and worst-fitting (bottom half) villages in order of increasing loss. We also include Village 10 for the baseline comparisons.

Village	$|V|$	Assort.	Tri.	Best α	Best β	Loss
4	239	-0.031	154	0.4375	0.5625	**0.027**
3	292	0.001	189	0.3750	0.5000	**0.037**
8	94	0.013	165	0.3125	0.2500	**0.041**
6	114	-0.017	107	0.1875	0.3750	**0.042**
2	195	0.102	113	0.3750	0.5625	**0.043**
10	77	0.041	152	0.2500	0.4375	**0.123**
75	172	0.123	661	0.3750	0.0000	**0.177**
70	205	0.048	860	0.3125	0.0000	**0.182**
52	327	0.040	1698	0.5625	0.0000	**0.226**
51	251	0.044	1440	0.5626	0.0000	**0.233**
69	180	0.077	1197	0.5000	0.0000	**0.247**

We show results from fitting our model to the microfinance dataset in Figure 4. We only show the five best-fitting and five worst-fitting villages, along with Village 10, due to space constraints. In our initial coarse-grained fitting, we found the optimal κ value was 10. Note that since we also ran our simulations with $\kappa = 10$, Figure 3 gives an idea of how metrics change with values of α and β.

Recall that α and β give the proportion of homophilic and social capital agents respectively. We expect that for villages with high assortativity, α will be high, and for villages with a high proportion of triangles compared to network size, β will be low. Reviewing Figures 3a and 3b reminds us that α and β do not cause independent variation in triangle count and assortativity, so we observe some seemingly "sub-optimal" values of α and β. For example, note the five worst-fit villages. They all have been fit with a β value of 0.000, indicating a high triangle count. Indeed, the triangle counts for these villages are high. However, the α value does not vary as predictably with assortativity, leading to higher loss values.

8.6 Baseline Models' Experimental Results

Table 5: Comparison to Baseline Model Fitting Results. We are unable to fit baselines to all villages due to long running times. We report mean (X) and standard deviation (Y) in the last row. (* statistics reported for 30 villages)

Village	Our Loss	Chris. [16] Loss	Mele [37] Loss
6	**0.042**	0.168	0.601
10	**0.123**	0.135	0.503
All $X(\pm Y)$	**0.110 ± 0.042**	0.139 ± 0.021*	–

We compare our model's fit to the baselines for a small village our model fit well (Village 6), and a small village our model fit less well (Village 10). We were unable to fit all villages because the running time of the baseline models was prohibitively long. To (locally) complete a search over 200 parameter sets to produce the networks in Figure 5, our simulation takes 45 minutes, Christakis's 40 minutes, and Mele's over 14 hours. We show results for Village 6 in Figure 5, and full loss results in Table 5. Though the Christakis

(a) Actual Network   (b) Our Simulation

(c) Christakis Simulation   (d) Mele Simulation
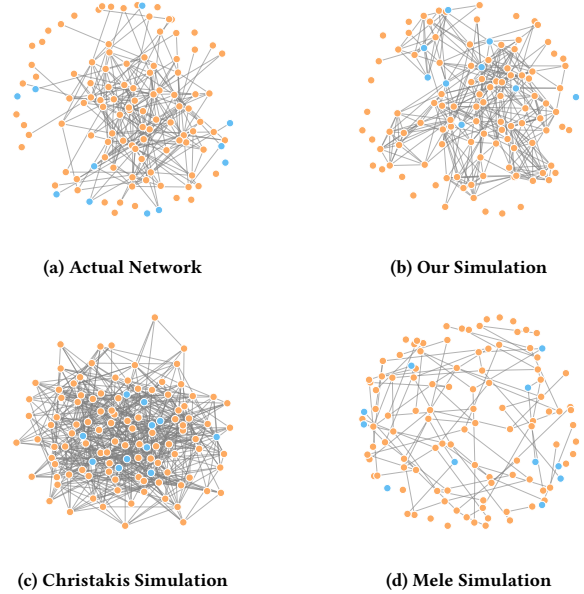
Figure 5: Comparison of actual network (Village 6) to our network and two baseline models. Color represents the attribute. Our model performs best, followed by Christakis [16] and then Mele [37].

model performs fairly well (Figure 5c), the resulting network is too dense and underestimates the number of isolated nodes in Figure 5a. On the other hand, the Mele model performs poorly, clearly underestimating the density and triangle count of the true network, while properly capturing the isolated nodes. Note that the Christakis model is parameterized by seven variables, and the Mele model by nine. It is difficult to make claims regarding populations of people using these parameter sets, even if they are well fit.

9 LIMITATIONS AND FUTURE WORK

Scalability: Agents in our model consider candidates two hops away; this can be as many as κ^2 candidates. Scalability of the model could be improved by restricting this candidate set. An implementation optimizing for sparse matrix operations could also improve computation time, especially on very large datasets.

Attribute dimension: We only consider agents with a single scalar categorical attribute. This could be resolved by increasing the dimension of attribute preferences as noted in Section 7.

Strategy space: Our definition of the social capital structural utility function (requiring completely disconnected neighbors) may be overly strict, and make it difficult for social capital agents to achieve high utility. Instead, a modification where utility is derived per connected component in the neighborhood induced subgraph may more accurately capture real payoffs.

Convergence rate: We stop our network formation simulations when the stable triad count stabilizes (see Section 5.3). While we prove that this occurs in a finite number of iterations, we only conjecture a convergence rate. Determining when the probability of substantial structural change is low would be a stronger result.

10 CONCLUSION

We present an ecologically valid model of community formation with boundedly rational, attributed agents. The agents can maintain a limited number of friendships, and have local network knowledge. They employ core strategies related to local network structure and preference for homophily to form friendships. We consider a population-level mixing of strategies, including attribute and structural preferences. Every parameter value in our model is meaningful — a particular choice of α (homophily), β (social capital), ω (type distribution), and κ (resource constraint) is a statement on the preferences of the population. We prove that the number of stable triads in our network converges; thus, our network must become structurally stable. We show via simulation that this convergence occurs quickly in practice. Our model is well-fit to the observed networks, and ablations demonstrate ecological validity. Our model differs from the prior work in that agent behaviors are ecologically valid and interpretable. The significance of our ecologically valid model lies in creating counterfactual scenarios and in examining policy interventions such as those designed to embed people in different parts of a social network [15].

REFERENCES

[1] Réka Albert and Albert-László Barabási. "Statistical mechanics of complex networks". In: *Reviews of modern physics* 74 (2002), p. 47.

[2] E Anshelevich et al. "The price of stability for network design with fair cost allocation". In: *SIAM Journal on Computing* 38.4 (2008), pp. 1602–1623.

[3] Elliot Anshelevich, F Shepherd, and Gordon Wilfong. "Strategic Network Formation through Peering and Service Agreements". In: *FOCS* 6 (2006).

[4] Joseph B. Bak-Coleman et al. "Stewardship of global collective behavior". In: *Proceedings of the National Academy of Sciences* 118.27 (2021). ISSN: 0027-8424. DOI: 10.1073/pnas.2025764118. eprint: https://www.pnas.org/content/118/27/e2025764118.full.pdf. URL: https://www.pnas.org/content/118/27/e2025764118.

[5] Venkatesh Bala and Sanjeev Goyal. "A noncooperative model of network formation". In: *Econometrica* 68.5 (2000), pp. 1181–1229.

[6] Abhijit Banerjee et al. *The Diffusion of Microfinance*. Version V9. 2013. DOI: 10.7910/DVN/U3BIHX. URL: https://doi.org/10.7910/DVN/U3BIHX.

[7] Abhijit Banerjee et al. "The diffusion of microfinance". In: *Science* 341.6144 (2013), p. 1236498.

[8] Albert-László Barabási and Réka Albert. "Emergence of Scaling in Random Networks." In: *Science* 286.5439 (1999), p. 509. ISSN: 00368075. URL: http://search.ebscohost.com/login.aspx?direct=true&db=tfh&AN=2405932&site=ehost-live.

[9] Vincent Boucher. "STRUCTURAL HOMOPHILY". In: *International Economic Review* 56.1 (Jan. 2015), pp. 235–264. DOI: 10.1111/iere.12101.

[10] Ronald S Burt. "Structural holes and good ideas". In: *American journal of sociology* 110.2 (2004), pp. 349–399.

[11] Ronald S. Burt. "The Social Capital of Opinion Leaders". English. In: *Annals of the American Academy of Political and Social Science* 566 (1999), pp. 37–54. ISSN: 00027162. URL: http://www.jstor.org/stable/1048841.

[12] S. W. Campbell. "Mobile communication and network privatism: A literature review of the implications for diverse, weak, and new ties". In: *Review of Communication Research* 3.1 (2015), pp. 1–21.

[13] Damon Centola. "An Experimental Study of Homophily in the Adoption of Health Behavior". In: *Science* 334.6060 (2011), pp. 1269–1272. ISSN: 0036-8075. DOI: 10.1126/science.1207055. eprint: https://science.sciencemag.org/content/334/6060/1269.full.pdf. URL: https://science.sciencemag.org/content/334/6060/1269.

[14] Damon M Centola. "Homophily, networks, and critical mass: Solving the start-up problem in large group collective action". In: *Rationality and Society* 25.1 (2013), pp. 3–40.

[15] Raj Chetty, Nathaniel Hendren, and Lawrence F Katz. "The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment". In: *American Economic Review* 106.4 (2016), pp. 855–902.

[16] Nicholas Christakis et al. "An empirical model for strategic network formation". In: *The Econometric Analysis of Network Data*. Elsevier, 2020, pp. 123–148.

[17] Sergio Currarini, Jesse Matheson, and Fernando Vega-Redondo. "A simple model of homophily in social networks". In: *European Economic Review* 90 (Nov. 2016), pp. 18–39. DOI: 10.1016/j.euroecorev.2016.03.011.

[18] Jean Derks et al. "Local dynamics in network formation". In: *Proc. Third World Congress of The Game Theory Society*. Citeseer. 2008.

[19] R.I.M. Dunbar. "The social brain hypothesis and its implications for social evolution". In: *Annals of Human Biology* 36.5 (Jan. 2009), pp. 562–572. DOI: 10.1080/03014460902960289.

[20] Paul Erdős and Alfréd Rényi. "On the evolution of random graphs". In: *Publ. Math. Inst. Hung. Acad. Sci* 5 (1960), pp. 17–60.

[21] Alex Fabrikant et al. "On a network creation game". In: *Proceedings of the twenty-second annual symposium on Principles of distributed computing - PODC '03*. ACM Press, 2003. DOI: 10.1145/872035.872088.

[22] Santo Fortunato. "Community detection in graphs". In: *Physics Reports* 486.3-5 (2010), pp. 75–174. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2009.11.002.

[23] Tobias Friedrich et al. "Efficient best response computation for strategic network formation under attack". In: *Algorithmic Game Theory: 10th International Symposium, SAGT 2017, L'Aquila, Italy, September 12–14, 2017, Proceedings 10*. Springer. 2017, pp. 199–211.

[24] Jacob K Goeree, Arno Riedl, and Aljaž Ule. "In search of stars: Network formation among heterogeneous agents". In: *Games and Economic Behavior* 67.2 (2009), pp. 445–466.

[25] Bryan S. Graham. "An Econometric Model of Network Formation With Degree Heterogeneity". In: *Econometrica* 85.4 (2017), pp. 1033–1063. DOI: 10.3982/ecta12679.

[26] Mark Granovetter. "Economic Action and Social Structure: The Problem of Embeddedness". English. In: *American Journal of Sociology* 91.3 (1985), pp. 481–510. ISSN: 00029602. URL: http://www.jstor.org/stable/2780199.

[27] Mark S. Granovetter. "The Strength of Weak Ties". In: *American Journal of Sociology* 78.6 (May 1973), pp. 1360–1380. DOI: 10.1086/225469.

[28] Paul W; Holland et al. "Stochastic blockmodels: First steps". In: *Social Networks* 5.2 (1983), pp. 2021–2026. DOI: 10.1016/0378-8733(83)90021-7.ISSN0378-8733.

[29] Mohammad Azhar Hussain. "The sensitivity of income polarization". In: *The Journal of Economic Inequality* 7.3 (Jan. 2008), pp. 207–223. DOI: 10.1007/s10888-007-9072-x.

[30] M Jackson. "A survey of network formation models: stability and efficiency". In: *Group formation in economics: Networks, clubs, and coalitions.* Vol. 664. 2005, pp. 11–49.

[31] Jon M. Kleinberg. "Navigation in a small world". In: *Nature* 406 (Aug. 2000), 845 EP -. URL: http://dx.doi.org/10.1038/35022643.

[32] Gueorgi Kossinets and Duncan J. Watts. "Origins of Homophily in an Evolving Social Network". In: *American Journal of Sociology* 115 (2009), pp. 405–450. URL: http://www.journals.uchicago.edu/doi/abs/10.1086/599247.

[33] J.M. Kumpula et al. "Model of community emergence in weighted social networks". In: *Computer Physics Communications* 180.4 (Apr. 2009), pp. 517–522. DOI: 10.1016/j.cpc.2008.12.016.

[34] Jure Leskovec et al. "Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters". In: *Internet Mathematics* 6.1 (2009), pp. 29–123.

[35] Yu-Ru Lin et al. "Community Discovery via Metagraph Factorization". In: *ACM Trans. Knowl. Discov. Data* 5.3 (Aug. 2011), 17:1–17:44. ISSN: 1556-4681. DOI: 10.1145/1993077.1993081. URL: http://doi.acm.org.proxy2.library.illinois.edu/10.1145/1993077.1993081.

[36] Miller McPherson, Lynn Smith-Lovin, and James M. Cook. "Birds of a Feather: Homophily in Social Networks". English. In: *Annual Review of Sociology* 27 (2001), pp. 415–444. ISSN: 03600572. URL: http://www.jstor.org/stable/2678628.

[37] Angelo Mele. "A structural model of homophily and clustering in social networks". In: *Journal of Business & Economic Statistics* 40.3 (2022), pp. 1377–1389.

[38] M. E. J. Newman and M. Girvan. "Finding and evaluating community structure in networks". In: *Phys. Rev. E* 69 (2 Feb. 2004), p. 026113. DOI: 10.1103/PhysRevE.69.026113. URL: http://link.aps.org/doi/10.1103/PhysRevE.69.026113.

[39] Kirby Nielsen and John Rehbeck. "When Choices Are Mistakes". In: *American Economic Review* 112.7 (July 2022), pp. 2237–2268. DOI: 10.1257/aer.20201550.

[40] Jan Overgoor, Austin Benson, and Johan Ugander. "Choosing to Grow a Graph: Modeling Network Formation as Discrete Choice". In: *The World Wide Web Conference on - WWW '19.* ACM Press, 2019. DOI: 10.1145/3308558.3313662.

[41] Gergely Palla et al. "Uncovering the overlapping community structure of complex networks in nature and society". In: *Nature* 435.7043 (2005), pp. 814–818.

[42] M Rivera, S Soderstrom, and B Uzzi. "Dynamics of dyads in social networks: Assortative, relational, and proximity mechanisms". In: *annual Review of Sociology* 36 (2010), pp. 91–115.

[43] Arnaud Sallaberry, Faraz Zaidi, and Guy Melançon. "Model for generating artificial social networks having community structures with small-world and scale-free properties". In: *Social Network Analysis and Mining* 3.3 (2013), pp. 597–609.

[44] Thomas C. Schelling. "Dynamic Models of Segregation". In: *Journal of Mathematical Sociology* 1.2 (1971), pp. 143–186. ISSN: 0022250X. URL: http://login.ezproxy1.lib.asu.edu/login?url=http://search.ebscohost.com.ezproxy1.lib.asu.edu/login.aspx?direct=true&db=sih&AN=13985129&site=ehost-live.

[45] Thomas C. Schelling. *Micromotives and Macrobehavior.* New York and London: WW Norton, 1978.

[46] R Selten. "Bounded rationality". In: *Journal of Institutional and Theoretical Economics* 146.4 (1990), pp. 649–658.

[47] Herbert A. Simon. "A Behavioral Model of Rational Choice". In: *The Quarterly Journal of Economics* 69.1 (Feb. 1955), p. 99. DOI: 10.2307/1884852.

[48] Duncan Watts and Steven Strogatz. "Collective dynamics of 'small-world'networks". In: *nature* 393 (1998), pp. 440–442.

[49] Jaewon Yang, Julian McAuley, and Jure Leskovec. "Community detection in networks with node attributes". In: *Data Mining (ICDM), 2013 IEEE 13th international conference on.* IEEE. 2013, pp. 1151–1156.