What Keeps a Vibrant Population Together?

Jayati Deshmukh
Srinath Srinivasa
Sridhar Mandyam

International Institute of Information Technology Bangalore
Bangalore, Karnataka, India

Managing diversity is a challenging problem for organizations and governments. Diversity in a population may be of two kinds—acquired and innate. The former refers to diversity acquired by pre-existing social or organizational environments, attracting employees or immigrants because of their wealth and opportunities. Innate diversity, on the other hand, refers to a collection of pre-existing communities having to interact with one another and to build an overarching social or organizational identity. While acquired diversity has a prior element of common identity, innate diversity needs to build a common identity from a number of disparate regional or local identities. Diversity in any large population may have different extents of acquired and innate elements. In this paper, innate and acquired diversity are modeled in terms of two factors, namely: insularity and homophily, respectively. Insularity is the tendency of agents to act cooperatively only with others from the same community, which is often the primary challenge of innate diversity; while homophily is the tendency of agents to prefer members from their own community to start new social or business connections, which is often the primary challenge in acquired diversity. The emergence of network structure is studied when insularity and homophily are varied. In order to promote cooperation in a diverse population, the role played by a subset of agents called "global" agents who are not affected by homophily and insularity considerations is also studied. Simulation results show several interesting emergent properties. While the global agents are shown to acquire high betweenness, they are by no means the wealthiest or the most powerful in the network. However, the presence of global agents is important for the regional agents whose own wealth prospects increase because of their interaction with global agents.

Keywords: diversity; multi-agent systems; networks and prisoner's dilemma

1. Introduction

Diversity is an administrative ideal for building a vibrant society or organization. Diversity may be of several kinds, such as linguistic, gender, ethnicity or race. A diverse population is made up of several

https://doi.org/10.25088/ComplexSystems.30.3.347

subgroups or *local identities*, each of which represents its own unique perspective on several issues. As a result, the collective insight from a diverse population tends to be rich and nuanced, resulting in several forms of collective benefits [1–5].

Yet at the same time, diversity also brings its own challenges. Managing diverse subgroups greatly increases strategic and operational costs. In addition, disparate local identities often conflict with one another on several issues, contributing to increased social strife [1, 4, 6]. Against this backdrop, studying the dynamics of diverse populations is a problem of increasing relevance that is being tackled using computational modeling.

Diversity in a population can be of two extreme kinds: *acquired* and *innate*. Acquired diversity refers to a pre-existing social or organizational framework, which attracts a large immigrant population because of its wealth and opportunities. This is typically the case with rich, inclusive countries like the US, Canada, Australia and different countries of Europe, or with large multinational corporations. These countries and organizations had an existing sense of national or organizational structure, practices, identity and cultural framework in place before they became an attractive destination for incoming immigrants or employees. In such cases, despite the diversity, a common sense of prior identity permeates across the population that is a result of the immigrants' or employees' conscious decision to affiliate with the country or the organization.

However, immigrants in such an environment also often experience a loss of their own sense of former cultural identity over time. Their cultural imports often get reduced to exotic or esoteric expressions of mere ornamental value. This often leads immigrants to seek to associate and develop relationships with others of their own kind, so as to be able to express their own ethnic identity and worldview at deeper levels. This forms cultural subgroups and hyphenated identities like for example, African-Americans, Asian-Americans, Indian-Americans, etc. Such a tendency to seek out and associate with others who are similar to us is called *homophily*.

The other kind of diversity is innate diversity. Here, the population is comprised of several pre-existing communities with their local cultures, identities and practices that have existed for a long time, and that need to work together to build an overarching social, cultural, national or organizational framework. This kind of diversity is characteristic of the European Union trying to build an overarching identity across the disparate cultures of Europe; or the diverse nation of India, comprised of several regional cultures, languages and worldviews. In business settings, cooperative movements of farmers and small businesses, or trade unions across several small businesses, have to contend with innate diversity.

In such populations, local identities or "regionalism" tend to be much stronger than the global identity. Regional subgroups fight to preserve their local cultural identity and often overtly resist implementation of a common identity or worldview across the population. Activists who fight to preserve regional identities often call such a population the "rainbow" model of diversity [7] where different regional subcultures collectively form a rich, collective culture, yet retaining their distinct regional identities, without blending into one another.

The challenge with innate diversity is to create an overarching framework of unity or oneness that enables the disparate subgroups to trust and act cooperatively with one another. Strong regionalism leads to entrenchment and distrust across the disparate subcultures. Entrenched communities lead to clustered network models that end up with a large diameter, and greater costs for transactions and bookkeeping, lower overall trust across entrenched clusters, and greater levels of distortion in information spread. Entrenched networks also typically have low connectivity and are not resilient against targeted attacks or even large-scale random failures [8–11].

A characteristic property leading to entrenchment is *insularity*. This is the property of an agent to only trust other members from the same community, and by default, distrust connection requests coming from members of other communities.

Any large and diverse population would contain a mix of innate and acquired elements. Varying degrees of innate and acquired diversity lead to different kinds of emergent properties.

In this paper, we model diversity in a population in terms of homophily and insularity. To differentiate between the two similar concepts, we note the following. Homophily is a characteristic displayed by an agent in initiating a new social connection, while insularity is a characteristic displayed by an agent when responding to a social connection request. We study how varying degrees of insularity and homophily affect the emergent network structure.

A common approach to managing diversity is to create a "globalized" subculture. This represents a subset of the population that is trained to adopt the larger worldview of the diverse country or organization and be resilient against homophily and insularity considerations. We study the role played by such a group of globalized agents and observe how the proportion of their population affects the overall network structure, distribution in payoff and power across the population.

2. Related Literature

The idea of acquired and innate diversity is more popularly known as "melting pots" and "salad bowls." The term melting pot was coined by Israel Zangwill in 1907, who wrote a play with the same name (The Melting Pot (play) en.wikipedia.org/wiki/The_Melting_Pot_(play)). The play portrayed America as the great cultural melting pot having the ability to create a society free of ethnic divisions and hatred. The metaphor of a salad bowl came much later, when societies were studied along multiple points of view like communities, markets and energy, cultural differences and their effect on housing prices or investments.

Homophily [12–14] and insularity [15, 16] have also been studied in the field of social psychology and are identified as important characteristics of populations. Our assertion is that these two characteristics form the foundations of diversity in populations, and studying network structures at various extents of the two characteristics would give us important insights into managing diversity.

Agent-based architectures have been an attractive paradigm for modeling societies and culture. They have been used to study emergent properties of societies based on individual agent-level characteristics. Some of the work using agent-based architectures has been done by [17–21] to build generic utility-based models of societies. These agent-based models can also be used to simulate multiple scenarios in the emergent societies, such as power dynamics and distribution of various metrics.

Convention formation is also a relevant area in this context. Emergence of conventions is studied using agent-based network models by Airiau et al. [22, 23], using multiple interactions among agents over time. It is found that the convention formation process is faster when the agents have the ability to reorganize their neighborhood based on the payoffs they get when interacting with their neighbors.

Group formation in networks has been addressed in detail in [24–26], where emergent groups in a network are studied in detail. In most of these works, it is found to be beneficial for agents to interact with other agents belonging to the same group or congregation, as it involves greater trust and familiarity. Another interesting takeaway is that although the overall network keeps changing dynamically at an individual level, it is found that the global properties of the network remain more or less the same.

Network creation has also been addressed from a game theoretic point of view to compute the Nash equilibrium of the emergent networks [27, 28]. In some cases, the network remains the same and the agents evolve their strategies over time [9], while in other cases the agents change their connections, leading to different emergent networks [27, 28]. The network structure has an impact on individual agents and on the resultant network as well [29].

Game theory is useful for modeling interactions among the agents. For example, game theoretic concepts to build models of heterogeneous networks have been used [30, 31]. They prove the existence of Nash equilibrium in different conditions. Also they have some variants of evolutionary models, which are used to study the dynamics of populations of agents over multiple generations. Based on the payoffs of their interactions, the agents decide which connections to make or break and which actions to perform.

Diversity is useful in a variety of contexts like networks, teams and problem solving, like predictive analysis and information and preference aggregation [32–34]. It can be of various kinds, like diversity of paradigms, heuristics, interpretations, values, processes and so on, apart from the usual interpretations of diversity, like diversity in terms of identity and cognitive ability. Diversity needs to be sustained, especially in heterogeneous networks in order to advocate pro-social behavior, for example, by encouraging social and economic interactions among people belonging to diverse groups [35].

To the best of our knowledge, the literature surveyed is relevant to, or is closest to, our model of diversity. In our proposed model, we use a combination of game theory and network analysis to model the dynamics of innate and acquired diversity. The underlying question pertains to studying how varying levels of insularity and homophily, as well as varying proportions of a globalized subgroup, affect the prospects of regional subgroups as well as the population as a whole.

3. Approach and Model

Our model of the diverse population comprises a set of agents A, where $|A| = n$. Of these agents, a subset $G \subseteq A$ of agents represents "global" agents, who associate themselves with the *global identity*—that of the entire network. The term $g_p = |G| / |A|$, also called *perc-global*, represents the proportion of global agents in the population.

The rest of the agents $R = A \backslash G$ are called "regional" or "local" agents, who associate themselves with one from a set of *regional identity* labels, called R_l. A labeling function $\rho : R \to R_l$ is defined, which associates a local identity with every regional agent.

Each agent may make up to k connections with other agents. A connection represents a social or business relationship, potentially leading to mutual value addition. Since each such relationship is expensive to initiate and maintain, the term k represents a cost constraint.

A relationship is represented as an iterated prisoner's dilemma (IPD) game [36]. The game matrix used in our experiments is shown in Table 1.

		Player 1	
		Cooperate	Defect
Player 2	Cooperate	(3, 3)	(−2, 5)
	Defect	(5, −2)	(−1, −1)

Table 1. Payoff matrix of prisoner's dilemma.

The prisoner's dilemma represents a situation where participating agents have two options: act cooperatively, or "defect." There is a clear collective good to be achieved by acting cooperatively. However, the game also provides a *temptation*—where an agent can reap a much higher reward by defecting, as long as the other agent is cooperating. In such cases, while the defecting agent obtains a high payoff, the cooperating agent suffers a penalty. When both agents defect, there is a collective penalty. This is also called the price of anarchy (PoA).

When this game is played as a one-off transaction, a decision to defect emerges as the dominant strategy for any player.

The prisoner's dilemma is one of the most studied games in game theory that accentuates the dilemma between optimizing for one's own prospects versus working toward a global good. While the prisoner's dilemma has no rationale for cooperation when played as a one-off transaction, it can be shown that cooperation can emerge when this game is played in an iterated fashion between the same set of players [36, 37]. IPD brings an element of memory into the game, where knowledge about the opponent's past choices may affect the other player's future choices. As a result, strategies in IPD need to consider future repercussions of a particular choice, in addition to its current expected payoff.

Future repercussions are modeled using a decay parameter indicating the importance attached to expected payoffs over time: δ, $0 \le \delta < 1$. Hence if the expected payoff for player i were to remain u_i for the rest of the game, the overall payoffs factoring future repercussions would be:

$$u_i^* = \lim_{n\to\infty} \sum_{j=0}^{n} \delta^j u_i = \frac{u_i}{1-\delta}. \tag{1}$$

Given such a formulation, it is shown that a strategy based on reciprocity, called tit-for-tat (TFT) emerges as a stable strategy that is both cooperative as well as resilient [36]. The TFT strategy is detailed as follows:

1. Begin by offering cooperation in the first iteration.

2. For subsequent iterations, choose the choice made by the other player in the previous iteration.

Global agents implement the TFT strategy when they interact with any player. Among regional agents, the TFT strategy is somewhat modified to account for insularity. A parameter called i_p or *insularity-prob*, models the extent of insularity in a regional society. This represents the proportion of the regional population that is insular. Every agent in a given regional population is marked as either "insular" or "open-minded" with a probability i_p and $(1 - i_p)$, respectively.

Insular agents implement the following modified version of the TFT strategy (called "distrustful TFT" or DTFT) when they interact with other agents:

1. If another agent interacting with me belongs to the same regional group as me, then:
   - Begin by offering cooperation for the first iteration;

   Else
   - Choose the distrustful choice of *defect*.

2. For subsequent iterations, choose the choice made by the other player in the previous iteration.

There are a variety of strategies that have been studied for IPD. However, TFT and DTFT represent stable and resilient strategies that are also compact. They do not need to store and process a large amount of information to make a decision, and other agents do not have a dominant counter strategy even when they know the TFT or DTFT choices of any given agent.

The limit to which an iterated game is played is called an "epoch" (denoted by the symbol τ). After an epoch, if the overall payoff accrued by an agent is negative, it means that the business or social connection was not beneficial. In such cases, it ends its existing connection and searches for a new connection.

A subtle difference between the formulation of evolutionary games and that of the presented model may be noted. In evolutionary games, agents evolve their strategies after an epoch, based on the accrued payoffs across different strategies. In this case, we allow the network topology to evolve instead, until the topology reaches a state of equilibrium.

This process of rewiring the network is described as follows:

Searching for new connections is influenced by the levels of homophily latent in the population. This is represented by the symbol h_p $(0 \leq h_p \leq 1)$ or *homophily-prob*. With probability h_p an agent connects randomly with other agents from their own group. And with a probability $\left(1 - h_p\right)$ they connect "rationally" to any agent in the population based on the logic of *preferential attachment*. Preferential attachment is based on the degree of the target node. A connection is proposed to node v with a probability: $d_v \big/ \sum_v d_v$, where d_v is the degree of the node v. The degree of a node is used as a marker for the local influence or bargaining power possessed by the corresponding agent.

Global agents are not affected by the homophily probability. While they can end existing connections based on their accrued payoffs, their new connections are always based on preferential attachment.

We can see that when all the agents end up with a net positive payoff from all their relationships, the network topology is in a state of equilibrium, which is Pareto optimal. At this point, none of the agents can change their strategies/connections without hurting their neighbors. We stop the simulation at this stage. Measurements and interpretations are performed after the network has reached equilibrium.

The overall simulation model in the form of pseudocode is presented in Algorithm 1. Based on the simulation model given, the following metrics are computed:

- num-epochs: Number of epochs after which equilibrium is reached.
- modularity (mod_p): Modularity measures the amount of segregation or community structure in the population. It is defined based on the number of in-community links versus the number of across-community links. This is calculated as:

  $$Q = \sum_{i=1}^{c} (e_{ii} - a_i^2),$$

  where i iterates over all the communities, which in the case of our experiments is four: global agents and the three regional groups. The fraction of edges e_{ii} is where both ends of the edge are part of the same community i, and a_i is the fraction of ends of edges that are attached to vertices in community i, that is, either it is an in-edge or an out-edge with respect to community i.
- bet-cen (β): Betweenness centrality of an agent is calculated as the proportion of shortest paths between all pairs of agents that pass through the current agent. Formally betweenness of a node v is given by:

  $$\beta_v = \sum_{\forall s \neq v} \sum_{\forall t \neq v} \frac{|\sigma_{st}(v)|}{|\sigma_{st}|},$$

  where σ_{st} is the set of all shortest paths between s and t, and $\sigma_{st}(v)$ is the set of all shortest paths between s and t passing through v.

```
Generate n agents
Mark each agent as "global" with a probability g_p
Assign a label chosen uniformly at random from R_l to each of the
agents not marked "global"
For each non-global agent, mark agent with label "insular" with proba-
bility i_p
Global agents and non-insular regional agents adopt strategy TFT
Insular agents adopt strategy DTFT
Every agent connects with k agents
while TRUE do
  Agents play IPD with all their neighbors for τ iterations
  count = 0
  for each agent a do
    let neighbor be the set of neighbors of agent a
    if payoff from neighbor[i]  < 0 then
      count = count + 1
      if agent a is global then
        Disconnect with neighbor[i]
        Connect using preferential attachment with any random agent
      else
        Disconnect with neighbor[i]
        With probability h_p within the group and with probability
          (1 − h_p) connect preferentially to any random agent
      end
    end
end
  if count == 0 then
    System reached equilibrium
    Update all metrics and plots
    STOP
  end
end
```

Algorithm 1. Pseudocode of model.

The terms β_g and β_r represent the average betweenness centrality of global and regional agents, respectively. Similarly β_c and β_d represent the average betweenness centrality of cooperating (non-insular) regional agents and defecting (insular) regional agents, respectively. These group-level betweenness centrality measures are useful to characterize the relative importance of groups with respect to each other. Similarly, payoff_g and payoff_r represent the average payoffs of global and regional agents, respectively.

- Dominance over neighborhood (DON): DON represents a new centrality measure indicating local influence, proposed by [38]. A standardized version of the measure, with certain parameters fixed, is expressed as:

$$\mathrm{DON}_i = 1 + \frac{1}{n\rho} \sum_{j \in N(i)} \ln\left(\frac{d(i)}{d(j)}\right). \tag{2}$$

  Here, $d(i)$ represents the degree of node i and $N(i)$ represents the set of neighbors of node i. The distribution of DON values in a network can be shown to represent a "fair" allocation of a conserved resource (in our case, attention) among agents in a network, and is shown to be a generalized form of the Shapley value [39, 40]. The term ρ is a scaling parameter to produce the best dispersion in the distribution of DON values. The sum of DON values over all agents is n and yields a balance property. The log sum of ratios of the degrees ensures that DON has a high positive value whenever its degree is high, and additionally, the degree of its neighbors is low. This captures the notion of dominance of node over its neighbors, which is used in our study as a measure of the bargaining power of an agent.

  Normalizing the values of DON by dividing by n yields a distribution of power or influence values that sums to 1. The entropy of this distribution presents an insight into how bargaining power is distributed across the population. This is calculated as follows:

$$H_{\mathrm{DON}} = -\sum_i \left(\frac{\mathrm{DON}_i}{n}\right) * \log_2\left(\frac{\mathrm{DON}_i}{n}\right). \tag{3}$$

  Smaller values of entropy imply the emergence of power centers of influence, while larger values indicate a more even distribution of bargaining power in the population.

4. Experiments and Results

All the experiments have been done using NetLogo (ccl.northwestern.edu/netlogo). The network of agents is initialized as discussed in Section 3, and Table 3 shows the initial values of the input parameters. The network is updated until it reaches equilibrium, that is, when all the agents get a positive payoff from all their neighbors and cannot change their strategies/connections without hurting others. At equilibrium, all the metrics defined in Section 3 are measured and reported. We conduct 500 independent runs for each of the network configurations in order to minimize the effect of outliers. We then take the mean value of the scores across all the runs. Error bars in all the graphs depict standard error.

In all the network visualizations, the color of the agent represents the group it belongs to. Regional agents are represented in different colors: red, green and dark blue; while global agents are shown in cyan blue. The shape of the agents also represents their characteristic.

A circle represents a regional "open-minded" (non-insular) agent that plays TFT. The triangle represents regional insular agents that play DTFT and the squares represent global agents that are not affected by homophily or insularity. Network visualizations are for one sample run of the case being discussed. The networks are plotted using *spring layout* [41] where the links in the graph act as springs and the agents connected by a link repel each other. All the spring layout–related hyperparameters remain the same across all the plots.

conf	epochs	mod_p	β_g	β_r	β_c	β_d	payoff_g	payoff_r
glo 0%								
LILH	4.24	0.22	-	87.16	102.60	26.03	-	7289.66
LIHH	3.12	0.56	-	109.59	123.86	52.48	-	5387.37
HILH	11.13	0.57	-	107.06	369.93	45.57	-	17851.21
HIHH	4.15	0.65	-	96.73	252.63	57.55	-	6956.79
glo 20%								
LILH	4.58	0.19	66.24	91.87	108.90	25.09	7055.11	7969.64
LIHH	3.59	0.48	79.20	108.56	124.39	45.49	3943.50	6659.00
HILH	11.81	0.52	59.20	114.64	411.72	45.22	13332.32	19558.17
HIHH	7.18	0.63	191.80	130.86	429.15	60.37	4735.78	13323.76
glo 80%								
LILH	5.02	0.07	83.74	97.95	119.38	13.49	8758.00	8437.12
LIHH	3.73	0.21	89.63	100.64	124.61	10.68	6113.55	7167.35
HILH	8.01	0.15	80.55	92.25	352.41	34.66	14341.43	7920.86
HIHH	6.25	0.27	92.85	75.37	340.07	13.61	10426.93	9438.20

Table 2. Metrics at four extremes: all scores presented are averaged over 500 independent runs.

Parameter	Value
♯ of agents	100
♯ of regional identities R_l	3
♯ of edges per agent k	3
epoch length (τ)	50 ticks
insularity-prob	[0.2 0.8] (two extreme cases)
homophily-prob	[0.2 0.8] (two extreme cases)

Table 3. Initial values of input parameters.

Simulation runs are performed for different "cases," where each case represents a specific diversity configuration for the population.

Each such case is also characterized by one or more hypotheses, which are tested on the simulation runs. The different cases are introduced below.

4.1 Case A: No Global Agents

The first case we study is of a population that contains only regional identities, that is: $G = \{\}$, in order to see how critical the population of global agents is to keep a diverse population together.

Hypothesis (HA1): Fostering a subculture of “global” agents in a diverse population is critical for keeping the population functioning as one network.

In order to test this hypothesis, first we build a network model with only regional agents. Insularity and homophily probability of the regional agents are varied, and simulation runs are conducted for all four extremes and the resultant network structures. These extreme cases correspond to the following: LILH, LIHH, HILH, HIHH, representing combinations of low and high values of insularity and homophily, respectively. A low value is represented as a probability of 0.2 and a high value is represented as a probability of 0.8. Resultant network structures for all the four extremes are shown in Figure 1. In the case of LILH where both insularity and homophily are low, the network forms a tightly knit melting pot. But when either insularity or homophily increases, the network segregates into regional clusters, with the open-minded agents from each cluster forming the bridges to other clusters. The modularity values in either of these cases are significantly higher than in the LILH case (shown in Table 2, first section).

When both homophily and insularity are high (HIHH), the network disintegrates into disparate clusters with little or no interaction between the clusters.

Inference: Without the presence of global agents, there is more responsibility on regional cooperators. And in the absence of sufficient regional cooperators, the population can actually cease to function as a single network.

Hypothesis (HA2): The non-insular agents earn a higher payoff than insular agents since they play the central role of holding the network together in the absence of global agents.

Figure 2 shows the relative payoffs of non-insular and insular agents obtained at equilibrium. Non-insular agents end up with higher payoffs than insular agents in all four configurations. And interestingly, their payoffs are significantly higher in the case of HILH configurations comprised of a large number of insular agents, where they play the most prominent role.

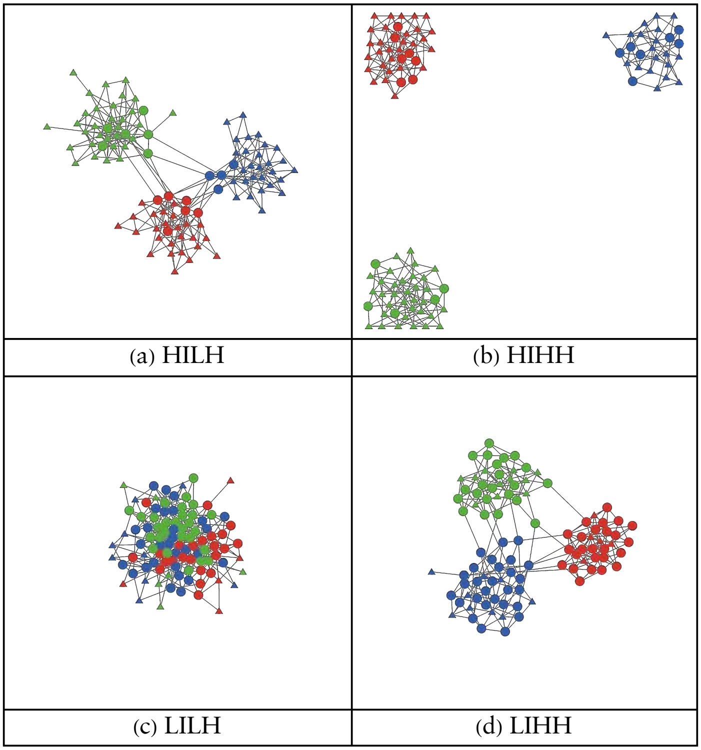

Figure 1. Sample network graphs with only regional agents, where colors represent the type of agent and shape represents the strategy of the agent (see the text).

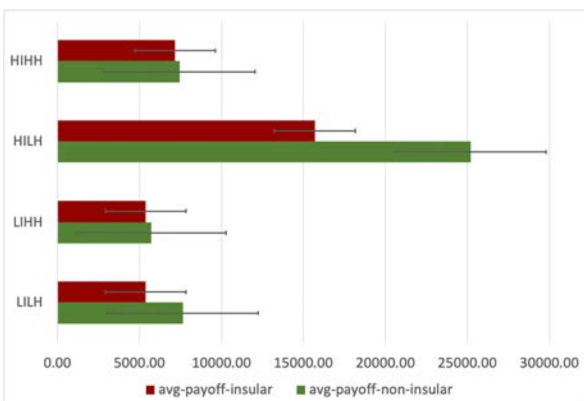

Figure 2. Average payoffs of insular and non-insular regional agents with 0% global agents.

Inference: In HILH, the non-insular agents play a critical role in connecting the network together into a single unit, and they have higher payoffs in that configuration both as compared to other network configurations as well as higher payoffs than insular agents in the same configuration.

4.2 Case B: Small (20%) Ratio of Global Agents

In the second case, we introduce a small number (20%) of agents who are trained with the "global" identity, in order to understand their role. Figure 3 shows the network configurations at equilibrium for low (0.2) and high (0.8) values of homophily and insularity.

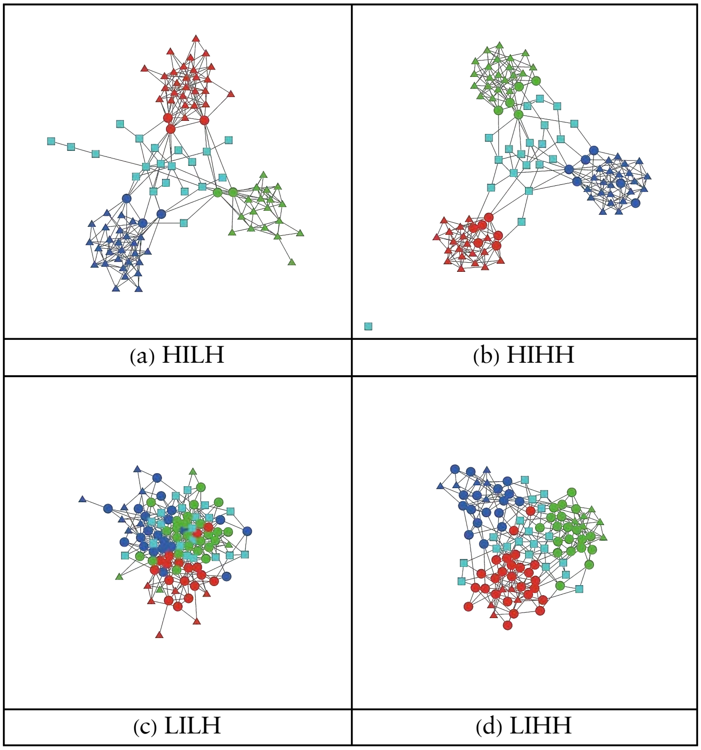

Figure 3. Sample network graphs with both global (20%) and regional agents, where colors represents the type of agent and shape represents the strategy of the agent, as described in the text.

Hypothesis (HB1): In a diverse population with increasing insularity and homophily, global agents play a central role in keeping the network functioning as one unit.

In Table 2, in row HIHH of 20% global agents, we observe that global agents have average betweenness centrality of 191.8, whereas regional agents have an average betweenness centrality of 130.86. This shows that global agents play the central role of holding the agents together, especially in the case of HIHH, where the network gets split into multiple parts in the absence of global agents.

Inference: Global agents have significantly higher betweenness centrality than regional agents in HIHH configuration.

Hypothesis (HB2): Global agents play a central role in keeping the network together, and hence they have the highest payoffs.

In Figure 4, the average payoffs of global and local agents are compared across all four network configurations. We find that the average payoff of global agents is always lower than that of regional agents in all the configurations! The payoff for both regional and global agents is highest in the case of HILH. Although HILH takes longer to reach equilibrium, we observed that even average per epoch payoff (not reported separately as it follows similar trends as overall average payoffs) for both global and regional agents is also the highest in case of HILH.

Inference: We fail to find support for HB2 in the test runs. Contrary to the hypothesis, we found that global agents get an overall lower payoff as compared to regional agents in all network configurations, despite their critical role of holding the network together.

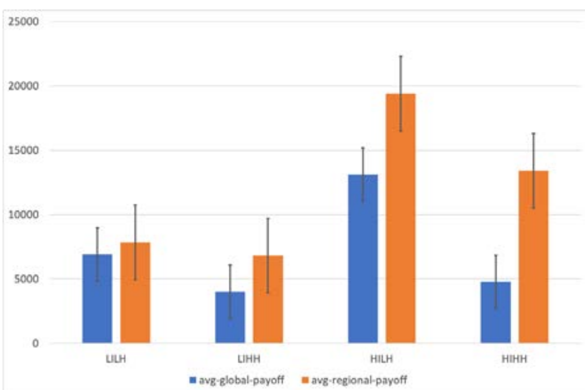

Figure 4. Average global versus regional payoffs.

Hypothesis (HB3): Since global agents play a central role in keeping the network together, they will have high bargaining power on average.

Figure 5 shows the scatter plot of DON versus payoff at equilibrium for one sample run of simulation in HIHH configuration with 20% global agents. The x axis represents the DON values of agents and the y axis represents the payoffs of the agents. Both DON and payoff values have been normalized to lie in the range $[0, 1]$. Blue dots represent global agents, red dots represent regional defectors and green dots represent regional cooperators. All the global agents appear in the lower-left corner of the graph, indicating that despite their high betweenness, global agents also do not enjoy high bargaining power! Regional, non-insular agents fare much better in terms of their payoffs and bargaining power, as compared to global agents.

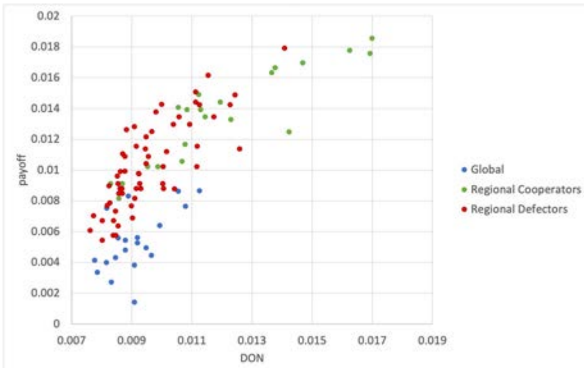

Figure 5. DON versus payoff of agents in HIHH with $\rho = 0.1$ and 20% global agents.

Inference: Again contrary to the hypothesis, we found that regional cooperators have a higher DON and payoff as compared to all other categories of agents. Global agents have both lower payoff and DON, especially when their ratio is less. Regional defectors have similar DON values, but with higher payoffs.

Case B suggests that with a small number of global agents, not only does the population function as one network, the regional subgroups also thrive. The non-insular agents from regional subgroups end up having the highest payoffs and bargaining power, while the global agents play the critical role of keeping the network functioning as one unit. The only caveat is that there is no rational incentive for any agent to be trained into the global agent worldview, since global agents end up being highly stressed (due to their high betweenness) and yet fare lower in payoffs and bargaining power than the regional cooperating leaders.

4.3 Case C: High (80%) Ratio of Global Agents

In the next case, we introduce a large number (80%) of agents who are trained with the "global" identity, in order to understand what happens when the global identity becomes so powerful so as to mask regional identities. Figure 6 shows the network configurations at equilibrium for low (0.2) and high (0.8) values of homophily and insularity.

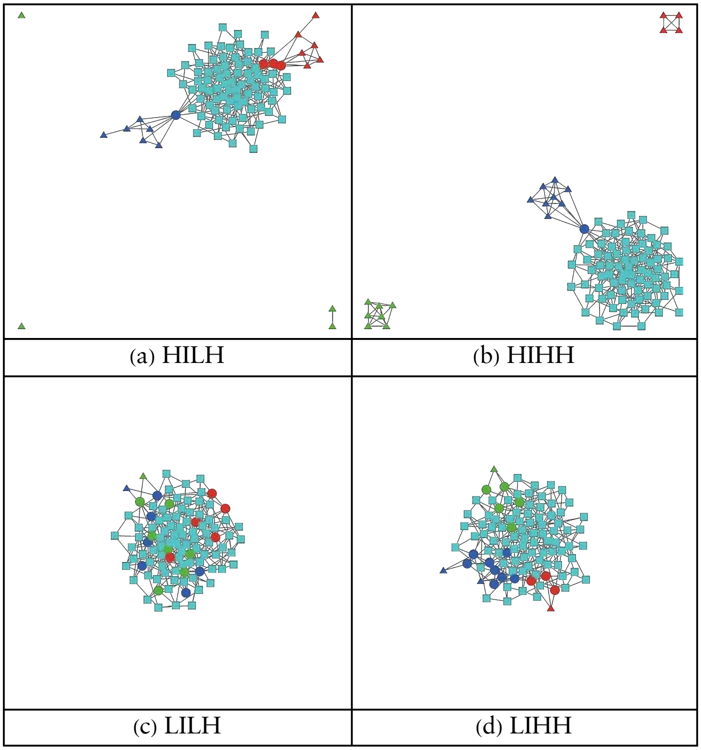

Figure 6. Sample network graphs with high (80%) global agents, where colors represent the type of agent and shape represents the strategy of the agent, as described in the text.

Hypothesis (HC1): High percentage of global agents in a network will be good for regional agents.

Figure 6 shows the four extreme resultant network configurations, with 80% global agents. We observe that in HIHH and HILH configurations, some of the regional agents might get completely disconnected from the network. In LIHH configurations, although they are a part of the larger group, they cluster together with other regional agents of their type. Also as shown in Table 2, when the number of global agents is very high, high levels of homophily (and low insularity) is the only case when the regional agents are better off than the global agents. A large percentage of global agents is not good for the prospects of regional agents specifically in high insularity configurations, and they risk getting alienated from the network. The small number of cooperators that act as a bridge between the regional agents and global agents are single-point failures because if they fail, then the whole community that was connected via them to the bigger network gets disconnected.

Inference: Hypothesis HC1 fails to be supported by the test runs. Having a large number of global agents might be detrimental for the regional agents, as they may get disconnected or completely sidelined in the network.

When the number of global agents is small, they play critical roles in keeping the network together, but end up having lower payoffs and bargaining power. With high numbers of global agents, there is a risk of alienating regional identities. This leads us to the question of what would be the "ideal" proportion of global agents in a diverse population.

4.4 Case D: Varying Ratio of Global Agents

In this case, we vary the proportion of global agents to see how the resultant change in demographics affects the prospects of both global agents and regional agents.

Hypothesis (HD1): As the number of global agents increases in HIHH, their payoffs increase.

Figure 7 shows the payoff trends of global and regional agents as the percentage of global agents increases in HIHH configuration. We observe that average payoffs of global agents monotonically increase, whereas average payoffs of regional agents increase until the network has around 45% global agents and then the average payoff starts decreasing. At 75% global agents, both global and regional agents have the same average payoffs, after which global agents have higher payoff than regional agents.

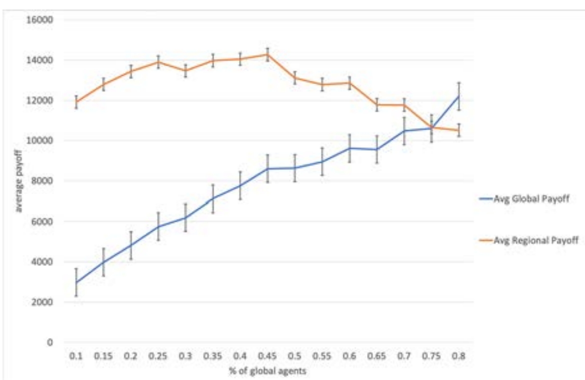

Figure 7. Average global versus regional payoffs with varying global agents in HIHH.

Inference: It is beneficial for regional cultures to participate in a diverse network until the proportion of global agents is about 45% of the overall population. After this point, an increase in the ratio of global agents is detrimental for regional agents in terms of payoff. But a truly "equal" society (solely in terms of payoffs) that includes the welfare of the global agents too happens when the percentage of global agents is about 75%.

Hypothesis (HD2): As the number of global agents increases, bargaining power is more evenly spread across the population.

In Figure 8, for the HIHH configuration, we observe that with an increasing percentage of global agents, entropy of DON or bargaining power of global agents monotonically increases, whereas the bargaining power of local agents monotonically decreases. An increase in entropy of DON values indicates a more even distribution of bargaining power, while decreasing entropy indicates emergence of power centers. Increasing entropy of global agents indicates that with increasing population of global agents, they evenly spread their influence over the network. In contrast, among the regional agents, bargaining power tends to get concentrated in a few power centers, as the percentage of global agents increases.

Inference: In a population with high levels of insularity and homophily, increasing the proportion of global agents tends to create flatter power structures among global agents, but also tends to create hierarchical power centers among the regional populations.

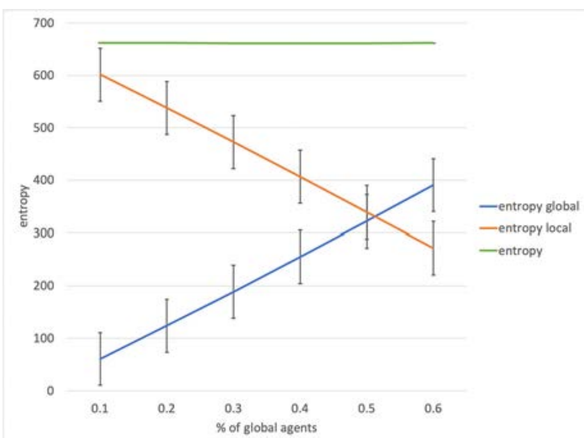

Figure 8. Entropy of DON in HIHH with $\rho = 0.1$.

4.5 Case E: Generic Results

Hypothesis (HE1): Populations with high insularity take longer to reach equilibrium.

In Figure 9, we observe that across different levels of global agents, it takes longer to reach equilibrium in configurations where insularity is high. HILH takes the longest to attain equilibrium among all four configurations, followed by HIHH for cases where the percentage of global agents is nonzero. Another interesting point to note in Figure 9 is that LIHH reaches equilibrium faster than even LILH! LILH is like the perfect melting pot with people cooperating and trusting others and open to making new connections with others, whereas in LIHH some agents are not open to making connections with others, yet it reaches equilibrium faster. But this equilibrium would be in the form of segregated regional clusters, which results in a much lower overall payoff (depicted in Figure 4) as compared to LILH.

Inference: Insularity defers the network from attaining equilibrium faster than homophily. Homophily might actually speed up the process of equilibrium in LIHH.

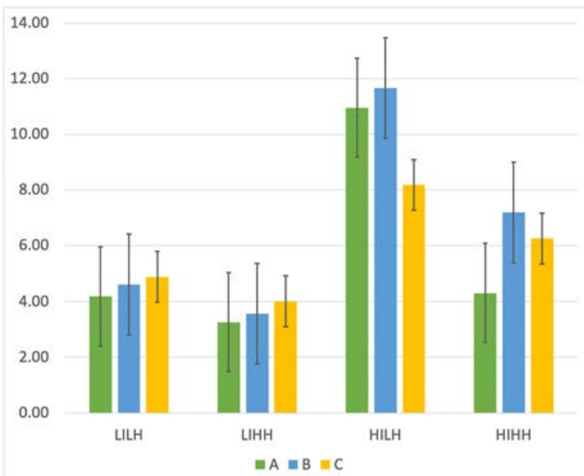

Figure 9. Average number of epochs to reach equilibrium in all four configurations where A represents 0% global agents, B represents 20% global agents and C represents 80% global agents.

Hypothesis (HE2): There are significant differences in resultant network characteristics of high insularity and high homophily configurations.

In the observed betweenness centrality of cooperators and defectors, we find that there is a much higher difference among the two in the case of HILH configuration than the LIHH configuration. This shows that cooperating agents play a more central role in HILH as compared to LIHH, and that the network is more evenly spread in the case of LIHH as compared to HILH. Also betweenness centrality of global agents lies in between the betweenness centrality of defectors and cooperators.

Inference: Insularity and homophily lead to characteristically different emergent networks, and these concepts are not interchangeable with one another.

5. Discussion and Conclusion

The set of experiments presented in this paper presents various facets of diversity in populations. While some aspects of the experiments can be inferred from individual agent behavior, simulation results show complexities and insights that are not directly apparent. The primary insight in this paper was to model diversity in terms of two underlying elements: insularity and homophily. Insularity is seen in homogeneous populations that are forced to interact and become part of a larger heterogeneous group. It is characterized by distrust of the "other," resulting in close-knit groups. Some examples include incoming refugee populations that relocated as a group and have been forced out of their original dwellings and are compelled to coexist in an alien culture.

In contrast to insularity, homophily is more pronounced in populations with acquired diversity comprising of willing immigrants. Here, immigrant subgroups have much less distrust of the other, given their conscious choice to be part of the diverse population. However, they may prefer establishing connections within their subgroup due to higher familiarity with each other. Examples include intra-ethnic business relationships and marriages among immigrants in a diverse country.

Insularity and homophily are characteristically different. In case A (Section 4.1), we can see that although insularity and homophily both result in a segregated population, they are not symmetric. High homophily and low insularity (representing acquired diversity) results in segregated clusters, with several intercluster links among the different communities. But with high insularity, the number of intercluster links is far smaller (reducing to zero as homophily increases). High insularity has much more drastic variation in betweenness centrality of insular and non-insular agents as compared to high homophily.

Similarly, HILH configurations take the longest amount of time to achieve equilibrium. Yet, HILH results in the highest average per-epoch payoff for all regional groups, including global agents. Our best explanation for this is that this configuration represents a balanced interplay between an ethnic group being conservative (high insularity) as well as open-minded (low homophily).

Yet another insight from the experiments was the critical nature of global agents in keeping the population together. In cases of both high homophily and high insularity, the global agents formed the critical link between disparate communities. However, despite their central role (in terms of betweenness) in the population, in case B (Section 4.2), they were neither the richest nor the most powerful of all the different groups. With a small percentage of global agents acting as the social glue, the biggest beneficiaries are the regional agents who interface between the global backbone and regional clusters.

It would be interesting to model evolutionary dynamics to understand how demographic changes may manifest due to variations in payoffs and bargaining power. Would the lower payoffs and bargaining power for global agents create a disincentive for a global agent to continue with a global worldview? Or would the high payoffs and bargaining power obtained by open-minded regional agents encourage other regional agents to become less insular? We relegate such questions to future research directions for this work.

Yet another insight from this paper in case C (Section 4.3), is that while global agents form the social glue to keep the network together even when there are high levels of insularity and homophily, too high a percentage of global agents may itself contribute to fragmentation of the network. Combining this with case D (Section 4.4), we see that an increase in the number of global agents increases the prospects of regional agents up to a point, after which their payoffs start reducing. In addition, we also see that increasing the percentage of global agents only increases power disparity among the regional clusters, resulting in a small number of power centers, when homophily and insularity levels are high.

This might serve as a cautionary note in showing that it is possible to go overboard with efforts toward integration and assimilation, which may ironically lead to more fragmentation and alienation. Cases C and D (Sections 4.3 and 4.4) only reinforce the challenging nature of diversity management, calling for nuanced approaches that balance between assimilation and retention of regional identities.

5.1 Real-World Examples of Managing Diversity

In this paper, we looked at insularity and homophily in population and the role of global agents who do not act based on these aspects. These factors are not just modeled in the context of simulations but are observed in the real world as well. As discussed, insularity and homophily are related to social trust. Social trust is of two types: generalized trust and particularized trust [42]. Particularized trust is trusting people of a specific community or group, which is characteristic of both homophily and insularity. Generalized trust is trusting people irrespective of their association with groups, which is similar to the behavior of global agents in our model. Our simulations show the crucial role played by global agents, and it has also been shown that generalized trust (which is similar to global agents) is positively associated with physical health, happiness and life satisfaction in societies [43]. Similar to our argument that managing diversity cannot be a one-size-fits-all approach but rather requires a nuanced approach, Hamamura et al., [43] find that managing generalized trust requires different measures in developed versus developing societies. There have been studies to check how welfare benefits contribute to form high generalized trust in Nordic countries like Finland [44]. Generalized trust (similar to global agents) is crucial both at the individual and community level and can build expansive communities [42, 45].

Acknowledgments

This work was supported by Karnataka Innovation & Technology Society, Dept. of IT, BT and S&T, Govt. of Karnataka vide GO No. ITD 76 ADM 2017, Bengaluru; Dated 28.02.2018.

References

[1] T. H. Cox and S. Blake, "Managing Cultural Diversity: Implications for Organizational Competitiveness," *Academy of Management Perspectives*, **5**(3), 1991 pp. 45–56. www.jstor.org/stable/4165021.

[2] E. Kearney, D. Gebert and S. C. Voelpel, "When and How Diversity Benefits Teams: The Importance of Team Members' Need for Cognition," *Academy of Management Journal*, **52**(3), 2009 pp. 581–598. doi:10.5465/amj.2009.41331431.

[3] J. F. Milem, "The Educational Benefits of Diversity: Evidence from Multiple Sectors," *Compelling Interest: Examining the Evidence on Racial Dynamics in Colleges and Universities* (M. J. Chang, D. Witt, J. Jones and K. Hakuta, eds.), Stanford, CA: Stanford Education, 2003 pp. 126–169.

[4] W. R. Poster, "The Challenges and Promises of Class and Racial Diversity in the Women's Movement: A Study of Two Women's Organizations," *Gender & Society*, **9**(6), 1995 pp. 659–679. doi:10.1177%2F089124395009006002.

[5] F. G. Stevens, V. C. Plaut and J. Sanchez-Burks, "Unlocking the Benefits of Diversity: All-Inclusive Multiculturalism and Positive Organizational Change," *The Journal of Applied Behavioral Science*, **44**(1), 2008 pp. 116–133. doi:10.1177%2F0021886308314460.

[6] T. Kochan, K. Bezrukova, R. Ely, S. Jackson, A. Joshi, K. Jehn, J. Leonard, D. Levine and D. Thomas, "The Effects of Diversity on Business Performance: Report of the Diversity Research Network," *Human Resource Management*, **42**(1), 2003 pp. 3–21. doi:10.1002/hrm.10061.

[7] B. Kupka, A. Everett and S. Wildermuth, "The Rainbow Model of Intercultural Communication Competence: A Review and Extension of Existing Research," *Intercultural Communication Studies*, **16**(2), 2007 pp. 18–36.

[8] D. S. Callaway, M. E. J. Newman, S. H. Strogatz and D. J. Watts, "Network Robustness and Fragility: Percolation on Random Graphs," *Physical Review Letters*, **85**(25), 2000 pp. 5468–5471. doi:10.1103/PhysRevLett.85.5468.

[9] J. Deshmukh and S. Srinivasa, "Evolution of Cooperation under Entrenchment Effects," in *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Istanbul, Turkey* (E. Bordini and Y. Weiss, eds.), International Foundation for Autonomous Agents and Multiagent Systems, 2015 pp. 1717–1718.

[10] S. Patil, S. Srinivasa and V. Venkatasubramanian, "Classes of Optimal Network Topologies under Multiple Efficiency and Robustness Constraints," in *IEEE International Conference on Systems, Man and Cybernetics, SMC 2009*, San Antonio, TX, (Piscataway, NJ: IEEE, 2009 pp. 4940–4945. doi:10.1109/ICSMC.2009.5346272.

[11] D. J. Watts and S. H. Strogatz, "Collective Dynamics of 'Small-World' Networks," *Nature*, **393**(6684), 1998 pp. 440–442. doi:10.1038/30918.

[12] M. McPherson, L. Smith-Lovin and J. M. Cook, "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology*, **27**(1), 2001 pp. 415–444, 2001. doi:10.1146/annurev.soc.27.1.415.

[13] P. F. Lazarsfeld, R. K. Merton, "Friendship as a Social Process: A Substantive and Methodological Analysis," *Freedom and Control in Modern Society* (M. Berger, T. Abel and H. Charles, eds.), New York: Van Nostrand, 1954 pp. 18–66.

[14] D. Easley and J. Kleinberg, *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*, New York: Cambridge University Press, 2010.

[15] H. Tajfel, M. G. Billig, R. P. Bundy and C. Flament, "Social Categorization and Intergroup Behaviour," *European Journal of Social Psychology*, **1**(2), 1971 pp. 149–178. doi:10.1002/ejsp.2420010202.

[16] M. Billig and H. Tajfel, "Social Categorization and Similarity in Intergroup Behaviour," *European Journal of Social Psychology*, **3**(1), 1973 pp. 27–52. doi:10.1002/ejsp.2420030103.

[17] P. Collard and S. Mesmoudi, "How to Prevent Intolerant Agents from High Segregation?," in *Advances in Artificial Life, ECAL 2011*, Paris, France (T. Lenaerts, M. Giacobini, H. Bersini, P. Bourgine, M. Dorigo and R. Doursat, eds.), Cambridge, MA: MIT Press, 2011, pp. 168–174.

[18] P. Collard, S. Mesmoudi, T. Ghetiu and F. Polack, "Emergence of Frontiers in Networked Schelling Segregationist Models," *Complex Systems*, **22**(1), 2013 pp. 35–59. doi:10.25088/ComplexSystems.22.1.35.

[19] R. V. Mendes and C. Aguirre, "Cooperation, Punishment, Emergence of Government, and the Tragedy of Authorities," *Complex Systems*, **20**(4), 2012 pp. 363–374. doi:10.25088/ComplexSystems.20.4.363.

[20] A. Chliaoutakis and G. Chalkiadakis, "Agent-Based Modeling of Ancient Societies and Their Organization Structure," *Autonomous Agents and Multi-Agent Systems*, **30**(6), 2016 pp. 1072–1116. doi:10.1007/s10458-016-9325-9.

[21] S. Mascarenhas, N. Degens, A. Paiva, R. Prada, G. J. Hofstede, A. Beulens and R. Aylett, "Modeling Culture in Intelligent Virtual Agents," *Autonomous Agents and Multi-Agent Systems*, **30**(5), 2016 pp. 931–962. doi:10.1007/s10458-015-9312-6.

[22] S. Airiau, S. Sen and D. Villatoro, "Emergence of Conventions through Social Learning," *Autonomous Agents and Multi-Agent Systems*, **28**(5), 2014 pp. 779–804. doi:10.1007/s10458-013-9237-x.

[23] M. R. Hasan, A. Raja and A. Bazzan, "A Context-Aware Convention Formation Framework for Large-Scale Networks," *Autonomous Agents and Multi-Agent Systems*, **33**(1), 2019 pp. 1–34. doi:10.1007/s10458-018-9397-9.

[24] C. H. Brooks and E. H. Durfee, "Congregation Formation in Multiagent Systems," *Autonomous Agents and Multi-Agent Systems*, **7**(1), 2003 pp. 145–170. doi:10.1023/A:1024133006761.

[25] S. Phelps, "Emergence of Social Networks via Direct and Indirect Reciprocity," *Autonomous Agents and Multi-Agent Systems*, **27**(3), 2013 pp. 355–374. doi:10.1007/s10458-012-9207-8.

[26] K. Oishi, T. Shimada and N. Ito, "Group Formation through Indirect Reciprocity," *Physical Review E*, **87**(3), 2013 030801. doi:10.1103/PhysRevE.87.030801.

[27] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou and S. Shenker, "On a Network Creation Game," in *Proceedings of the Twenty-Second Annual Symposium on Principles of Distributed Computing (PODC '03)*, Boston, MA, New York: ACM, 2003 pp. 347–351. doi:10.1145/872035.872088.

[28] A. Cabrales, A. Calvó-Armengol and Y. Zenou, "Social Interactions and Spillovers," *Games and Economic Behavior*, **72**(2), 2011 pp. 339–360. doi:10.1016/j.geb.2010.10.010.

[29] J. Hay and D. Flynn, "The Effect of Network Structure on Individual Behavior," *Complex Systems*, **23**(4), 2014 pp. 295–311. doi:10.25088/ComplexSystems.23.4.295.

[30] A. Advani and B. Reich, "Melting Pot or Salad Bowl: The Formation of Heterogeneous Communities," IFS Working Paper W15/30, Institute for Fiscal Studies, 2015. doi:10.1920/wp.ifs.2015.1530.

[31] S. Goyal, P. Hernández, G. Martinez-Cánovas, F. Moisan, M. Muñoz-Herrera and A. Sánchez, "Integration and Diversity," *Experimental Economics*, **24**(2), 2021 pp. 387–413. doi:10.1007/s10683-020-09676-6.

[32] R. Mehrotra, V. Soni and S. Jain, "Diversity Sustains an Evolving Network," *Journal of the Royal Society Interface*, **6**(38), 2009 pp. 793–799. doi:10.1098/rsif.2008.0412.

[33] L. S. Marcolino, A. X. Jiang and M. Tambe, "Multi-agent Team Formation: Diversity Beats Strength?," in *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI '13),* Beijing, China (F. Rossi, ed.), Palo Alto, CA: Association for the Advancement of Artificial Intelligence Press, 2013 pp. 279–285.

[34] S. E. Page, *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies-New Edition*, Princeton: Princeton University Press, 2008.

[35] D. Baldassarri and M. Abascal, "Diversity and Prosocial Behavior," *Science*, **369**(6508), 2020 pp. 1183–1187. doi:10.1126/science.abb2432.

[36] R. Axelrod and W. D. Hamilton, "The Evolution of Cooperation," *Science*, **211**(4489), 1981 pp. 1390–1396. doi:10.1126/science.7466396.

[37] R. Axelrod, "The Evolution of Strategies in the Iterated Prisoner's Dilemma," *Genetic Algorithms and Simulated Annealing* (L. Davis, ed.), Los Altos, CA: Morgan Kaufman, 1987 pp. 32–41.

[38] S. Mandyam and U. Sridhar, "DON and Shapley Value for Allocation among Cooperating Agents in a Network: Conditions for Equivalence," *Studies in Microeconomics*, **5**(2), 2017 pp. 143–161. doi:10.1177%2F2321022217702257.

[39] K. V. Aadithya, B. Ravindran, T. P. Michalak and N. R. Jennings, "Efficient Computation of the Shapley Value for Centrality in Networks," in *International Workshop on Internet and Network Economics (WINE 2010)*, Stanford, CA (A. Saberi, ed.), Berlin, Heidelberg: Springer, 2010 pp. 1–13. doi:10.1007/978-3-642-17572-5_1.

[40] P. L. Szczepanski, T. Michalak and T. Rahwan, "A New Approach to Betweenness Centrality Based on the Shapley Value," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '12)*, Valencia, Spain, Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2012 pp. 239–246.

[41] T. M. J. Fruchterman and E. M. Reingold, "Graph Drawing by Force-Directed Placement," *Software: Practice and Experience*, **21**(11), 1991 pp. 1129–1164. doi:10.1002/spe.4380211102.

[42] E. M. Uslaner, *The Moral Foundations of Trust*, New York: Cambridge University Press, 2002.

[43] T. Hamamura, L. M. W. Li and D. Chan, "The Association between Generalized Trust and Physical and Psychological Health across Societies," *Social Indicators Research*, **134**(1), 2017 pp. 277–286. doi:10.1007/s11205-016-1428-9.

[44] A. Kouvo, T. Kankainen and M. Niemelä, "Welfare Benefits and Generalized Trust in Finland and Europe," *The Future of the Welfare State* (H. Ervasti, J. G. Andersen, T. Fridberg and K. Ringdal, eds.), Northampton, MA: Edward Elgar Publishing, 2012. doi:10.4337/9781781001271.00015.

[45] D. Stolle, "Trusting Strangers: The Concept of Generalized Trust in Perspective," *Österreichische Zeitschrift für Politikwissenschaft*, **31**(4), 2002 pp. 397–412. nbn-resolving.org/urn:nbn:de:0168-ssoar-60076.