

The Influence of Social Status on Consensus Building in Collaboration Networks

Iliire Hasani-Mavriqi, Florian Geigl, Subhash Chandra Pujari, Elisabeth Lex, Denis Helic
 Knowledge Technologies Institute, KTI
 Graz University of Technology
 Inffeldgasse 13/VI 8010, Graz, Austria
 Telephone: +43316 87330803, Fax:+43316 87330802
 Email: ihasani@know-center.at, {florian.geigl, s.pujari, elisabeth.lex, dhelic}@tugraz.at

Abstract—In this paper, we analyze the influence of social status on opinion dynamics and consensus building in collaboration networks. To that end, we simulate the diffusion of opinions in empirical collaboration networks by taking into account both the network structure and the individual differences of people reflected through their social status. For our simulations, we adapt a well-known Naming Game model and extend it with the *Probabilistic Meeting Rule* to account for the social status of individuals participating in a meeting. This mechanism is sufficiently flexible and allows us to model various situations in collaboration networks, such as the emergence or disappearance of social classes. In this work, we concentrate on studying three well-known forms of class society: *egalitarian*, *ranked* and *stratified*. In particular, we are interested in the way these society forms facilitate opinion diffusion. Our experimental findings reveal that (i) opinion dynamics in collaboration networks is indeed affected by the individuals' social status and (ii) this effect is intricate and non-obvious. In particular, although the social status favors consensus building, relying on it too strongly can slow down the opinion diffusion, indicating that there is a specific setting for each collaboration network in which social status optimally benefits the consensus building process.

I. INTRODUCTION

It is our natural predisposition to interact with people who have a high social status in our social communities. Customarily, our social interactions and, to some extent, our behavior are influenced by actions of individuals with a high social status. In the field of social psychology, the social status theory attempts to explain this phenomenon [1, 2, 3]. According to it, people tend to form their connections in a social network to maximize their perceived social benefits arising from the social status of their connections. Also, in the work of Guha et al. [4] the authors relate social status to the mechanism of link formation in a social network, hypothesizing that people with a lower social status are more likely to create (directed) links with people of a higher social status.

In this paper, however, we are not interested in the relation between the social status and the process of link formation, but rather in the relation between social status and *dynamical processes* that may take place in a social or collaboration network (i.e., a special case of social network, in which users collaborate). One example of such dynamical process is a so-called opinion dynamics process. In our daily lives, we interact with our peers, discuss certain problems, exchange opinions and try to reach some kind of consensus. The question we want to answer in this paper is how social status influences

such processes in a collaboration network. For example, in a university class there is a lively discussion between a student and her mentor regarding their newest experimental results and their interpretation. The mentor has a higher social status than the student, due to a superior education, a broader experience and a higher position in the organizational hierarchy. Undoubtedly, while trying to reach a consensus, the student will be influenced by opinions of her mentor because of the latter's convincing power [5, 6]. The literature [5] identifies this process as dynamics of agreement/disagreement between persons belonging to a social group. For clarity, in this paper we will refer to it as opinion dynamics.

Problem. The aim of this work is to investigate the influence of social status on the process of reaching consensus within a social community that has a heterogeneous distribution of social status. In particular, we are interested in social communities in which interaction between the community members is empowered by social media. While there is a substantial body of work on opinion dynamics (see Section V) in general settings, we focus on a more specific and more realistic situation in which the dynamics are influenced not only by the network structure and the relevant parameters but also by the intrinsic properties of every single node in the network, such as e.g., social status. In other words, we study the interplay between structure, dynamics and exogenous node characteristics and how these complex interactions influence the process of consensus building. To the best of our knowledge, this is the first study that analyzes the effect of all those three aspects on opinion dynamics.

Approach & methods. In the field of statistical physics [5], opinion dynamics are commonly studied by applying mathematical models and analytic approaches. To make these complex problems tractable for mathematical analysis, researchers make simplifications, such as presenting opinions as sets of numbers, ignoring the network structure (a typical approach from e.g. mean-field theory) and neglecting the individual differences between nodes. Simplifications narrow the scope of research down to theoretical models, which typically do not consider empirical data. Even so, statistical physics constitutes important basics for the state-of-the-art research on social dynamics in collaboration networks. In this paper, we build upon these basics.

In our work, we take a computational approach and analyze opinion dynamics by simulating the diffusion of opinions in empirical collaboration networks (specifically, we study

datasets from a Q&A site StackExchange). In our simulations, we consider the network structure, apply a set of simple rules for opinion diffusion and take into account people’s individual differences (e.g., their social status). In particular, we simulate scenarios of peer interactions in empirical datasets assuming that the status theory holds and observe the consequences. We model the dynamics of opinion spreading by adapting a well-known *Naming Game* model [7] and extending it by incorporating a mechanism to configure the degree of the influence of social status on the network dynamics. We termed this mechanism the *Probabilistic Meeting Rule*. Through parametrization, we are able to explore various scenarios from the opposite sides of the spectrum: (i) we can completely neglect the status by allowing any two individuals to exchange their opinions regardless of their social status (an *egalitarian* society) [8], (ii) we can have opinions flowing only in one direction – from individuals with a higher social status to those with a lower social status (a *stratified* society) [9], (iii) we can probabilistically model any situation in between these two extreme cases, i.e. a case in which opinions are very likely to flow from individuals with a higher social status to those with a lower social status but with small probability they can also flow into the other direction (a *ranked* society) [9].

Contributions. The main contributions of our work are twofold. Firstly, with our paper we contribute to the field of opinion dynamics *methodologically*. Secondly, with our work we also make an *empirical* contribution.

Our methodological contribution can be summarized as follows. To model various scenarios of how social status may influence the opinion dynamics, we have invented the Probabilistic Meeting Rule (see subsection II-B) and extended a standard Naming Game model with that rule. The extension is flexible and may reflect a variety of interesting scenarios, such as the emergence or disappearance of social classes in collaboration networks.

From the empirical point of view, we made a much-needed contribution to the limited body of research on Naming Game and empirical data [10] and obtained very interesting empirical experimental results. For example, based on the status theory it can be expected that consensus can be reached faster when social status plays a role. However, our results only partially confirm this expectation. In particular, if an opinion flows only in one high-to-low status direction, opinions do not converge at all since there are always a few people who do not adopt the common opinion from the network (cf. Figure 1). However, with only a low influence of social status convergence is reached faster than with no status at all (as in a standard Naming Game). These results suggest that finding the optimal process of consensus reaching is a tuning act of how to integrate social status in the opinion dynamics.

II. METHODOLOGY

A. Naming Game

Naming Game [7, 11, 12, 13, 14] is a networked agent-based topology, in which agent-to-agent interactions take place based on predefined gaming rules. In particular, agents exchange their opinions and try to reach a consensus about the name of an unknown object. When all agents in the

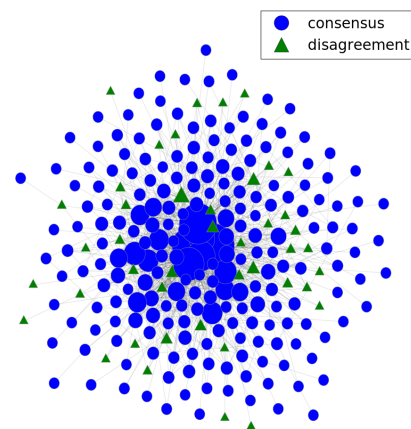


Figure 1: **Consensus building in a network.** Blue (circle) nodes reached consensus (have a single common opinion) whereas green (triangle) nodes did not adopt a common opinion (in a typical case they kept two or more opinions).

network agree on the name, the network is considered to have established a common opinion.

Agents in the game are represented as nodes of a network and edges between two agents allow them to interact with each other. Names are represented with an inventory of words and each agent has her own inventory to store the words. Technically, an inventory is a set, i.e., a bag of words. In the initial state, the inventories are empty. Two random adjacent agents are chosen in each simulation step to interact through a meeting, one agent is declared as a speaker and the other as a listener. In the course of the meeting, the speaker selects a word from her inventory and communicates it to the listener (note that if the speaker’s inventory is empty, a new unique word is created and stored in the inventory). After communicating the word to the listener, two scenarios are possible (see Figure 2):

- 1) the word is not in the listener’s inventory – the word is added to listener’s inventory,
- 2) otherwise, both speaker and listener agree on that word and remove all other words from their inventories – they agree on the selected word.

B. Naming Game and Social Status

We modify the Naming Game to account for social status. As before, the agents are represented as network nodes, edges denote whether two agents can interact or not and names (opinions) are represented as word inventories.

The first difference between our model and a standard Naming Game is the simulation initialization. We initialize the inventories with a given number of selected words from a given vocabulary. The words are selected (with replacement uniformly) at random from the vocabulary. This results in an initial state where each opinion occurs with the same probability.

Secondly, we adopt the social status that governs how agent interactions are turned into meetings – not every agent interaction is turned into a meeting. During each interaction

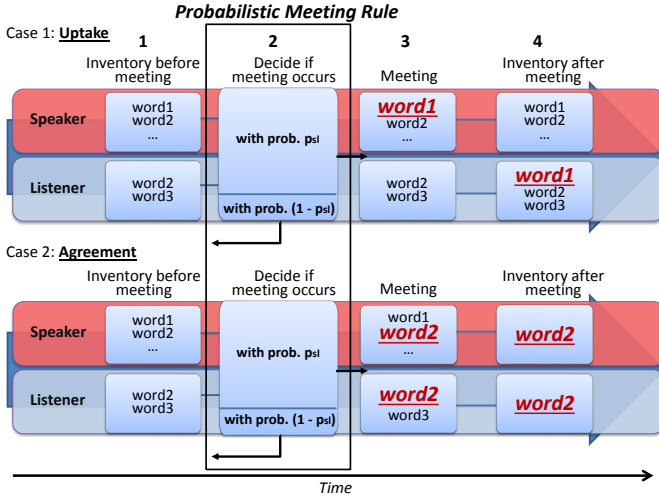


Figure 2: **Naming Game meeting.** The classical Naming Game consists of steps 1, 3 and 4, whereas our extension also includes the step 2. In step 2 we decide whether the meeting between two agents occurs by evaluating Probabilistic Meeting Rule (equation 1). For illustration, consider a ranked society with stratification factor $\beta = 0.0001$. *Example 1:* Speaker’s status $s_s = 101$ and listener’s status $s_l = 7967$. The meeting probability evaluates to $p_{sl} = 0.45$. We then draw a number from $[0, 1]$ uniformly at random, e.g., 0.93 and compare it with p_{sl} – the meeting does not take place. *Example 2:* Let $s_s = 576$ and $s_l = 865$, which leads to the meeting probability $p_{sl} = 0.97$. We again draw a random number from $[0, 1]$, e.g., 0.77 – in this case the meeting takes place. If the meeting takes place two scenarios are possible. 1) If the speaker transmits a word (red) that is unknown by the listener, the listener adds it to her inventory (*uptake*). 2) If the word chosen by the speaker is also known to the listener, they both agree on this word. In this case they both remove all other words from their inventories and keep only the transmitted one (*agreement*).

a random agent and a random neighbor are chosen to have a meeting. Then, the speaker and the listener are assigned randomly. Based on the difference between the speaker’s and the listener’s statuses, we randomly decide if the meeting occurs.

To decide if a meeting takes place, we introduce the Probabilistic Meeting Rule. Basically, the Probabilistic Meeting Rule is a function that takes the agents’ social statuses as input and, based on the difference between the speaker’s and listener’s status, calculates the probability of the meeting taking place. The rule is defined by the following equation:

$$p_{sl} = \min(1, e^{\beta(s_s - s_l)}), \quad (1)$$

where s_s is the speaker’s status, s_l is the listener’s status and $\beta \geq 0$ is the *stratification factor*. The stratification factor β , which can be viewed as a measure of conformance to the agent’s social status, is a tuning parameter in our model. The above equation results in the following probabilities. If the speaker’s status is higher than the listener’s status, p_{sl} has the value of 1, i.e., such a meeting always takes a place. If




Figure 3: **Naming Game and social status.** The application of the *Probabilistic Meeting Rule* and the emergence of social classes based on the stratification factor β are illustrated for the English dataset (see Section III). $\beta = 0$ indicates an *egalitarian* society, in which each agent can meet every other agent (depicted by the black bar). With an increase in β , our society becomes more conservative (as represented with the blue bar) and becomes a *ranked* society. Already at $\beta = 1$ we observe a two-class society (red bar), i.e., a *stratified* society.

the opposite is true, various scenarios are possible, depending on the value of the stratification factor. For example, $\beta = 0$, indicates an *egalitarian* society and p_{sl} is always equal to 1. However, if we slowly increase the stratification factor, p_{sl} will start to decay and in general will take a value between 0 and 1, which signifies a *ranked* society (see the running example in Figure 2). If we continue to increase β , we will soon (because of the exponential term in the equation) reach a situation where p_{sl} for all practical matters is equal to 0. In other words, we have reached a *stratified* society where meetings take place only if the speaker’s status is higher than the listener’s status but never in the opposite case.

The application of our Probabilistic Meeting Rule to one of our datasets (English) is depicted in Figure 3. The probability of a meeting taking place is shown in correlation with the percentage of pairs of agents participating in that meeting. The above mentioned scenarios are represented as follows: *egalitarian* society with $\beta = 0$ – black bar (no texture), *ranked* society with $\beta = 0.0001$ – blue bar (line texture) and *stratified* society with $\beta = 1$ – red bar (star texture).

III. DATASETS AND EXPERIMENTS

A. Datasets

In our experiments, we use datasets from a Q&A site (StackExchange¹), in which users collaborate, ask questions and give answers on particular problems. After an iterative discussion process users exchange their opinions, find solutions to a problem and agree on the best suggested solutions [15].

¹<http://stackexchange.com/>

Such Q&A sites have a reputation system which rewards users via reputation scores based on their contributions [16, 17]. Based on the policies of this reputation system, users get appropriate reputation scores for giving good answers, asking good questions or for voting on questions/answers of other users. It is evident that high reputation users contribute high quality answers [16]. We expect that high reputation users also demonstrate high convincing power during the agreement process, influencing opinions of other (low reputation) users. In our experiments, we apply reputation scores as a proxy for the social status and these two terms are used interchangeably throughout the paper. The StackExchange platform does not indicate associations between users or friendship links. For that reason, we turn our attention to collaboration networks which we extract by analyzing co-posting activities of users in order to have social ties between them [17, 18, 19]. In Q&A sites, a co-posting activity between two users refers to a scenario under which two users comment on the same post. Thus, if two users contributed in any way to a same post, they are connected via an edge in the collaboration network. We analyze the following StackExchange language datasets: French, Spanish, Chinese, Japanese, German and English. They are available for downloading for research purpose from the StackExchange dataset archive.

B. Datasets Statistics

The details of our empirical networks (derived from the above-mentioned datasets) and their properties are shown in Table I, with the number of nodes (n), number of edges (m), mean (μ), median ($\mu_{1/2}$), standard deviation (σ) of the reputation scores, assortativity coefficient (r) and modularity (Q).

Among our datasets, the English network is the largest one with 30,656 nodes and 192,983 edges, whereas the French is the smallest one with 1,478 nodes and 6,668 edges in the network. The German, Japanese, Chinese and Spanish networks lie in between the English and French networks in terms of network size.

A negative assortativity coefficient r [20] indicates a negative correlation between reputation scores over the network edges. In other words, users with lower reputation scores are more likely to connect to users with higher reputation scores. In particular, a typical post in our datasets has many users with low scores, e.g., who post a question, and only a few or even only a single user with a high score, e.g., who answers

Table I: **StackExchange language datasets.** Description of StackExchange datasets with the number of nodes (n), number of edges (m), mean (μ), median ($\mu_{1/2}$) and standard deviation (σ) of the reputation scores, assortativity coefficient (r) and modularity (Q).

Dataset	n	m	μ	$\mu_{1/2}$	σ	r	Q
French	1,478	6,668	298	111	1,273	-0.23	0.31
Spanish	1,584	6,908	196	101	554	-0.19	0.38
Chinese	1,985	8,556	160	61	477	-0.15	0.41
Japanese	2,069	11,155	328	77	1,535	-0.23	0.34
German	2,316	12,825	285	103	1,219	-0.16	0.32
English	30,656	192,983	199	48	1,654	-0.19	0.33




Figure 4: **Distribution of reputation scores.** The plot on the right shows the heterogenous distribution of reputation scores in the English network. The plot on the top presents the heterogenous distribution of node degrees. In the middle, the scatter plot of reputation scores vs. node degrees is shown. The Pearson correlation coefficient between the degree and the reputation score is 0.88. All other datasets have comparable distributions and correlation coefficients.

the question. This finding is in line with the assumptions from the social status theory. The Chinese network has the lowest assortativity coefficient among our networks indicating that in this network there is a smaller chance of connection with a dissimilar reputation score. The Japanese and French networks have the highest assortativity coefficient.

The modularity score is a measure of strength of the community structure in a network. A high modularity score indicates the existence of strong communities in the network, while a low modularity score means that the community structure is not that strong [21]. In our networks, we observe low modularity values corresponding to a very weak or almost nonexistent community structure. As previously shown in a network without communities, in general Naming Game converges quickly to a single opinion [7].

The distribution of reputation scores and node degrees resemble a heterogenous distribution for all networks, which indicates that the majority of users in our collaboration networks have low reputation scores. Figure 4 shows the English network, in which the correlation between the reputation scores and the node degrees is a linear correlation with a Pearson correlation coefficient of 0.88. All other datasets have comparable properties.

C. Simulations

In our experiments, we simulate Naming Game extended with the Probabilistic Meeting Rule. The simulation framework is provided as an open source project². Our experiments consist of the following steps:

- 1) Depending on the network size, we define the number of user interactions (iterations) for the simulations. We perform 2 million interactions for the largest network (English) and 1 million interactions for the five other networks.

²https://github.com/floriangeigl/reputation_networks




Figure 5: **Inventory size evolution.** Mean values of the agent’s inventory size in relation to the number of interactions for English (a), German (b) and Japanese (c) networks. In an *egalitarian* society ($\beta = 0$) a common opinion is reached and the convergence rate is fast. In a *stratified* society ($\beta = 1$), the opinions do not converge (mean number of opinions lies between 1 and 2). *Ranked* societies (e.g. $\beta \in \{0.0001, 0.00013, 0.00016\}$) also reach a common opinion with the highest convergence rate. Thus, the consensus building depends on the status but in a non-obvious way, indicating that there is a specific setting at which the influence of the social status reaches the optimal state.

- 2) We investigate various values of the stratification factor β . For all networks, we perform simulations with the following values of β : (0, 0.00001, 0.0001, 0.00013, 0.00016, 0.0002, 0.001, 1).
- 3) During the simulations, we store important information such as the appearance of agents as listeners/speakers, their participation in overall interactions versus successful meetings and the evolution of the agent’s inventory size.
- 4) Each agent’s inventory is initialized with a fixed number of three opinions (represented through numbers from 0 to 99). These opinions are selected uniformly at random from a bag of opinions to ensure that each opinion occurs with the same probability.

IV. RESULTS AND DISCUSSION

Figure 5 summarizes the results of our experiments by depicting the agent’s inventory size as a function of the simulation progress. We show results for three largest networks (i.e., English, German, and Japanese). The simulation results for the three remaining networks (French, Spanish, Chinese) are comparable to the largest three.

In the case of *egalitarian* society ($\beta = 0$), the networks converge to a single opinion. This is in line with the previous experiments with Naming Game – in networks without a strong community structure we always reach a consensus. In the case of *stratified* society we do not observe convergence – consensus cannot be reached. This seems slightly counter-intuitive – an intuition would be that consensus building would benefit from the presence of agents with a high social status and their influence on agents with a lower social status.

Finding 1: Opinion dynamics in collaboration networks are affected by the individual’s social status. If, due to the social status, opinions flow only in the high-low direction, the consensus building process is disturbed and consensus cannot be achieved, as opposed to when the status does not play any role at all.

The simulation results for *ranked* societies indicate that the impact of the social status on opinion dynamics is a complex one. In all our networks, we observe the following situation. By starting at $\beta = 0$ and slowly increasing the stratification factor, we are at first still able to reach consensus. Moreover, the convergence rate increases with a slightly increased stratification factor (cf. Figure 5 for e.g. stratification factor 0.0001, 0.00013 and 0.00016). However, by further increasing the stratification factor, we reach a tipping point after which a further increase of the stratification factor results firstly in slower convergence rates before we again reach a state of no convergence at all (within e.g. *stratified* society and $\beta = 1$). The optimal value for the stratification factor is very similar in all networks and lies in range between 0.0001 and 0.00016.

Finding 2: The relation between the opinion dynamics and the stratification factor of a society is intricate. Low values of stratification tend to favor consensus reaching – in such societies, consensus is always reached at a very fast convergence rate, which is higher than in *egalitarian* societies. However, if the stratification factor becomes too large, the consensus reaching process is hindered.

To further analyze these findings, let us investigate in more details the direction and intensity of opinions flow in our networks. To that end, we separate the agents into two classes: high (agents with the status above 90th percentile) and low (agents below 90th percentile) class. All reputation distributions are skewed to right and resemble a heterogenous distribution and the division into classes results in a reputation boundary of e.g. 364 for English network with all agents having reputation above 364 belonging to the high class and all agents below 364 belonging to the low class (for comparison the highest reputation score in English dataset is around 37,000). All other networks are comparable to English and our analysis produces similar results. For that reason, we henceforth discuss only the English network.

An important question is what happens when agents interact and how the Probabilistic Meeting Rule evaluates depending on the classes of agents participating in a meeting. In other words, we want to investigate the fraction of interactions that turn into a successful meeting (which consequently results in an opinion flow and increases the likelihood of two agents agreeing on a single word). We therefore classify each interaction according to the agent classes into four possible pairs: (i) low-low, (ii) low-high, (iii) high-low, (iv) and high-high where the first class corresponds to the speaker's class and the second corresponds to the listener class. Figure 6 depicts the fractions of successful meetings among all interactions in the English dataset for three values of the stratification factor—*egalitarian* society with $\beta = 0$, *stratified* society with $\beta = 1$ and *ranked* society with the optimal value for this dataset $\beta = 0.0001$.

In the case of *stratified* society (red bars), opinions flow without restrictions only in high-low direction. Thus, the agents with a higher status can pass over their opinions to the agents with a lower status. The flow in the opposite direction is completely prohibited and therefore agents with a lower status cannot influence the opinions of the agents with a higher status. However, the Probabilistic Meeting Rule in this case is so strict and prohibitive that it greatly inhibits the opinion flow within the agents of the same status, i.e. high-high and low-low pairs. Because of the skewed nature of the reputation distributions, the inhibition in the low-low group (which is considerably larger than the high-high group) is more severe – the agents with a lower social status cannot efficiently exchange their opinions with each other and must rely on the agents with a higher social status to inject opinions into the low group by meeting each low agent separately. Since there are few high and many low status agents, consensus is never reached.

On the other hand, in the case of *egalitarian* society (black bars), opinions flow without any restrictions in all directions. This results in the convergence of opinions and a rather fast convergence rate. However, the convergence rate is slightly slower as compared to the optimal case (*ranked* society). In our opinion, the explanation for this phenomenon lies in the dynamics of the low-high group meetings. Since everybody can impose her opinion onto everybody else, low status agents very often change the opinions of high status agents. Thus, low status agents increase the variance in the inventories of high status agents and they need additional meetings to eliminate these opinions. This results in slower convergence rates.

A particular dynamics of low-high meetings also explains faster convergence rates in *ranked* societies (blue bars). In this case, the opinion flow from the agents of low status to the agents of high status is strongly slowed down. Therefore, the disturbances in the opinions of high status agents are not substantial any more. On the other hand, as opposed to the *stratified* society, the opinion flow within the low-low group is not impaired at all. Thus, the injected opinions from the high status agents can be diffused among the low status agents themselves without need to address each low status agent separately. This, combined with the reduced disturbances flowing from low to high status agents, results in optimal opinion convergence rates.

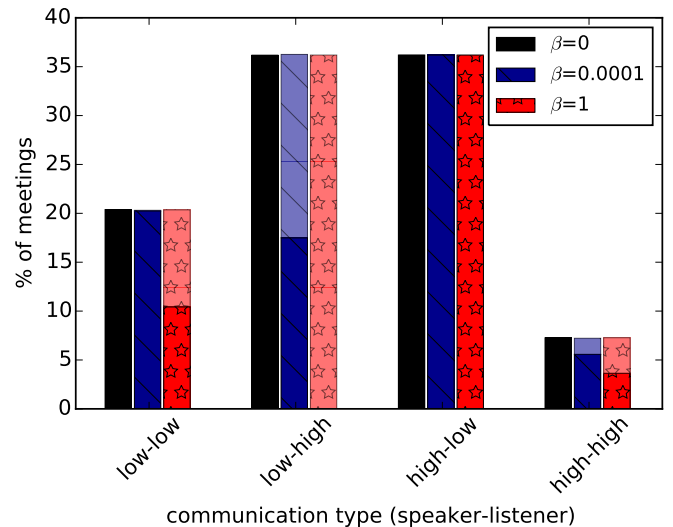


Figure 6: **Participation of agents in meetings across status groups.** The percentage of interactions resulting in meetings as a function of reputation classes in the English network. The high class comprises agents with the status above 90th percentile and the low class all other agents. In the *stratified* society (red bars), a common opinion cannot be reached because the meeting rule is so strict that even communications between low agents (low-low pairs) are severely impaired. In the *egalitarian* society (black bars), the convergence is slower because low status agents disturb high status agents by inflicting their opinion upon them (low-high pairs). In the *ranked* society (blue bars), the optimal convergence is achieved because low status agents can diffuse opinions among themselves (low-low pairs). At the same time, since the communications between low and high status agents are inhibited (low-high pairs), low status agents' opinions cannot disturb those of high status agents.

Finding 3: The optimal convergence of opinions is achieved when low status agents can exchange their opinions among themselves without any restrictions. In addition, there must be a barrier that prohibits low status agents to inflict their opinions on high status agents so that disturbances in the opinions of high status agents are minimized.

V. RELATED WORK

At present, we identify three main lines of research related to our work: opinion dynamics, social status theory and naming game.

A. Opinion Dynamics

Opinion dynamics is a process characterized with a group of individuals reaching a consensus (i.e., the majority of a group share the same opinion). In opinion dynamics, the focus is on modeling the opinion state of an individual in particular and a population in general. Opinion dynamics has been tackled in the past in the context of statistical physics (see e.g., [5], [22]). As discussed in [5], if opinion dynamics is viewed from a perspective of statistical physics, an individual is analogous to a particle with properties that may or may not change

over a period of time. Thus, the social process of interaction among individuals can be designed as a mathematical model that represents a change in the local and global state of an individual and a group. One of the examples of such a process is the Naming Game model, a variant of which we are using in our work, that models how individuals behave during a meeting and exchange their opinions. In our experiments, the meeting process is further enhanced by taking reputation scores of individuals into account. Constraining the system to favor high reputation nodes resulted in reaching consensus later as compared to an unconstrained model.

B. Social Status Theory

Research on how the position and status of a node influence a network is mostly carried out in the context of network exchange theory (e.g., [1, 2, 3]). This theory states that connections and a position in a network lead to a power condition that is based on how the nodes are connected and which position they take in the network [2]. For example, in [1], researchers differentiate between weak and strong powers network in terms of node positions and network properties. The authors give a theoretical extension to the network exchange theory to explain why in sparsely connected networks a stronger power effect is observed than in densely connected networks. They found that in densely connected networks, weak position nodes have an advantage since they have a higher connectivity, which enables them to short circuit the structural advantages of strong position nodes. This is related to our work, as we concentrate on investigating how the reputation of a node in a network affects the spread of opinion that leads to establishing consensus in the network. Also, we define various classes of nodes based on reputation and determined how their interaction affects their overall process of consensus building.

C. Naming Game

The Naming Game has been introduced in the context of linguistics [23] and the emergence of a shared vocabulary among agents [7] with the aim to demonstrate how autonomous agents can achieve a global agreement through pair-wise communications without central coordination [24]. With that regard, we present a selection of variations of the Naming Game that are relevant to our work.

Similarly to our approach, the work of Brigatti et al. [25] describes a variation of the Naming Game that incorporates the agents' reputation scores. In the beginning, reputation is randomly distributed (Gaussian distribution) among the agents. Successful communication increases the agents' reputation and during each iteration, the agent with a higher reputation score acts as a teacher and the one with the lower score as a learner. The main difference from our work is that in [25], they use synthetic data for the simulations and that the assigned reputation scores are random numbers that change during iterations. In our work, we employed empirical collaboration networks from StackExchange with reputation scores that were assigned by the community. As opposed to the work of [25] where there is an open-ended game with unlimited number of words, the inventory of our agents consists of predefined sets of three opinions.

Other examples for the Naming Game variations include the work of Li et al. [26] who studied the impact of spatial structures, e.g., geographical distances, have on meetings between individuals in a network, and [27], who proposed a Naming Game that follows an asymmetric negotiation strategy and investigated the influence of hub effects on the agreement dynamics with specific focus on how quickly consensus could be achieved. Each agent in the network is assigned a weight defined by the agent's degree and a tuneable parameter α . During iteration, two nodes are randomly selected and based on their degree and the configuration of the parameter α , they are either the speaker or the listener (i.e., if $\alpha > 0$, high degree agents have more chances to be speakers and vice versa). This way, the dynamics of the game can be investigated in light of the varying influence of high degree agents. Our work is somewhat related to this work as we also use a parameterized probability function to define the probability of a meeting taking place between two nodes, in our case depending on their reputation score. The main difference to our work is that agents' selection is unbiased and empirical data with explicitly provided reputation scores are used.

The diffusion of opinions across networks and the potential of reaching consensus are strongly influenced by the availability of communities and, specifically, by the presence of strong community boundaries [28]. To investigate this effect, Lu et al. [28] assigned a group of nodes in a network as a committed fraction, i.e., nodes that are not influenced by other nodes in a network and don't ever change their opinion. In our dataset, however, no strong community structures are present.

VI. CONCLUSION AND FUTURE WORK

Understanding opinion dynamics and how consensus is reached in social networks has been an open and complex challenge for our community for years. In this work, we addressed a sub-problem related to this challenge by investigating a specific case of collaboration networks in which individual nodes have a certain social status.

To that end, we presented an extension (Probabilistic Meeting Rule) to the standard Naming Game model of opinion dynamics and computationally analyzed six large empirical collaboration networks. This extension constitutes our methodological contribution. Apart from this methodological contribution to the field, we have experimentally evaluated various real world scenarios such as the emergence and disappearance of social classes in collaboration networks. From the empirical point of view, our investigations revealed interesting facts about the influence of social status on the diffusion of opinions. Our main finding indicates that social status strongly influences the opinion dynamics in a complex and intricate way. More specifically, weakly stratified societies reach consensus at the highest convergence rate, whereas completely stratified societies do not reach consensus at all. The most important issue in this process is related to low status agents and how their communication is controlled. In particular, the optimal convergence is achieved when (i) low status agents are allowed to freely exchange opinions between themselves (since this reduces the need for high status agents to interact with low status agents) and (ii) simultaneously there is a communication barrier reducing the number of interactions of low status agents

towards high status agents (since this reduces the variance in opinions of high status agents).

Limitations. In our opinion, our work has the following limitations. Firstly, we represent social status with a single number – for certain scenarios this representation may be too simplistic. For example, people often play different roles in social networks and a non-simple interplay between the roles and status may exist. Secondly, a more finely grained classification of agents into various groups (e.g., low, mid and high groups or even finer divisions) may shed more light on the opinion dynamics. Finally, in our work we consider only static snapshots of networks and reputation scores. However, not only opinions but also networks are dynamic, as new agents may arrive to the network, new edges may form and inactive edges may disappear from the network. Moreover, reputation itself is very dynamic and depends on the agent’s activity and the current perception of an agent by her peers.

Future work. In our future work, we plan to address some of the limitations of our current work and extend our approach and experiments to other scenarios. For example, one interesting avenue for further research are the networks with a strong community structure. As communities tend to slow down the consensus reaching process, it would be interesting to investigate how status and/or network structure can be adjusted to support the process. Apart from social status, the influence of trust is of utmost importance in various social systems and in particular in social media. Thus, adapting the presented approach to analyzing how trust relates to opinion dynamics is another promising research direction for the future.

ACKNOWLEDGMENTS

This work is supported by the Know-Center, the EU funded project Learning Layers (Grant Agreement 318209) and the Austrian Science Fund (FWF) research project ”Navigability of Decentralized Information Networks” (P 24866-N15). The Know-Center is funded within the Austrian COMET Program - Competence Centers for Excellent Technologies - under the auspices of the Austrian Ministry of Transport, Innovation and Technology, the Austrian Ministry of Economics and Labor and by the State of Styria. COMET is managed by the Austrian Research Promotion Agency (FFG).

REFERENCES

- [1] B. Markovsky, J. Skvoretz, D. Willer, M. J. Lovaglia, and J. Erger, “The seeds of weak power: An extension of network exchange theory,” *American Sociological Review*, vol. 58, no. 2, pp. 197–209, 1993.
- [2] H. A. Walker, S. R. Thye, B. Simpson, M. J. Lovaglia, D. Willer, and B. Markovsky, “Network exchange theory: Recent developments and new directions,” *Social Psychology Quarterly*, vol. 63, no. 4, pp. 324–337, 2000.
- [3] D. Willer, *Network Exchange Theory*. Westport, CT: Praeger, 1999.
- [4] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins, “Propagation of trust and distrust,” in *Proceedings of the 13th International Conference on World Wide Web*, ser. WWW ’04. New York, NY, USA: ACM, 2004, pp. 403–412.
- [5] C. Castellano, S. Fortunato, and V. Loreto, “Statistical physics of social dynamics,” *Rev. Mod. Phys.*, vol. 81, pp. 591–646, May 2009.
- [6] B. Latané, “The psychology of social impact,” *American Psychologist*, vol. 36, pp. 343–365, 1981.
- [7] A. Baronchelli, M. Felici, E. Caglioti, V. Loreto, and L. Steels, “Sharp transition towards shared vocabularies in multi-agent systems,” *J. Stat. Mech.*, no. P06014, 2006.
- [8] R. Arneson, “Egalitarianism,” in *The Stanford Encyclopedia of Philosophy*, summer 2013 ed., E. N. Zalta, Ed., 2013.
- [9] M. Weber, *The theory of social and economic organization*. The Free Press, 1964.
- [10] Y. Gao, G. Chen, and R. H. M. Chan, “Naming game on networks: Let everyone be both speaker and hearer,” *CoRR*, vol. abs/1310.4169, 2013.
- [11] A. Baronchelli, L. Dall’Asta, A. Barrat, and V. Loreto, “Strategies for fast convergence in semiotic dynamics,” in *Artificial Life X*, L. M. Rocha and et al., Eds. MIT Press, 2005, pp. 480–485.
- [12] A. Baronchelli, L. Dall’Asta, A. Barrat, and V. Loreto, “Topology-induced coarsening in language games,” *Phys. Rev. E*, vol. 73, p. 015102, Jan 2006.
- [13] L. Dall’Asta, A. Baronchelli, A. Barrat, and V. Loreto, “Agreement dynamics on small-world networks,” *EPL (Europhysics Letters)*, vol. 73, no. 6, p. 969, 2006.
- [14] L. Dall’Asta, A. Baronchelli, A. Barrat, and V. Loreto, “Nonequilibrium dynamics of language games on complex networks,” *Phys. Rev. E*, vol. 74, p. 036105, Sep 2006.
- [15] Y. R. Tausczik, A. Kittur, and R. E. Kraut, “Collaborative problem solving: A study of mathoverflow,” in *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, ser. CSCW ’14. New York, NY, USA: ACM, 2014, pp. 355–367.
- [16] D. Movshovitz-Attias, Y. Movshovitz-Attias, P. Steenkiste, and C. Faloutsos, “Analysis of the reputation system and user contributions on a question answering website: Stackoverflow,” in *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, ser. ASONAM ’13. New York, NY, USA: ACM, 2013, pp. 886–893.
- [17] A. Halavais, K. Kwon, S. Havener, and J. Striker, “Badges of friendship: Social influence and badge acquisition on stack overflow,” in *System Sciences (HICSS), 2014 47th Hawaii International Conference on*, Jan 2014, pp. 1607–1615.
- [18] J. Tang, T. Lou, and J. Kleinberg, “Inferring social ties across heterogeneous networks,” in *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining*, ser. WSDM ’12. New York, NY, USA: ACM, 2012, pp. 743–752.
- [19] L. Adamic and E. Adar, “Friends and neighbors on the web,” *Social Networks*, vol. 25, pp. 211–230, 2001.
- [20] M. E. J. Newman, “Mixing patterns in networks,” *Phys. Rev. E*, vol. 67, p. 026126, 2003.
- [21] M. Newman, “Modularity and community structure in networks,” *Proceedings of the National Academy of Sciences*, vol. 103, no. 23, pp. 8577–8582, 2006.
- [22] G. Iniguez, J. Török, T. Yasseri, K. Kaski, and J. Kertesz, “Modeling social dynamics in a collaborative environment,” *EPJ Data Science*, Jan 2014.
- [23] L. Dall’Asta, A. Baronchelli, A. Barrat, and V. Loreto, “Nonequilibrium dynamics of language games on complex networks,” *Phys. Rev. E*, vol. 74, p. 036105, 2006.
- [24] W. Zhang, C. C. Lim, G. Korniss, and B. K. Szymanski, “Opinion dynamics and influencing on random geometric graphs,” *Sci. Rep.*, vol. 4, 07 2014.
- [25] E. Brigatti, “Consequence of reputation in an open-ended naming game,” *Phys. Rev. E*, vol. 78, no. 4, p. 046108, Oct. 2008.
- [26] R.-R. Liu, W.-X. Wang, Y.-C. Lai, G. Chen, and B.-H. Wang, “Optimal convergence in naming game with geography-based negotiation on small-world networks,” *Physics Letters A*, vol. 375, no. 3, pp. 363 – 367, 2011.
- [27] H.-X. Yang, W.-X. Wang, and B.-H. Wang, “Asymmetric negotiation in structured language games,” *Phys. Rev. E*, vol. 77, p. 027103, Feb 2008.
- [28] Q. Lu, G. Korniss, and B. Szymanski, “The naming game in social networks: community formation and consensus engineering,” *Journal of Economic Interaction and Coordination*, vol. 4, no. 2, pp. 221–235, 2009.