

# Engineering Agreement: The Naming Game with Asymmetric and Heterogeneous Agents

Jie Gao,<sup>1</sup> Bo Li,<sup>2</sup> Grant Schoenebeck,<sup>2</sup> Fang-Yi Yu<sup>2</sup> \*

<sup>1</sup>Stony Brook University; <sup>2</sup>University of Michigan

## Abstract

Being popular in language evolution, cognitive science, and culture dynamics, the Naming Game has been widely used to analyze how agents reach global consensus via communications in multi-agent systems. Most prior work considered networks that are symmetric and homogeneous (e.g., vertex transitive). In this paper we consider asymmetric or heterogeneous settings that complement the current literature: 1) we show that increasing asymmetry in network topology can improve convergence rates. The star graph empirically converges faster than all previously studied graphs; 2) we consider graph topologies that are particularly challenging for naming game such as disjoint cliques or multi-level trees and ask how much extra homogeneity (random edges) is required to allow convergence or fast convergence. We provided theoretical analysis which was confirmed by simulations; 3) we analyze how consensus can be manipulated when stubborn nodes are introduced at different points of the process. Early introduction of stubborn nodes can easily influence the outcome in certain family of networks while late introduction of stubborn nodes has much less power.

## 1 Introduction

The analysis of shared conventions in multi-agent systems and complex decentralized social networks has been the focus of study in several diverse fields, such as linguistics, sociology, cognitive science, and computer science. The problem of how such conventions can be established, from among countless options, without a central coordinator has been addressed by several disciplines (Nowak and Krakauer 1999; Brighton and Kirby 2001). Among them, the multi-agent models and mathematical approaches gain the most attention by accounting for both the network topology and opinion change over time (Steels 2005; Nowak, Plotkin, and Jansen 2000; Baronchelli, Loreto, and Steels 2008; Pickering and Lim 2016; Franks, Griffiths, and Jhumka 2013). It has been shown that the emergence of new political, social, economic behaviors, and culture transmission are highly dependent on such convention dynamics (Backstrom

et al. 2006; Hurford 1989; Nowak, Plotkin, and Krakauer 1999).

In order to analyze the social dynamics in multi-agent systems in depth, we focus on one stylized model, the Naming Game, in which agents negotiate conventions through local pairwise interactions (Steels 1995; Baronchelli et al. 2006a). The Naming Game captures the generic and essential features of an agreement process in networked agent-based systems. Briefly speaking, when two agents wish to communicate, one agent, the speaker, randomly selects one convention from her list of current conventions and uses this convention to initiate communication with the listener. If the listener recognizes that convention, both the speaker and listener purge their lists of current conventions to only include that “successful” convention. If the listener does not recognize that convention, she adds it to her list of known conventions.

This simple model is able to account for the emergence of shared conventions in a homogeneous population of agents. Both simulations and experiments have been conducted on various network topologies. However many key questions, especially those related to asymmetric and heterogeneous agents, remain open. For example: what network topologies enable the fastest convergence? Does community structure help or harm convergence? Does homogeneity or heterogeneity help or harm convergence? How robust are the dynamics to possible manipulations by a small number of agents? Moreover, rigorous theoretical analysis is almost entirely absent in previous work on the Naming Game. In this paper we aim to fill in the literature in the following aspects:

1. We discovered that the star graph empirically converges faster than all previously considered graphs for the Naming Game. This network differs from previously analyzed topologies in that it is not symmetric (vertex transitive). In some sense, it is not too surprising that the star graph, an asymmetric graph, works so well to reach consensus, which is a symmetry breaking problem. Though, from first principles, this is far from obvious, and other asymmetric graphs, for example a multi-level tree, perform extremely poorly.
2. To understand network topologies that inhibit fast convergence of the Naming Game, we study two networks with community structures: agents divided into two dis-

\*J. Gao would like to acknowledge support through NSF DMS-1418255, CCF-1535900, CNS-1618391 and AFOSR FA9550-14-1-0193. The remaining authors are pleased to acknowledge the support of the NSF via AitF-1535912 and CAREER-1452915.

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

connected communities; and a multi-level tree. For the first network, it is clear that it cannot converge to consensus (it is disconnected). We investigate how much inter community communication needs to be added in order to facilitate convergence. Empirically we observe a sharp threshold on the level of inter community communication: above this threshold, fast convergence is guaranteed, and below it the dynamics fail to converge before time out. We give theoretical justifications for this threshold by showing that convergence takes exponentially long if inter community communication is insufficient (below the threshold). For the second network, the multi-level tree, we observe via simulations that it converges exceedingly slowly—we conjecture that it takes exponential time. For this network, we perform the same simulation tests for adding homogeneity and obtain similar results.

We show that with added communication, the community divisions that thwart consensus can be overcome. Perhaps surprisingly, the amount of intercommunity communication required after disagreement is entrenched, is not substantially more than the amount of communication needed to avoid such division in the first place.

3. Finally, we analyze a third way of introducing asymmetry and heterogeneity: including “stubborn” nodes that do not follow the standard Naming Game protocol. Our experimental results suggest the following hypothesis: in some graphs (e.g. cliques) even a small constant (e.g. 5) number of stubborn nodes can assure convergence to a particular name. However, in others networks (e.g. star graphs, grid graphs, Kleinberg’s small world models), the number of nodes required seems to grow with the size of the graph. Additionally, we prove that in a complete graph, manipulation after convergence is much harder than before: there exists a value  $p$  such that if an adversary controls more than a  $p$  fraction of the nodes, consensus results can be easily manipulated; otherwise it takes exponential time to manipulate the consensus.

The results on stubborn nodes have implications for the use of the Naming Game in distributed systems. In Steels and McIntyre (1998) it was assumed that the protocol would be robust to manipulation. We confirmed this claim if the stubborn nodes appear after the system has converged. But in certain networks these protocols are immensely vulnerable to rogue agents that appear from the start.

## Related Work

Baronchelli et al. (2006b) proposed the Naming Game as a simple multi-agent framework that accounts for the emergence of shared conventions in a structured population. One of the most important problems for Naming Game is to understand how fast the global consensus can be reached and what factors affect it. Some research has been conducted to analyze the effect of network topology on the Naming Game dynamics (Dall’Asta et al. 2006). Lu, Korniss, and Szymanski (2009) show via simulations on real-world graphs that communities show speedy convergence of the dynamics. Centola and Baronchelli (2015), using human-subject study,

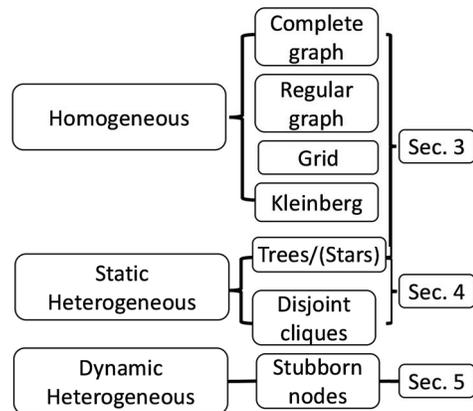


Figure 1: Overview of considered graph structures.

empirically demonstrate the spontaneous creation of universally adopted social conventions and show simple changes in a population’s network structure can greatly change the dynamics of norm formation. Baronchelli et al. (2007) show that finite connectivity, combined with the small-world property, ensures superior performance in terms of memory usage and convergence rate to that of the grid or complete network. Additionally, a dynamically evolving topology of co-evolution of language and social structure has been studied by Gong et al. (2004), for a more complex language game.

One common way to influence the social dynamics and facilitate the converging process towards the consensus is to break the symmetry. Lu, Korniss, and Szymanski (2009) Lu et al. have made use of a special kind of agents called “committed” nodes, who will stick to a preferred opinion without deviating, and show that such agents often reduce the time needed to reach consensus. However, in their work they did not evaluate how these nodes might influence which name was converged upon. Additionally, they did not study how the network topology interacted with stubborn nodes or how robust the communication protocol is.

## 2 Preliminaries

We present here the version of the *Naming Game* introduced in Baronchelli et al. (2006a) in which agents negotiate conventions (names), i.e. associations between forms and meaning. The process stops when all agents reach consensus on a single ‘name.’ The Naming Game is played by agents on a (weighted) graph  $G = (V, E, w)$  and proceeds in steps. At each step  $t$ , each agent  $v$ , is characterized by its inventory (list of names)  $A_t(v) \subseteq \mathbb{S}$ . At time 0 each agent has an initial inventory  $A_0(\cdot)$  which is possibly empty. At each time step  $s = 1, 2, \dots$

- An edge is randomly chosen with probability proportional to its weight; and with equal chance one vertex incident to the edge is considered as the speaker and the other as the listener.
- The speaker  $v$  selects a word  $c$  uniformly at random from its inventory  $A_t(v)$  and sends  $c$  to the listener  $u$ . If the

speaker’s inventory is empty, the speaker invents a new word  $c$  (one that is not in the list of any other agent).

- If the word is in the listener’s inventory,  $c \in A_s(u)$ , the interaction is a “success”, and both the speaker and listener remove all words besides  $c$  from their inventories.
- If the word is not in the listener’s inventory,  $c \notin A_s(u)$ , the interaction is a “failure” and the listener adds  $c$  to its inventory.

The process stops when all the inventories are a singleton of the same name, and we say the process has reached consensus. Notice that the only time a node can have an empty inventory is if it starts that way and has yet to engage in any interaction.

The way in which agents may interact with each other is determined by the topology of the underlying contact network. Here we will introduce the models considered in this paper.

1. *Complete graphs*: all agents are mutual connects.
2. *Regular random graph  $G_{n,k}$*  (see Bollobás (1998)): every node has degree  $k = 8$  and the connection is randomly sample under this constrain.
3. *Kleinberg’s small world model* (Kleinberg 2000): in standard Kleinberg’s model the nodes are on two dimensional grid. Each node  $u$  connects to every other node within Manhattan distance  $p$  as strong ties, and there are  $q$  weak ties which connects to other nodes  $v$  proportional to  $d(u, v)^\alpha$ . In our simulation, the each nodes has 4 strong tie which is  $p = 1$ , and 4 weak ties with  $\alpha = 2$ .
4. *Watts-Strogatz’s small world model* (Watts and Strogatz 1998): the nodes are on one-dimensional ring, and connect to 8 nearest nodes with respect to Manhattan distance, then we rewire the edges of independently with probability 0.5.
5. *Complete bipartite graph* is a bipartite graph such that every pair of graph vertices in the two sets are adjacent. If there are  $p$  and  $q$  graph vertices in the two sets, the complete bipartite graph is denoted  $K(p, q)$ .
6. The *trees* in this paper refer to perfect  $k$ -ary trees with height  $h$ —that is, a rooted tree with  $h$  levels where each node except leaf nodes has exactly  $k$  children and the leaf nodes are all at the level  $h$ . Note that a *star graph* with  $n$  leaves is the complete bipartite graph  $K_{1,n}$ . Alternatively, a star graph can also be defined as rooted tree of branching factor  $n - 1$  with depth 1.

### 3 Networks with Fast and Slow Convergence

In this section we study the convergence rate of various graphs. Here we show that a family of asymmetric graphs, the star graphs, empirically converge faster than previously proposed graphs. Next, we point out, perhaps surprisingly, that the convergence time of a multi-level tree is extremely slow. We will engineer and analyze fast converge versions of trees by adding random edges in Section 4.

We first examine the convergence time for different graphs on a large scale. Here we calculate the time in

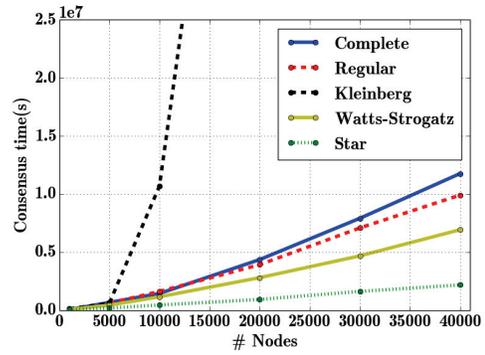


Figure 2: Evaluation of the consensus time for different graphs with size growing until 40000.

terms of the number of communication steps denoted as “s”. We look at complete graphs, random regular graphs ( $G_{n,k}$  graphs), Kleinberg’s small world graphs, Watts-Strogatz graphs, as well as star and tree graphs. Unless mentioned otherwise, we will use the same setting defined above in Section 2. From Figure 2 we can see that the star graph converges the fastest. The tree graph is in fact the slowest. If the tree has two levels with 5000 nodes, after  $10^7$  steps the nodes still cannot reach consensus. Therefore we did not present the consensus time of the tree in the figure. Among the rest of the graphs, the Kleinberg’s small world model is the second slowest, while the other graphs have convergence rate roughly a constant factor of each other.

The network topology’s impact on the Naming Game’s consensus time is fairly intriguing. To better understand the results, let us consider the best and worst topology scenarios for multiple agents to reach consensus. The best (quickest) way to reach consensus is to have a specific node to inform all the other nodes of the name. In other words, it is represented by a star graph and the center node is always the speaker. In the naming game framework, even when the speaker/listener role assignment is uniformly random, the star graph is still the fastest in reaching global consensus. This is partly attributed by the asymmetry inherent in the star graph topology.

To analyze the effect of asymmetry, we simulate the graph morphing from a balanced complete bipartite graph to a star by increasing the number of vertices in the larger side of a complete bipartite graph. Figure 3 shows the converge time for various complete bipartite graphs. Moving to the right in the figure, the graph becomes more asymmetric and we see that the convergence time decreases. Note that at  $m = n$  ( $m/n = 1$ ), this is a balanced bipartite network, and at  $m = 2n - 1$  ( $m/n \approx 2$ ) this is a star graph. This finding is also aligned with the idea that breaking symmetry can improve consensus efficiency for naming game via “stubborn” agents (Lu, Korniss, and Szymanski 2009) (and see Section 5).

On the other hand, the worst graph topology for reaching global consensus is the multi-level tree graph. We hypothesize that this is due to the “community structure” embedded

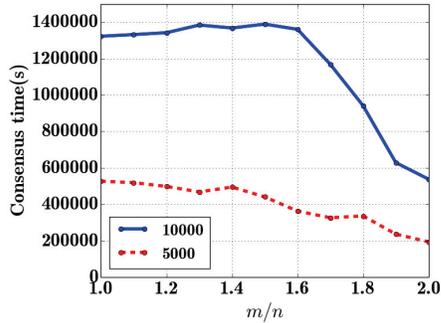


Figure 3: Evaluation for converging time for various complete bipartite graphs  $K_{m,2n-m}$  where  $m$  is the cardinality of the larger partition of vertices.

in the tree that converge fast by themselves. In a two-level tree, the subtrees of the main tree are themselves star graphs. Such community structure enables fast “local” convergence of the dynamics within the communities, but face challenges in reaching global convergence — the communities are trying to influence each other but each community has more internal influence than external influence. This phenomena is the topic of the next section, where we give both empirical and rigorous theoretical analysis.

#### 4 Effects of Community Structure

In this section we study the effects of community structure using two network models, one of them is a dense graph and the other one a sparse graph. The first is a graph of heterogeneous agents divided into two disconnected communities. The simplicity of this model permits theoretical analysis of precisely how and when community structure can exhibit convergence. The second is a multi-level tree introduced in the previous section.

Given a weighted graph  $G$  where the sum of the weights is  $W$  we construct  $\text{Hom}(G, p)$  by adding  $\frac{p}{1-p} \frac{W}{\binom{n}{2}}$  mass to each edge (creating a new edge if it does not exist). This effectively samples the complete graph with probability  $p$  and the graph  $G$  with probability  $1-p$ .

For each network, we first examine the convergence rate of  $\text{Hom}(\cdot, p)$  using simulations. We show that adding a sufficient amount of homogeneity overcomes the heterogeneity. For the first network, we will provide a theoretical analysis which predicts, supports, and explains the empirical results.

##### Disjoint Cliques

Naturally, a graph  $G$  of  $2n$  heterogeneous agents divided into two equally sized disconnected communities will not converge to consensus. As  $p$  increases from 0 toward 1  $\text{Hom}(G, p)$  becomes a network of increasingly interconnected communities.

Additionally, the behavior of the Naming Game depends on the initial states, i.e., the collection of names at these nodes at the beginning. We consider two situations for the initial states. 1) “Empty” start, where all nodes start with

empty lists  $\forall v \in V, A_0(v) = \phi$ . 2) “Segregated” start, in which the two groups have different initial opinions,  $\forall v \in V_1, A_0(v) = \{0\}$  and  $\forall v \in V_2, A_0(v) = \{1\}$ . Clearly it is more challenging for the Naming Game to reach global convergence under the segregated initial state.

**Simulation Results.** Figure 4 (row 1 (a)) shows the convergence time for different values of  $p$  under different initial scenarios on graphs of size  $n$ . For each setting we run the simulation multiple times and plot the time to reach consensus for each run as a dot in the figure. In certain situations it is hard to reach consensus even after a long time. Therefore we set  $10^7$  as the time-out criteria – i.e., if no consensus is reached after  $10^7$  rounds and we stop the simulation. From Figure 4 (row 1 (a)) we can see that when  $p$  is smaller it is harder to reach consensus for all situations. When  $p$  is sufficiently small all situations may hit the timeout condition before consensus is reached. In addition, the threshold of  $p$  which allows this happen is larger for the “segregated” initial setup compared to the empty initial setup. Similarly, for graphs of larger size it is easier to hit the time out condition. When  $p > 0.2$  the time to reach consensus for all situations is small so we chose not to plot it.

To further analyze the naming game behavior when  $p$  is in between  $[0, 0.25]$ , we show in Figure 4 (row 2 (a)) the fraction of trials failing to reach consensus (before timing out) with different values of  $p$ . It is clear that for the empty start initial condition, the game will time out at about  $p = 0.24$ , while for the segregated start case, the game will time out when  $p$  is around 0.26. This threshold value changes with the size of the local community.

Curiously, for the “empty” start, graphs with smaller sizes are more likely to encounter timeouts than their larger counterparts. This may be because the smaller size of each community results in a greater chance of quickly reaching local consensus, which resembles the segregated start scenario. Therefore, it takes longer for graphs with smaller sizes to break the local consensus and escape the so called “stuck” situation.

However, for the segregated start, it immediately starts with the “worst” case setting where the two communities have diverging opinions, so overall it takes longer to leave “stuck” situation compared with graphs of the same size in the “empty” start scenario. Additionally, graphs with larger sizes in the segregated setting more easily encounter a timeout. This may be because larger graphs occasionally time out even if they are not really “stuck” because they take longer to reach consensus in any event.

**Theoretical Analysis.** Next we will analyze the consensus time for the naming game on  $\text{Hom}(G, p)$  where  $G$  has  $2n$  agents divided into two equally sized disconnected communities with segregated start.

**Theorem 4.1.** *Let  $G$  be the disjoint union of two  $n$  cliques, each of size  $n$ . Then for the segregated start naming game on  $\text{Hom}(G, p)$ , there exists a constant  $p_0 \approx 0.110$  such that if  $0 \leq p < p_0$  the expected consensus time is  $\exp(\Omega(n))$ .*

Here we sketch a proof of theorem. A full proof appears in the full paper.

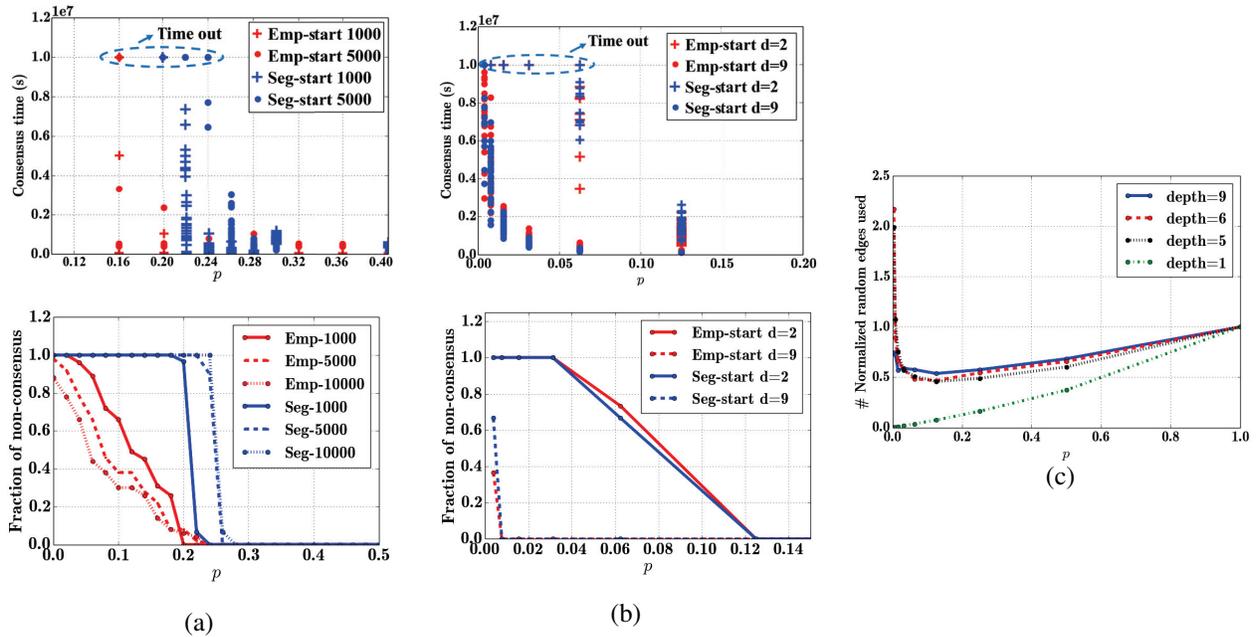


Figure 4: Evaluations of reaching global consensus for different initialization scenarios and sizes of graphs. Row 1: Consensus time for (a) disjoint cliques (b) tree structure; row 2: fraction of nodes failing to reach consensus, based on different probability of random edges  $p$  for (a) disjoint cliques (b) tree structure; (c) normalized number of random edges used for communication as a function of the probability of random edges  $p$ .

To prove this theorem, we formulate the naming game as a nonhomogenous random walk on  $\mathbb{Z}^4$  and relate this non-homogenous random walk to a corresponding autonomous system in  $\mathbb{R}^4$ .

In the segregated start scenario, every node has an initial opinion, therefore no new name will be generated, and nodes inventory will be either  $\{0\}$ ,  $\{1\}$ , or  $\{0, 1\}$ . Due to the symmetry among nodes, at each step  $t$  we only need to keep track of the number of nodes in the two groups whose inventory corresponds to the three categories above. Moreover, because the total number of two communities are  $n$ , we can use four variables to describes this random process: fraction of  $\{0\}$ ,  $\{1\}$  nodes in two communities.

As the size of community increase, the above process is closely related to its mean field which can be seen as a autonomous system in  $\mathbb{R}^4$ . We show that this system has a stable fixed point as long as  $0 \leq p < p_0$ . To proof Theorem 4.1 we show two things with the help in autonomous system:

1. Global behaviour: the random process  $X(t)$  will initially “converge” to a point corresponding to the stabile fixed point of the autonomous system.
2. Local behaviour: the random process  $X(t)$  takes exponential time to leave the regions corresponding to the regions around stable fixed point of the autonomous system.

### Tree Structure

In this section, we systematically study the Naming Game on trees and examine how the naming game converges when applying  $\text{Hom}(\cdot, p)$  to the tree structure. We show that con-

vergence is substantially sped up for random edges added with small probabilities.

In Figure 4 (row 1 (b)) we evaluate the time to reach consensus based on the probability  $p$  of choosing random edges. It is clear that for trees with smaller depth ( $d$ ) and more branches, the time to reach consensus is larger. Compared with Figure 4 (row 1 (a)) we see that by adding random edges, the tree graph is much less likely to encounter a time-out than the densely connected community graph. Figure 4 (row 2 (b)) show the fraction of agents failing to reach consensus as a function of  $p$ . Additionally, though the additional pairs can break up the sparse community structure and help to accelerate the converging process, redundant communications may be introduced at the same time. Therefore in Figure 4 (c) for various  $p$ , we present the total number of time the dynamics choose a homogenous edge before consensus is reached, normalized by time it takes homogeneous graph (clique) to reach consensus. We can see that there is actually an tipping point where the homogenous edges are used the least, which implies the edges of the original tree actually help towards consensus. Above this point, the homogeneous edges provide unnecessary communication redundancy.

## 5 Stubborn Nodes

In this section we introduce another aspect of asymmetry and heterogeneity. We introduce special agents called “stubborn” nodes, which never change their own opinions and aim to influence the whole network. The topic is also related to the robustness of the naming game in the real world setting, in which a small number of nodes can be malicious and

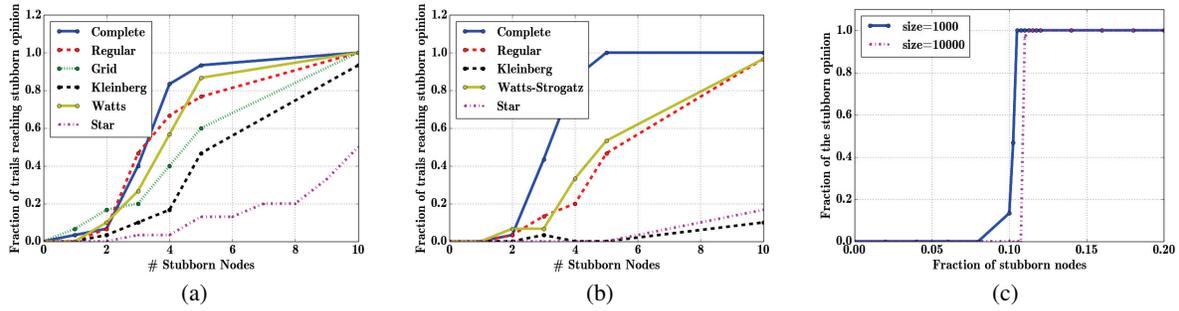


Figure 5: Evaluation for early stage coordinated stubborn nodes within different graphs. (a) fraction of trials converging to the stubborn nodes’ opinion, as a function of the number of stubborn nodes with size 1000; (b) fraction of trials converging to the stubborn nodes’ opinion, as a function of the number of stubborn nodes with size 10000; (c) fraction of nodes converging to the stubborn nodes’ opinion in the late stage situation for complete graph of size 1000 and 10000.

not follow the protocol. The primary question we want to ask is: how and when can such nodes affect the opinion/name to which the dynamics converge? There are two important factors to consider here – the network topology and the time when the stubborn nodes are activated. Here we consider two situations: 1) the stubborn nodes join at the beginning of the game; 2) the nodes become “stubborn” after the graph has converged to one global opinion. Figure 5 (a) (b) show the fraction of trials converging to the stubborn nodes’ preference based on the number of stubborn nodes in situation 1) for graphs of size 1000 and 10000, respectively. From Figure 5 (a), it is clear that in some graphs (e.g. the clique) even a small constant (e.g. 5) number of stubborn nodes can guarantee convergence to a particular name. Note that as the number of nodes increase, the curve barely changes, and if anything, becomes a sharper threshold. However, in others networks (e.g. star, grid, Kleinberg model), the number of required nodes seems to grow with the size of the graph.

This shows that in certain networks these protocols are not robust to rogue/stubborn agents. By comparing Figure 5 (a) and (b)), we see that the complete graph is not affected much by its size in terms of the influence efficiency of the “stubborn” nodes. However, in the Kleinberg and star graphs the number of stubborn nodes needed greatly depends on the size of the network. Note that here we choose the same number of stubborn nodes because complete, regular and Watts-Strogatz graphs actually perform similarly with size 1000 on these number of stubborn nodes.

Additionally, we show that in the complete graph, manipulating the name after convergence is much harder than before: there exists a value  $p_0 \in [0, 1]$ , such that if an adversary controls more than  $p_0$  fraction of the node, consensus can be easily manipulated and otherwise it will take exponential time to manipulate the consensus. In Figure 5 (c) we verify this empirically by showing the fraction of trials converging to the stubborn nodes’ preference (before timing out) based on the fraction of “stubborn” nodes within the network. It shows that at least 10% such stubborn nodes are needed to manipulate the opinion of the original graph empirically.

We provide theoretical analysis on the lower bound for the number of “stubborn” nodes required to manipulate the

global consensus to align with the “stubborn” nodes’ in a complete graph. In completed graph if the naming game converges to opinion 1, we want to answer the following question: what fraction of stubborn nodes with opinion 0 are required in order to convert the graph’s consensus to 0 in polynomial time?

**Theorem 5.1.** *Given the naming game with  $p$  fraction of stubborn nodes defined above, there exists a constant  $p_0 \approx 0.108$  such that for all  $0 \leq p < p_0$  the expected consensus time is  $\exp(\Omega(n))$ . Additionally, if  $p_0 < p \leq 1$  for all  $\epsilon > 0$  the fraction of original opinion is smaller than  $\epsilon$  after  $O(n)$  steps.*

The proof appears in the full paper.

## 6 Conclusions

Our work sheds light on how asymmetric and heterogeneous agents, in both network topology and node types, affect the Naming Game with respect to convergence rate and the converged state. We show that in network structure complete asymmetry as in the star network is beneficial while asymmetry represented in community structures slow down global convergence. We rigorously studied how community structure can prohibit global convergence, and we also theoretically prove how much additional communication is required to facilitate convergence for graphs with community structures. Besides, we also analyze the robustness of different topologies (protocols) and discovered that only about 1% stubborn nodes can make a big difference in the behavior if they are introduced early in the game, while the introduction of stubborn nodes after convergence is achieved does not change much. Additionally, we theoretically prove how many stubborn nodes are needed to convert the global consensus for a complete graph, when they emerge after the convergence is reached. We believe these insights improve our foundational understanding of social dynamics in multi-agent systems and will spur on further insightful studies on this topic.

## References

- Backstrom, L.; Huttenlocher, D.; Kleinberg, J.; and Lan, X. 2006. Group formation in large social networks: membership, growth, and evolution. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 44–54. ACM.
- Baronchelli, A.; Felici, M.; Loreto, V.; Caglioti, E.; and Steels, L. 2006a. Sharp transition towards shared vocabularies in multi-agent systems. *Journal of Statistical Mechanics: Theory and Experiment* 2006(06):P06014.
- Baronchelli, A.; Loreto, V.; Dall'Asta, L.; and Barrat, A. 2006b. Bootstrapping communication in language games: Strategy, topology and all that. In *Proceedings of the 6th International Conference on the Evolution of Language*, volume 2006, 11–18. World Scientific Press.
- Baronchelli, A.; Dall'Asta, L.; Barrat, A.; and Loreto, V. 2007. The role of topology on the dynamics of the naming game. *The European Physical Journal Special Topics* 143(1):233–235.
- Baronchelli, A.; Loreto, V.; and Steels, L. 2008. In-depth analysis of the naming game dynamics: the homogeneous mixing case. *International Journal of Modern Physics C* 19(05):785–812.
- Bollobás, B. 1998. Random graphs. In *Modern Graph Theory*. Springer. 215–252.
- Brighton, H., and Kirby, S. 2001. The survival of the smallest: Stability conditions for the cultural evolution of compositional language. In *European Conference on Artificial Life*, 592–601. Springer.
- Centola, D., and Baronchelli, A. 2015. The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences* 112(7):1989–1994.
- Dall'Asta, L.; Baronchelli, A.; Barrat, A.; and Loreto, V. 2006. Agreement dynamics on small-world networks. *EPL (Europhysics Letters)* 73(6):969.
- Franks, H.; Griffiths, N.; and Jhumka, A. 2013. Manipulating convention emergence using influencer agents. *Autonomous Agents and Multi-Agent Systems* 26(3):315–353.
- Gong, T.; Ke, J.; Minett, J. W.; and Wang, W. S. 2004. A computational framework to simulate the coevolution of language and social structure. In *Artificial Life IX: Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems*, 158–64.
- Hurford, J. R. 1989. Biological evolution of the saussurean sign as a component of the language acquisition device. *Lingua* 77(2):187–222.
- Kleinberg, J. 2000. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, 163–170. ACM.
- Lu, Q.; Korniss, G.; and Szymanski, B. K. 2009. The naming game in social networks: community formation and consensus engineering. *Journal of Economic Interaction and Coordination* 4(2):221–235.
- Nowak, M. A., and Krakauer, D. C. 1999. The evolution of language. *Proceedings of the National Academy of Sciences* 96(14):8028–8033.
- Nowak, M. A.; Plotkin, J. B.; and Jansen, V. A. 2000. The evolution of syntactic communication. *Nature* 404(6777):495–498.
- Nowak, M. A.; Plotkin, J. B.; and Krakauer, D. C. 1999. The evolutionary language game. *Journal of Theoretical Biology* 200(2):147–162.
- Pickering, W., and Lim, C. 2016. Solution of the multistate voter model and application to strong neutrals in the naming game. *Physical Review E* 93(3):032318.
- Steels, L., and McIntyre, A. 1998. Spatially distributed naming games. *Advances in complex systems* 1(04):301–323.
- Steels, L. 1995. A self-organizing spatial vocabulary. *Artificial life* 2(3):319–332.
- Steels, L. 2005. The emergence and evolution of linguistic structure: from lexical to grammatical communication systems. *Connection science* 17(3-4):213–230.
- Watts, D. J., and Strogatz, S. H. 1998. Collective dynamics of small-world networks. *Nature* 393(6684):440–442.