

Coordination and Common Knowledge on Communication Networks

Gizem Korkmaz
Virginia Tech
gkorkmaz@vt.edu

Monica Capra
Claremont Graduate University
monica.capra@cgu.edu

Adriana Kraig
Claremont Graduate University
kraigadriana@gmail.com

Kiran Lakkaraju
Sandia National Laboratories
klakkar@sandia.gov

Chris J. Kuhlman
Virginia Tech
ckuhlman@vt.edu

Fernando Vega-Redondo
Bocconi University
fernando.vega@unibocconi.it

ABSTRACT

Protest is a collective action problem and can be modeled as a coordination game in which two or more people each take an action with the potential to achieve shared mutual benefits, only if their actions coincide. In the context of protest participation, successful coordination requires that people know each others' willingness to participate, and that this information is common knowledge. Social networks can facilitate the creation of common knowledge through the flow of messages. Although there is a rich experimental literature that documents behavior in coordination games with and without communication, little is known about how people coordinate behaviors within a social network and how different types of communication structures affect behavior.

In this paper, we develop a theoretically based on-line experiment with Amazon Mechanical Turk participants to characterize the emergence of common knowledge and coordination through interactions within a network. Our experiment is designed to identify the effects of both social network topology and communication and to falsify the game-theoretic predictions. Our data reveal that choices are affected by the network structure and they move towards the theoretical predictions with communication. We use our behavioral findings to simulate dynamics in more complex networks through agent-based modeling. Thus, we combine human behaviors identified in experiments with realistic social network structures to reveal patterns not previously observed.

KEYWORDS

coordination; common knowledge; social networks; Amazon Mechanical Turk; online experiments; agent-based modeling

ACM Reference Format:

Gizem Korkmaz, Monica Capra, Adriana Kraig, Kiran Lakkaraju, Chris J. Kuhlman, and Fernando Vega-Redondo. 2018. Coordination and Common Knowledge on Communication Networks. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018*, IFAAMAS, 9 pages.

1 INTRODUCTION

A single protester risks prosecution in authoritarian regimes and has little chance of success. This risk can be reduced if a million people successfully coordinate their actions. Thus the protest is a

collective action problem where an individual wants to protest only if joined by an "enough" number of other protesters. Game theorists refer to this type of collective action problem as a **coordination game** in which two or more people each take an action with the potential to achieve shared mutual benefits, only if their actions coincide [30, 37]. Critically, in the context of protest participation, successful coordination requires that people know each others' willingness to participate, and that this information is **common knowledge**. Common knowledge is defined as an infinite string of embedded levels of knowledge (i.e., I know they will participate, they know I will participate, I know that they know, they know that I know, I know that they know that I know, and so on).

Clearly, **communication technologies** play an important role in the creation of common knowledge. Social networks, such as Facebook, Twitter, and Youtube can facilitate information sharing and the generation of common knowledge within groups of users. Coordination within social network sites was a distinctive feature of the infamous uprisings against authoritarian regimes such as the Arab spring and Gezi protests in Istanbul. In contrast to traditional media that broadcast widely to a crowd, social networks can facilitate actionable common knowledge through local interactions. These interactions generate patterns of contact that allow for the flow of messages among communicators through time and space [33]. Here, message should be understood in its broadest sense to refer to information and knowledge that can flow from one point in a network to another and can be co-created by network members. The ways in which individuals can share messages (i.e., the communication mechanisms such as electronic mail, social media, nonverbal: eye-contact, blushing, etc.) and the structural features of the local interactions (i.e., network topology) affect the generation and distribution of messages.

Previous studies by Chwe [13, 14] and Korkmaz et al. [28] combined social structure and individual incentives to provide a rigorous game-theoretic formalization of common knowledge on social networks (and the characterizing network structures) within the context of collective action (e.g., protests). The former emphasized simple node-to-node or bilateral communication, whereas the latter studied the effects of "richer" on-line communication mechanism, such as Facebook. However, there is little empirical evidence that supports these stylized common knowledge models, nor do these models incorporate behavioral factors, and psychological processes.

In this paper, we develop a theoretically based experimental framework to characterize the emergence of common knowledge through interactions among individuals in networks. We test novel

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

hypotheses derived from the mathematical model [14] at the individual level by conducting human subject experiments in an on-line setting. The objectives are (i) to characterize how communication can facilitate an individual’s awareness of common knowledge through local interactions, and (ii) to understand the effect of network structure. More specifically, we designed an experiment with Amazon Mechanical Turk (AMT) workers in which human subjects and bots participated in several rounds of one-shot coordination games within three different kind of network structures: Clique, Circle and Star. Treatments varied with respect to whether players could send messages, and whether they could observe the complete graph. Choices consisted of either participating in a group event or not participating. Participation was profitable only when a threshold number of others also participated in the event. In our experiments, human subjects were matched with bots to make decisions within a network. Bots were programmed to follow the theoretical prediction. We matched human subjects with bots, because this allowed us to isolate the network and message effects from other group dynamics that may arise from humans playing humans and having different decision rules.

Results of our experiment show that communication helped players coordinate their actions. The mechanism of coordination being the generation of common knowledge. We found that participation rates tended to be higher than predicted, in general, but mostly when subjects knew more about the network. Interestingly, communication did not always lead to more participation. Instead, communication had the effect of moving participation rates closer to the theoretical predictions, which sometimes prescribed no participation. Without communication and when the subject’s threshold for participating was high, Clique networks performed worse than other networks. In contrast, Star networks performed worst when participation threshold was low. Through agent-based modeling, we combined human behaviors identified in the experiments with realistic social network structures. This exercise revealed interesting effects from the fact that agents can participate in more than one common knowledge set.

The rest of the paper is organized as follows. The next section summarizes related work. In Section 3, we describe the theoretical model presented in Chwe [13, 14] that our experiments are built on. The experimental design, recruitment, and procedures are described in Section 4. We introduce the agent-based modeling framework and the real networks used for the simulations in Section 5. The experimental and simulation results are presented in Sections 6.1 and 6.2, respectively. Section 7 discusses our findings.

2 RELATED WORK

Common knowledge has come up in many different scholarly contexts; David Lewis [30], influenced by Thomas Schelling [37], first made it explicitly, and Robert Aumann [3, 4] developed the mathematical representation. The problem of coordination and common knowledge has been examined by many disciplines, including political science [34], philosophy [27, 30, 36], economics [16, 23], linguistics [17, 18, 39], sociology [42], legal theory [31], and computer science [1, 24].

In a series of six experiments, [20] found experimentally that 2-level mutual knowledge (i knows that j knows that they both

know X) produces coordination more frequently than common knowledge under particular conditions. Any level of knowledge that fall short of infinity is called “shared knowledge.”

Recent psychological research shows that people represent common knowledge as qualitatively distinct from shared knowledge and this distinction affects their strategic decisions [41]. Thomas et al. [41] placed participants in hypothetical scenarios with the payoffs of a Stag Hunt: they had to decide whether to work alone for a certain but lower payoff, or to try to work together with a partner for a higher payoff only if both made the same choice. They found that subjects were most likely to choose to coordinate, when they had common knowledge, in line with game-theoretic predictions. These experiments were conducted with subjects in pairs.

More recently, there has been rapid growth of experimental platforms and research on social networks, which are invaluable tools to validate existing theoretical findings [11]. Experiments reveal how individuals actually use network information, which may be full or partial, and are generating behavioral data that relates network structure to choices. These data, in turn, can serve as an input for novel theoretical developments. Although there are a number of recent papers that employ laboratory experiments to study decisions within networks (e.g., [6–10]), the experimental evidence concerning the role of network structure on collective action and the degree to which network structure and communication interact is not conclusive [40]. Much has been learned about behavior in coordination games on networks (e.g., [2, 12, 21, 22]), but none of these studies mainly focus on the formation of *common knowledge* embedded in a network.

3 THE MODEL

In this paper, we focus on the game-theoretic model presented in Chwe [13, 14]. Suppose there is a finite set of people $N = \{1, 2, \dots, n\}$ and each person $i \in N$ chooses an action $a_i \in \{0, 1\}$, where 0 is the safe action, and 1 is the risky action (e.g., participation in a protest). Each person i has an idiosyncratic private threshold $T_i \in \{1, 2, \dots, n\}$, which is the minimum number of people that must choose action 1 for i to benefit from choosing action 1. Individuals in N are connected by edges in the social network G , which denote pairwise interactions. Let d_i denote the number of connections, i.e., the degree, of person i .

Chwe [13, 14] models social structure as a communication network through which every person i tells her neighbors her willingness to participate, represented by her threshold T_i . The communication network helps coordination by creating common knowledge at each discrete time. Given person i ’s threshold T_i and everyone’s actions $a_s = (a_{1s}, a_{2s}, \dots, a_{ns})$, his utility at time $s \in \{0, 1, \dots, S\}$ can be formulated as

$$U_{is} = \begin{cases} 0 & \text{if } a_{is} = 0 \\ 1 & \text{if } a_{is} = 1 \wedge \#\{j \in N : a_{js} = 1\} \geq T_i \\ -z & \text{if } a_{is} = 1 \wedge \#\{j \in N : a_{js} = 1\} < T_i \end{cases} \quad (1)$$

where $-z < 0$ is the penalty he gets if he participates and not enough people join him. Thus, a person will participate as long as he is sure that there is a sufficient number of people (in the population) choosing the risky action. A person always gets utility 0 by staying at home regardless of what others do since we do not consider free-riding problems. When he participates, he gets utility 1 if the

total number of people participating is at least T_i . Each individual must take into account what she expects the other agents to do. As described before, in game-theoretic contexts, coordination requires that people know each others' willingness to participate, and that this information is *common knowledge* among a sufficient number of people. Common knowledge among a set of people implies that: *they know each others' thresholds and they know that they know their thresholds*. Therefore, they can count on each other.

The Chwe model [15] assumes that the network itself is common knowledge, so that agents know all communication channels that exist between all pairs of members of the population. One of the key features of large social networks prevailing in the real world is that agents only have local information about the network. Therefore, in our experimental setup we also test the network-knowledge and consider the cases when the network structure – who is connected to whom – is commonly known (referred to as global network-knowledge), and only locally known. In our set-up, common knowledge can also arise from the interaction of communication and network topology.

4 THE EXPERIMENT

4.1 Subjects

We recruited a total of 165 subjects from Amazon Mechanical Turk. The experiment consisted of a questionnaire to extract demographic information and personality traits, a set of hypothetical choice tasks to elicit preferences over lotteries and over future payments, and a series of decision trials. In each of these decision trials, each subject made a one-shot decision on whether to participate in a group event. Each group consisted of five decision-makers: four bots and one human subject. Subjects were compensated for their participation and for the earnings they got in one randomly chosen trial. On average, subjects made 78.3 Experiment Currency Units (ECU) where a single ECU was worth \$0.03. The experiment took, on average, 36 minutes to complete. Subject characteristics were broad: 60% of the subjects were between 18-35; close to 60% were male, more than 50% had a Bachelor's degree, and close to 35% had a household income in the range of \$25K-\$49K.

4.2 Design

The main objective of the study is to experimentally identify the effect of network topology (clique, circle, or star) and messages on coordination in threshold participation games under global information (where subjects observe the complete graph) and local network information (where subjects observe their own links only).

Our main outcome variable is participation rates. In each trial, subjects made decisions of whether to participate or not in a group project. The choice of no participation guaranteed a payoff of 50 units, whereas the payoffs from participation were 150 units or 0 units depending on the subject's threshold. The threshold represented the minimum number of other players needed for profitable participation. Low thresholds (L) required that at least one other person participated to earn 150 units. High thresholds (H) required that at least 3 other people participated in order to earn 150 units.

At the beginning of each trial, each subject was randomly assigned an avatar, a threshold (H or L) and a group that consisted of 4

bots with varying threshold levels. Unknown to the subject, the bots played the strategies that aligned with the theoretical predictions.

Our experiment consisted of a between and within subject mixed design. There were a total of 4 treatments that varied with respect to communication: *no communication or bi-lateral*, and whether the subject had local or complete knowledge of the network: *local or global*. In each of these treatments, subjects made a series of one-shot decisions under differing network topology: *star, circle or clique*. Table 1 summarizes the treatments and the number of observations in each condition. The degrees (d) of the nodes in each network are given in brackets.

Condition	Subjects	Starp (d=1)	Circle (d=2)	Star (d=4)	Clique (d=4)
Local – No Comm.	44	176	352	0	440
Local – With Comm.	39	156	312	0	390
Global – No Comm.	39	156	312	390	390
Global – With Comm.	43	172	344	430	430

Table 1: Summary of treatments and the number of trials and subjects per condition.

Communication. In the communication treatments, subjects had the option to send messages to and receive messages from those with whom they had a direct link. In other words, we implemented node-to-node or bi-lateral communication. Messages were chosen from a set of predetermined options: “I will Participate” or “I will NOT Participate”. In contrast, in the no communication condition, participants were unable to send or receive any messages.

Network Topology. As mentioned above, there were a total of 4 treatment conditions (2 communication conditions x 2 network knowledge conditions). In each of these treatments, the same subjects made a series of one-shot participation decisions under differing network topology and threshold distribution. We studied three graphs: circle, star, and clique, illustrated in Figure 1. The Circle block type is a condition where each of the 5 players is connected to 2 others by communication links (i.e., a pentagon). The Star block type is a condition where 1 of the 5 players is connected to each of the other 4 (i.e., a cross). Subjects were exposed to trials in the star condition as a periphery point (starp) and again as the center point (star). Lastly, the Clique block type is a condition where each of the 5 participants are connected to everyone else.

Knowledge of Network. In the global network knowledge treatments, subjects were given the complete network graph including the names and thresholds of all other players (represented by avatars) and all players' neighbors (see Figure 1). Thus, the subject knew his own position in the network, and who could communicate with whom. In the local network knowledge treatments, subjects did not observe the complete network and were only given threshold information for those with whom they were directly connected (see Figure 2); however, the avatars of all other players were revealed on a list.

4.3 Procedures

We recruited and rewarded subjects through Amazon Mechanical Turk (AMT). AMT is an on-line marketplace for tasks. Workers can login and perform tasks posted by requestors. Requestors, in turn, provide payment to the workers through the AMT interface.




Figure 1: Network topologies and arrangements for Global network-knowledge condition. The subject is shown in blue.




Figure 2: Network topologies and arrangements for Local network-knowledge condition. The subject is shown in blue.

AMT has become a popular site to deploy tasks, called "Human Intelligence Tasks (HITS)", that are easy for humans but difficult for machines. AMT has several benefits. It is relatively inexpensive, it provides a diverse subject pool ([26, 29, 35]), and it is fast. Replication of results between the lab and AMT ([5, 19]) suggests that AMT can be an useful proxy for laboratory experiments.

The experiment was run in two phases. In the first phase, a task was posted on Amazon Mechanical Turk asking for workers to complete the survey. We requested, and received, 575 responses to the task. In the second phase, we posted a task to participate in the experiment and restricted it to only subjects who had completed our survey. We received 167 responses to the HIT for the experiment.

After consenting to participate in the experiment, subjects were shown the instructions for the decision task. Upon finishing with the instructions, subjects were given a quiz to ensure that they understood the information presented throughout and they knew how to navigate through the game platform. No one could move on to the decision task before answering all questions correctly.

The timing of each trial started with all participants being assigned a new avatar and threshold that appeared on their screen or "web-page". They were also able to see diagrams of their network according to their "Network Knowledge" (i.e., local or global), and a list of their direct connections with their connections' thresholds (either H or L) (see Figure 3). In the "communication" sessions, players had the option to send a pre-determined message "I will participate" or "I will NOT participate" to his/her direct connections. After all messages were sent, they were able to read the messages sent to them. When the players were finished reading messages, they had to make the final decision to either "Participate (P)" or "Not Participate (NP)" for that trial. Once the decision to participate was made, the next trial began without any feedback, and the process repeated for a total of 22 or 32 trials in the local or global network knowledge treatments, respectively. After all trials in a session were completed, participants were paid privately according to the outcomes of a randomly chosen trial.

4.4 Bot Behavior

As mentioned above, in each trial, each human subject was matched with 4 bots to form a group. This was unknown to the subjects,




Figure 3: Subject messaging screen. In this trial the subject is in the "Bilateral" communication condition.

who made repeated decisions matched with bots only. Bots were programmed to follow the theoretical decisions and to send messages that truthfully reflected their intentions to either participate or not participate. Table 2 describes the bots' decision rules. The decision rule r_1 represents: "If there is at least one friend that has a Low threshold, I Participate (P) else I do Not Participate (NP)".

Network Type	Knowledge	Threshold H	Threshold L
Clique	Local	NP	r_1
	Global	P	P
Star	Local	NP	r_1
	Global	NP	r_1
Starp	Local	NP	r_1
	Global	NP	r_1
Circle	Local	NP	r_1
	Global	NP	r_1

Table 2: Bot decision making. Columns "T=H" and "T=L" refer to the bot threshold. The entries indicate whether the bot chose Participate (P) or Not Participate (NP). The decision rule of bots, denoted as r_1 , is described in Section 4.4.

5 AGENT-BASED MODELING

5.1 Two Types of ABMs and Simulations

To understand the implications of the human subjects experiments on larger, real-world populations, we performed simulations of the initiation and spread of CK using agent-based modeling (ABM) and social networks. We constructed simulations for two models. First, we constructed the agent-based simulation (ABS) for Chwe model of Section 3, called *theoretical simulation* (TS). Our second model, *experimentally augmented TS* (EATS), augments the Chwe model with data from experiments on cliques where players know the network structure (global knowledge). In these simulations, we used the experimental data for the case of no communication among agents. The algorithm is given immediately below.

5.2 Algorithm

We first focus on the algorithm for computing nodes that transition from state 0 to state 1 with a CK set M at time step s . We use the term agent “state” as a synonym for agent “action.” Also, we use states 0 and 1 to represent actions participate (P) and not participate (NP), respectively. Algorithm 1 is used in both the TS and EATS simulations for this purpose. Thereafter, we put this algorithm in the context of the overall simulations.

In Algorithm 1, we address the case where all nodes v_j in a CK set M are in state 0 ($a_j = 0$). In the **for** loop, starting with the maximum threshold $T_{max,0}(M)$ of all nodes of M , we determine whether this maximum is less than or equal to the number of nodes in state 0. If so, all nodes in state 0 (i.e., in set S_{k0T}) transition to state 1, according to the Chwe model. If not, then the nodes with the maximum threshold are removed, and the **for** loop repeats. If the nodes in S_{k0T} do transition for some loop index, and an EATS is being executed, then each node’s actual transition is determined probabilistically, using the distribution in Table 3, according to whether each node in M , in turn, has a low or high threshold. In essence, Table 3 contains probabilities of agent transition, conditioned on the TS model prediction of transition, for configurations of M . Thus, in this last step the theoretical model is augmented by a probabilistic model constructed with the experimental data of this study. By setting a flag, the simulations can be run in each of two modes: TS and EATS. Finally, we note that by conditioning the experimentally augmented transitions on the pure-theory prediction in Table 3, we extend the range of conditions (e.g., network structures, thresholds) over which these experimental data may be used in modeling. Algorithm 1 is executed within the context of a simulation that includes a user-specified number of diffusion instances or runs, and a user-specified number of time steps for each run.

Node Threshold	0.0	0.25	0.50	0.75	1.0
L	0.7949	0.9231	0.9487	0.8974	0.9487
H	0.6410	0.6154	0.6923	0.7436	0.7949

Table 3: Experimentally augmented transition probabilities, $\Pr(\text{state transition} \mid w_1, w_2, w_3, w_4, w_5)$, for a single node in a CK set M for EATS, based on the fraction of other nodes in M that have low threshold L (from 0 to 1.0 in 0.25 increment), where the conditions are w_1 : the particular CK substructure [here, clique]; w_2 : whether the nodes know the local or global clique substructure [here, global]; w_3 : communication type [here, no communication]; w_4 : theory prediction of state transition [here, for the case where the theory predicts that the node will transition]; and w_5 : whether the node evaluated has a low L or high H threshold. This probability is used in Algorithm 1. Linear interpolation is used for other values of the fraction of other nodes with L. These data were produced from the results in Section 6.1.

5.3 Networks

The two social networks used in ABS to evaluate CK spreading on realistic populations are given in Table 4. NRV is a human social contact network for a high school in the New River Valley (NRV), Virginia. AH is an mutual friendship network from Add-Health [25].

Algorithm 1 CK-based state transition algorithm for agents in one common knowledge set M .

```

1: Inputs: time  $t$ , CK set  $M$ ;  $\forall v_j \in M$ , tuples  $(v_j, T_j, a_j)$ ; simulation type (TS or EATS).
2: Outputs: Set  $S_{k0T} \subseteq M$  for TS ABS, or  $Q \subseteq S_{k0T}$  for EATS ABS, the set of agents changing to state 1. ( $S_{k0T}$  or  $Q$  may possibly be the empty set.)
3: Steps:
4:  $T_{max,0}(M)$  is the maximum threshold of all agents in  $M$  in state 0.
5:  $T_{min,0}(M)$  is the minimum threshold of all agents in  $M$  in state 0.
6: for ( $k = T_{max,0}(M)$ ;  $k \geq T_{min,0}$ ;  $--k$ ) do
7:    $S_{k0T} = \{v_j \in M : a_j = 0 \text{ and } T_j \leq k\}$ .
8:   if ( $|S_{k0T}| \geq k$ ) then
9:     // All nodes in  $S_{k0T}$  transition according to theory.
10:    if (simulation type == TS) then
11:      Return ( $S_{k0T}$ )
12:    else
13:      // Running EATS.
14:      From nodes  $v_j \in S_{k0T}$ , compute the set  $Q \subseteq S_{k0T}$  of experimentally augmented state transitions according to Table 3.
15:      return( $Q$ ).
16:    end if
17:  end if
18: end for
19: // No agent transitions; return the empty set.
20: return( $\emptyset$ ).
```

Network	Type	n	m	d_{ave}	d_{max}	n_c	C^*
NRV	High School	769	4551	11.8	20	1495	9
AH	Friendship	2448	5277	4.31	10	1140	9

Table 4: Networks used in simulations, which are 2 and 3 orders of magnitude larger than the experimental networks (in terms of n). n and m are numbers of nodes and edges; d_{ave} and d_{max} are average and maximum degrees; n_c is number of cliques in the graph and C^* is the maximum clique size.

6 RESULTS

6.1 Experimental Results

We collected data from 165 AMT participants: 82 of whom were randomly assigned to the bi-lateral “communication” condition and 83 of whom were in the “no communication” condition.

In the Clique with communication, participation rates were not statistically different from the theoretical predictions. Figure 4 illustrates the participation decision of subjects with low threshold, L, in the clique network. In this network, each player has 4 neighbors. Since the subject’s threshold is 1, s/he needs at least one other player to benefit from choosing P. The red bars correspond to the case in which none of the neighbors had low threshold, and the blue ones represent the case which at least one of the neighbors had low threshold. Based on the theoretical models, subjects should choose NP in the former, and P in the latter.

Yet, we observed that there was a bias towards participation that is most likely due to an “irrational” expectation that others would




Figure 4: Percentage of subjects with low threshold, L, that chose “Participate” in the clique network. Error bars represent standard error. The theoretical prediction is 0% participation under the local condition when all neighbors have high threshold (red bars), and 100% under the global condition for both threshold configurations.

participate. We found that 86% of the subjects choose P (under local - no communication condition), although they did not have any neighbors with low thresholds. Introducing communication reduced this percentage to 18%. The bias to participate disappeared with communication, because the bots sent truthful messages (as prescribed in the model) that they would not participate. Given that no participation was common knowledge, subjects had no reason to believe that others, all H, would participate.

In the global network-knowledge condition, high participation levels ranging from 79%-95% with no communication and participation rates increased to almost full participation, 98%-100%, with communication. These results follow the theoretical prediction.




Figure 5: Percentage of subjects with high threshold, H, that chose “Participate” in the clique network. The number of neighbors (nbrs) with threshold L is on the x-axis. The theoretical prediction is 0% participation under local, 100% under global conditions.

When subjects’ thresholds were high, communication also moved decisions towards theoretical predictions (Figure 5). A high threshold means that the subject needed at least three others to participate. Under the local-no communication condition, the theoretical prediction was “NP” for all threshold allocations. However, the participation probability ranged from 59%-79% (the red line in Figure 5). Under the global network knowledge condition on a clique, the subject observed everyone’s thresholds and that everyone was connected to everyone else (hence everyone knew everyone else’s

thresholds). Everyone ‘should have known’ that if everyone (regardless of thresholds) participated, everyone would benefit, because the threshold would be met. In comparison to the local condition (in which theory predicts zero participation for H-type), the model predicts full participation in the global network condition. However, the participation probability we obtained in the global network knowledge was in the range 61%-79% (the green line in Figure 5), which was not statistically different from the 59%-79% range.




Figure 6: Percentage of subjects that chose “Participate” in the star network. The theoretical prediction is 0% participation when subject has high threshold (right). When the subject’s threshold is T=L (left), the prediction is 100% participation when there is at least 1 neighbor with T=L (green bars), and 0% when all neighbors of the subject have T=H (red).

In the star network under the global network-knowledge condition (see Figure 1), the subject observed the exact same network as in the clique-local condition (see Figure 2). Not surprisingly, similar behavior was observed in the two networks when the subjects’ thresholds were low. This is shown in Figure 4 (the first four bars) and Figure 6 (left). Here, communication reduced the participation rate significantly. When the subject had a high threshold and there was no communication, behaviors were similar (red and green lines in Figure 5 and the green line in Figure 7, respectively). However, when communication was introduced, the participation rate was significantly lower in the star network compared to the clique. The purple line in Figure 7 is significantly lower than the purple line in Figure 5. This effect is in part due to the fact that messages were generated by truthful bots, and reflected the optimal decision.




Figure 7: Percentage of subjects with threshold H that chose “Participate” in the star network under global network knowledge condition.

In the circle network (Figure 8), when the subject had 2 neighbors, the results are similar to those for the star network as well as the ones for the star network, which we omitted due to space constraints. To summarize, communication in these conditions also

moved the participation rates towards the theoretical prediction. Communication yielded higher participation rates when at least one of the neighbors had low threshold or $T=L$ (green bars), and lower rate when both neighbors had high thresholds or $T=H$ (red).




Figure 8: Percentage of subjects that chose “Participate” in the circle ($d=2$) network. The theoretical prediction is 100% participation when at least 1 of the neighbors has low threshold (green), and 0% when both neighbors have $T=H$.

Figure 9 summarizes the findings for the networks of the study and for all conditions when everyone in the network (the subject and the four bots) had high threshold, $T=H$ (when $T=L$ for all nodes, the participation rate for all networks was high, and communication increased it further; this figure is omitted due to space constraints). When $T=H$, communication decreased the participation significantly for three networks (circle, star and starp) under both local and global conditions (Figure 9). In the clique network, communication increased participation rates under the global network-knowledge condition, and decreased under the local.




Figure 9: Percentage of subjects that chose “Participate” when $T=L$ and $T=H$ for all subjects, respectively.

We compared and contrasted these findings with the theoretical predictions in Figure 10. The comparison suggests that the participation rate is below the theoretical prediction for all networks when $T=L$. Communication moved the rate towards the horizontal line that represents the match between the predictions and the findings. When the threshold was high for all nodes, communication pushed behavior towards the theoretical predictions (although the effects on the networks vary).




Figure 10: Deviations from theoretical predictions when $T=L$ and $T=H$ for all subjects, respectively. The dashed horizontal line represents the match between the theoretical prediction and the experimental finding.

6.2 ABM Results

We performed over 80 simulations on the two networks in Table 4, with conditions summarized in Table 5. Within each simulation, each run used a different threshold assignment of L and H to nodes, with one-half the nodes possessing each threshold in expectation. Note that nodes with threshold $T > (C - 1)$, where C is clique size, will not transition to state 1.

n_n	n_s	n_r	t_{max}	L	H	Sim. Type
2	40	100	30	from 2 to H	$[0.2C^*, 1.1C^*]$	TS, EATS

Table 5: Summary of simulation characteristics. The number of networks $n_n = 2$; number of simulations per network $n_s = 40$; number of runs per simulation $n_r = 100$; number of time steps per run $t_{max} = 30$; low thresholds varied from 2 to H, in increments of 2; high threshold varied from a small fraction of the maximum clique size C^* to slightly greater than C^* .

A notable difference between the experiments and the simulations is that while the experiments and the theory-based simulations (TS) are one-shot events, the EATS occur over multiple rounds. This is because the theoretical model, as used here, is deterministic; thus, all nodes that change state do so in one time unit. However, the EATS are probabilistic, owing to the transition probabilities of Table 3.

Simulation results for the NRV network (left) and AH network (right) are given in Figure 11. The cumulative fraction of agents in the network that transitions to state 1 is shown as a function of time for different L,H threshold combinations in the legends.

The agreement between the theory-only model simulation (TS) results and experimentally-augmented theory simulation (EATS) results is somewhat surprising. There is much greater agreement than anticipated from the probabilities in Table 3 and in the experiments, where high threshold nodes only transition with probability

between 0.641 and 0.795. These results in Figure 11 hold even when H is as low as 3 ($< C$ for all but the cliques of size 3), so that even many H -threshold nodes can transition.




Figure 11: Simulation results showing the cumulative fraction of agents in state 1 as a function of time for the two networks: (a) NRV and (b) AH. The legend in each plot gives L, H and simulation type. The theoretical simulation (TS) results are dashed curves and they lay under the corresponding solid curves (EATS), indicating close agreement between the two types of simulations. TS terminate at time 1 because TS are deterministic. The EATS, because of the stochasticity, continue on for a few more time steps, transitioning relatively few nodes for $s > 1$. The qualitative results in these plots hold over all of the conditions in Table 5.

The observation is explained by the fact that roughly one-half of the nodes of AH are in at least two cliques, and one-quarter of the nodes are in at least three cliques; see Figure 12. Since each clique is evaluated at each time, a node that transitions with probability $p = 0.641$ in Table 3 (given that the theory computes that the node should transition) and is in $q = 2$ cliques, will transition with probability $\Pr(\text{state transition}) = 1 - (1 - p)^q = 0.902$; if $q = 3$, $\Pr=0.954$. That is, nodes in more than one clique have a much greater effective probability of transition to state 1. By comparison, almost all nodes of NRV in Figure 12 are in more than one clique. In fact, many social networks (we have evaluated over 10) have this feature that nodes are in multiple cliques, so this is a phenomenon that is applicable to many networks beyond those examined here. Consequently, these simulations illustrate the power of modeling—in capturing how human behavior from experiments and social network structures interact to produce interesting and somewhat counterintuitive results.




Figure 12: Cumulative density function showing the fraction of nodes that appear in the number of CK sets on the abscissa. In AH, one-half of the nodes are in more than one clique; in NRV, almost all nodes are in more than one clique.

7 DISCUSSION

Our main take-away from this experiment is that communication and network structure interact to produce common knowledge. Without communication and with global network knowledge, participation rates are higher than predicted in the Star and Circle networks and lower than predicted in the Clique. Communication does not always lead to more participation; instead, communication allows players to make decisions that are consistent with the theoretical predictions. Sometimes this means more participation, other times this means less participation.

Strategically and theoretically speaking, the participants in the local knowledge with no-communication conditions acted with an exceptionally high risk-accepting attitude. This attitude was even more pronounced among individuals with a high threshold, which artificially produced risk-avoiding attitude. Indeed, participating under these constraints would lead to a suboptimal outcome when other players use theoretically derived behaviors. Despite of inherent risk of a loss in this context, a majority of subjects chose “Participate”.

We can infer from this that subjects had an ‘irrational’ belief that others would participate. In psychology, this phenomenon is known as unrealistic optimism. People are considered to be unrealistically optimistic if they predict that future outcomes will be more personally beneficial, or less negatively harmful, than that suggested by a relevant, objective standard [38].

Indeed, during a pilot testing, we distributed a survey to our subjects to investigate some of the thought processes behind participant motivation. The majority reported basing their individual participation decisions on three sets of information: 1) “my threshold”, 2) “my friends’ threshold” and 3) “my network”. In the local knowledge with no-communication conditions, it was not longer possible to use these to make informed participation decisions. Thus, participants had to rely on their intuition, which can lead individuals to underestimate the control that others have in their lives and overlook the possible motives of others.

We believe that the illusion of control perceived by the personal threshold being the only information available, was the mechanism that led to the high participation rates in the local knowledge with no-communication conditions [32]. Further evidence of this can be seen in the comparisons to the other conditions that held much closer to the theoretical predictions. In the case of the local knowledge with communication condition, it became obvious that the subjects’ intention was still to participate initially. The fact that the other players all sent signals that they would “not participate”, is what ultimately persuaded the participants to change their decision.

8 ACKNOWLEDGMENTS

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525. This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-17-1-0378. Any opinions, finding, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force. This work is also supported by DTRA CNIMS Contract HDTRA1-17-D-0023.

REFERENCES

- [1] Luca Alberucci and Gerhard Jäger. 2005. About cut elimination for logics of common knowledge. *Annals of Pure and Applied Logic* 133, 1 (2005), 73–99.
- [2] Alberto Antonioni, Maria Paula Cacault, Rafael Lalive, and Marco Tomassini. 2013. Coordination on networks: Does topology matter? *PLoS One* 8, 2 (2013), e55033.
- [3] Robert J. Aumann. 1976. Agreeing to disagree. *The Annals of Statistics* (1976), 1236–1239.
- [4] Robert J Aumann et al. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 1 (1974), 67–96.
- [5] Adam J Berinsky, Gregory A Huber, and Gabriel S Lenz. 2012. Evaluating Online Labor Markets for Experimental Research: Amazon.com’s Mechanical Turk. *Political Analysis* 20, 3 (July 2012), 351–368.
- [6] Siegfried K Berninghaus, Karl-Martin Ehrhart, and Claudia Keser. 2002. Conventions and local interaction structures: experimental evidence. *Games and Economic Behavior* 39, 2 (2002), 177–205.
- [7] Alessandra Cassar. 2007. Coordination and cooperation in local, random and small world networks: Experimental evidence. *Games and Economic Behavior* 58, 2 (2007), 209–230.
- [8] Alessandra Cassar, Daniel Friedman, and Patricia Higinio Schneider. 2010. A laboratory investigation of networked markets. *The Economic Journal* 120, 547 (2010), 919–943.
- [9] Gary Charness, Margarida Corominas-Bosch, and Guillaume R Frechette. 2007. Bargaining and network structure: An experiment. *Journal of Economic Theory* 136, 1 (2007), 28–65.
- [10] Gary Charness, Francesco Feri, Miguel A Meléndez-Jiménez, and Matthias Sutter. 2014. Experimental games on networks: Underpinnings of behavior and equilibrium selection. *Econometrica* 82, 5 (2014), 1615–1670.
- [11] Syngjoo Choi, Edoardo Gallo, and Shachar Kariv. 2016. Networks in the Laboratory. *Oxford Handbook of the Economics of Networks* (May 2016).
- [12] Syngjoo Choi and Jihong Lee. 2014. Communication, coordination, and networks. *Journal of the European Economic Association* 12, 1 (2014), 223–247.
- [13] Michael Suk-Young Chwe. 1999. Structure and strategy in collective action. *Amer. J. Sociology* 105, 1 (1999), 128–156.
- [14] Michael Suk-Young Chwe. 2000. Communication and coordination in social networks. *The Review of Economic Studies* 67, 1 (2000), 1–16.
- [15] M. S.-Y. Chwe. 2000. Communication and Coordination in Social Networks. *Rev. of Econ. Stud.* 67 (2000), 1–16.
- [16] Michael Suk-Young Chwe. 2001. *Rational ritual: Culture, coordination, and common knowledge*. Princeton University Press.
- [17] Herbert H Clark. 1996. Using language. *Cambridge University Press: Cambridge* 952 (1996), 274–296.
- [18] Herbert H Clark and Catherine R Marshall. 1981. Definite reference and mutual knowledge. (1981).
- [19] Matthew J C Crump, John V McDonnell, and Todd M Gureckis. 2013. Evaluating Amazon’s Mechanical Turk as a Tool for Experimental Behavioral Research. *PLoS ONE* 8, 3 (March 2013), e57410.
- [20] Giovanna Devetag, Hykel Hosni, and Giacomo Sillari. 2013. You better play 7: mutual versus common knowledge of advice in a weak-link experiment. *Synthese* 190, 8 (2013), 1351–1381.
- [21] Daniel Enemark, Mathew D McCubbins, and Nicholas Weller. 2014. Knowledge and networks: An experimental test of how network knowledge affects coordination. *Social Networks* 36 (2014), 122–133.
- [22] Vincenz Frey, Rense Corten, and Vincent Buskens. 2012. Equilibrium selection in network coordination games: An experimental study. *Review of Network Economics* 11, 3 (2012).
- [23] John Geanakoplos. 1994. Common knowledge. *Handbook of Game Theory with Economic Applications* 2 (1994), 1437–1496.
- [24] Joseph Y Halpern and Yoram Moses. 1990. Knowledge and common knowledge in a distributed environment. *Journal of the ACM (JACM)* 37, 3 (1990), 549–587.
- [25] Kathleen Mullan Harris. 2008. The National Longitudinal Study of Adolescent Health (Add Health), Waves III. (2008). Carolina Population Center, UNC.
- [26] Connor Huff and Dustin Tingley. 2015. “Who are these people?” Evaluating the demographic characteristics and political preferences of MTurk survey respondents. *Research & Politics* 2, 3 (July 2015), 2053168015604648.
- [27] David Hume. 2003. *A treatise of human nature*. Courier Corporation.
- [28] Gizem Korkmaz, Chris J Kuhlman, Achla Marathe, Madhav V Marathe, and Fernando Vega-Redondo. 2014. Collective action through common knowledge using a Facebook model. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 253–260.
- [29] Kiran Lakkaraju. 2015. A Study of Daily Sample Composition on Amazon Mechanical Turk. In *Social Computing, Behavioral-Cultural Modeling, and Prediction*, Nitin Agarwal, Kevin Xu, and Nathaniel Osgood (Eds.). Springer International Publishing, 333–338.
- [30] David K Lewis. 1969. *Convention. A Philosophical Study*. Cambridge (MA). (1969).
- [31] Richard H McAdams and Janice Nadler. 2005. Testing the Focal Point Theory of Legal Compliance: The Effect of Third-Party Expression in an Experimental Hawk/Dove Game. *Journal of Empirical Legal Studies* 2, 1 (2005), 87–123.
- [32] Frank P McKenna. 1993. It won’t happen to me: Unrealistic optimism or illusion of control? *British Journal of Psychology* 84, 1 (1993), 39–50.
- [33] Peter R Monge and Noshir S Contractor. 2003. *Theories of Communication Networks*. Oxford University Press, USA.
- [34] Elinor Ostrom. 1990. *Governing the commons*. Cambridge University Press.
- [35] Gabriele Paolacci and Jesse Chandler. 2014. Inside the Turk Understanding Mechanical Turk as a Participant Pool. *Current Directions in Psychological Science* 23, 3 (June 2014), 184–188.
- [36] Jean-Jacques Rousseau. 1984. *A discourse on inequality*. Penguin.
- [37] Thomas C. Schelling. 1960. *The strategy of conflict*. Cambridge, MA (1960).
- [38] James A Shepperd, William MP Klein, Erika A Waters, and Neil D Weinstein. 2013. Taking stock of unrealistic optimism. *Perspectives on Psychological Science* 8, 4 (2013), 395–411.
- [39] Neilson Voyné Smith et al. 1982. *Mutual knowledge*. Academic Press London.
- [40] Siddharth Suri and Duncan J Watts. 2011. Cooperation and contagion in web-based, networked public goods experiments. *PLoS One* 6, 3 (2011), e16836.
- [41] Kyle A Thomas, Peter DeScioli, Omar Sultan Haque, and Steven Pinker. 2014. The psychology of coordination and common knowledge. *Journal of Personality and Social Psychology* 107, 4 (2014), 657.
- [42] Robb Willer, Ko Kuwabara, and Michael W Macy. 2009. The False Enforcement of Unpopular Norms. *Amer. J. Sociology* 115, 2 (2009), 451–490.