Collective Action Through Common Knowledge Using A Facebook Model

Gizem Korkmaz
Virginia Bioinformatics Institute, Virginia Tech
Blacksburg, VA 24061
gkorkmaz@vbi.vt.edu

Chris J. Kuhlman
Virginia Bioinformatics Institute, Virginia Tech
Blacksburg, VA 24061
ckuhlman@vbi.vt.edu

Achla Marathe
Virginia Bioinformatics Institute, Virginia Tech
Blacksburg, VA 24061
amarathe@vbi.vt.edu

Madhav V. Marathe
Virginia Bioinformatics Institute, Virginia Tech
Blacksburg, VA 24061
mmarathe@vbi.vt.edu

Fernando Vega-Redondo
Department of Decision Sciences, Bocconi University
20136 Milan, Italy
fernando.vega@unibocconi.it

ABSTRACT

We develop a dynamic game-theoretic model of information (contagion) propagation using a Facebook-type of communication. The model accounts for information posted on a member's wall or timeline by her friends, which can be read by all friends. This communication facilitates coordination by creating common knowledge among users. We illustrate subtle features of the model, which generalize a host of influence-based contagion mechanisms, and prove characteristics of its dynamics. We show that a complete bipartite graph within a certain group of agents is a necessary and sufficient condition for common knowledge of relevant information to arise among that group's members. Finally, we illustrate the behavior of our model through simulation, and compare it to the classic diffusion model and Chwe's model of common knowledge, using a real Facebook network, a high school social network and a friendship network.

Categories and Subject Descriptors

1.6.3 [**Simulation and Modeling**]: Applications

General Terms

Theory, Experimentation

Keywords

Facebook, Common knowledge, Collective action

1. INTRODUCTION

The use of social networking sites (especially Facebook, Twitter and Youtube) was a distinctive feature of the recent wave of uprisings against authoritarian regimes in the Arab world, and of social actions in Western countries following a recent financial crisis (e.g. Occupy Wall Street). Information sharing prior to, as well as during, mass demonstrations and street protests proved to be essential for the success of these protests [8]. The significant impacts of these events and others (e.g., political upheaval, regional instability) motivate the construction of models to understand the mechanisms by which information spreads through social media, and their consequences. In this paper, we develop a dynamic game-theoretic model of "on-set of revolutions" that focuses on the local spread of information through online social networks (we use this context for definiteness; the model is equally applicable to other contexts). We investigate how the flow of information through online communication channels can facilitate collective behavior. Our approach is that collective action is a coordination problem that takes place locally. In the proposed coordination game, we consider a population of heterogeneous agents who differ in their willingness to participate in protests. An individual wants to participate only if joined by others. The number of participants at or above which an individual would choose to participate defines the individual's *threshold*. Coordination requires that people know each others' thresholds and that this information is common knowledge among a sufficient number of people. We study how network structure affects common knowledge and participation decisions.

Works to date overwhelmingly rely on what we term *classic diffusion* models (e.g., [18,22]), which includes complex contagions (e.g. [4]). By way of comparison, in deterministic classic diffusion models, an agent *unilaterally* transitions state (e.g., from non-participating to participating) if at least a threshold number of its neighbors are participating. While our approach incorporates this model in a more generalized form, another of our contributions is a *joint* or *collective* mechanism whereby agents transition state together. One implication of this mechanism is that a contagion can germinate and spread organically without being explicitly seeded, whereas classic threshold models require seeds (for thresholds $\theta > 0$). The main contributions of this research are listed below.

1. A Model of Contagion Dynamics on Facebook. Our model incorporates the concept of a Facebook posting as the information-sharing mechanism in which friends post their their willingness—or threshold—to participate and current state on each others' timelines or "walls." A wall acts as

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

a localized repository where information can be shared by its owner and her friends. This allows individuals to learn about their direct neighbors as well as distance-2 neighbors. Consequently, two individuals can influence each other even though there is no edge between them, if they have a common neighbor. This model incorporates contagion propagation via (*i*) classic threshold diffusion [4, 7, 9], where an agent receives contagion from all of its direct neighbors; and (*ii*) common knowledge that can lead to collective action [5]. Crucially, however, our model goes beyond each of these mechanisms as currently utilized; these features are described in Section 2 on related work. Further, the mechanisms are fundamentally different from those arising in Twitter (e.g., [18]). The model also captures the dynamics produced by some genetic algorithms [12] for the appropriate network structures. We provide a formal description of this model and illustrative examples (see Section 3). To the best of our knowledge, this is the first model of its kind, both in terms of the mechanisms used and the modeling of Facebook wall postings.

2. Knowledge - Learning. A key feature of social networks prevailing in the real world is that agents only have local information about the network. Consequently, we restrict the information on the topology of the network to a local level, which is determined by the information sharing mechanism considered. In a dynamic extension, we examine how further information is revealed if agents observe the actions taken by their neighbors, and how this additional information is used to make inferences about the number of revolters. This allows us to study large scale movements that unfold over time despite a small number or no participants initially. We show a contagion paradox that results from this learning process. Specifically, if all agents possess the same threshold $\theta = d_{max} + 1$, where d_{max} is the maximum agent degree in the network, the contagion will propagate through the network. This is a consequence of the learning and inferences made by agents that is lacking in other classic diffusion and common knowledge models.

3. Characterization of Common Knowledge. In game-theoretic contexts, agents' behaviors depend on whether some piece of relevant information is common knowledge. In our framework, this information concerns not only the network, but also the types of the individuals and their actions. In order to address this issue, we use techniques from the epistemic approach developed within game theory. We formulate how information is diffused through networks and how common knowledge is obtained. In particular, we characterize the structure of sub-networks in which all agents have common knowledge about their types and actions (states). We prove necessary and sufficient conditions for the existence of common knowledge within a subgraph of a networked population. We find that for a given network if there exists a set of agents who share common knowledge, then there must exist a complete bipartite subgraph formed by these agents such that all edges in the complete bipartite graph are also edges in the original graph.

4. Comparisons of Contagion Spread by Various Mechanisms. We extend the theoretical formulations by introducing stochasticity through people's Facebook usage rates and times of participation, but also examine the deterministic model. We perform simulations of contagion propagation on a mined Facebook network, a High School social contact network and a mutual friendship network, with different structures. We compare contagion spreading via classic diffusion [4], Chwe's common knowledge model [5], and our model and show significant differences among them in spreading contagion. We describe the reasons for these differences. We also demonstrate effects of heterogeneous agent thresholds. Our common knowledge model, unique to Facebook, raises interesting speculations about the use of Facebook and Twitter in social unrest. Most news articles treat these two media as similar in their roles in social unrest. Our work raises the interesting possibility that these social media might have *somewhat distinct* and *complementary* roles.

2. RELATED WORK

Coordination and collective action have been widely studied in many contexts (see review [16]). Our model builds upon the seminal works of Granovetter [9] and Schelling [19], studying tipping or threshold behavior in groups. We propose that the more people participate, the more likely it is that a given individual will choose to participate in the collective action. We abstract from the free-riding issues addressed by Olson [17]. A broad and growing interdisciplinary literature suggests that participation in collective actions depends on the social structure regarding both pattern of connections and the position of individuals within the networks (see [21] and references therein).

Our work is closely related to Chwe [5, 6], who considers social structure and individual incentives together in order to study which network structures are conducive to coordination. He presents a coordination game of incomplete information and models social structure as a communication network through which people tell their willingness to participate. He finds structured networks that are built upon a hierarchy of cliques to be the minimal network structure that allows everyone to revolt when the network itself is common knowledge. First of all, Facebook-type communication allows for common knowledge to be attainable for a larger number of individuals and in a larger variety of network structures, rather than cliques. Second, and more importantly, in our model the network structure is only locally known by agents. Finally, the dynamics in Chwe [6] are such that agents learn thresholds and actions of agents at progressively greater geodesic distances, and in a few time steps, each agent knows the information of a large fraction of a population. In contrast, in our model, agents can only observe the actions taken by their neighbors, and make inferences. We show how the limited information on the network structure, and the rich process of updating, tailored to the topology of the network, affect common knowledge and generate different results than in Chwe [6].

There are many models of information diffusion and contagion, beyond common knowledge models. Among these are deterministic and stochastic threshold models (e.g., [4, 9]), and game theoretic approaches (e.g., [23]). Twitter [18] and Facebook [22] models have been devised, although the Facebook model employs a Twitter-like broadcast mechanism and does not account for Facebook walls. Hence it has no notion of common knowledge. A deterministic threshold model was used to explain a Spain 2011 protest event, based on Twitter data [8]. Our model accounts for classic diffusion through distance-2, which is a unique feature, and can be readily extended to incorporate other probabilistic models such as that in [18].

3. MODEL OF CONTAGION DYNAMICS

3.1 Preliminaries

There is a finite set of people $N = \{1, 2, ..., n\}$ and each person $i \in N$ chooses an action $a_i \in \{r, s\}$, where r is '*revolt*,' the '*risky*' action, and s is '*stay at home*,' the '*safe*' action. Each person i has an idiosyncratic private threshold $\theta_i \in \{1, 2, ..., n+1\}$; a person wants to revolt only if the total number of people who revolt is greater or equal to his threshold. Given person i's threshold θ_i and everyone's actions $a = (a_1, a_2, ... a_n)$, her utility is given by

$$U_i(\theta_i, a_i, a_{-i}) = \begin{cases} 0, & \text{if } a_i = s \\ 1, & \text{if } a_i = r \wedge \#\{j \in N \mid a_j = r\} \geq \theta_i \\ -z, & \text{if } a_i = r \wedge \#\{j \in N \mid a_j = r\} < \theta_i \end{cases}$$

where $-z < 0$ is the penalty he gets, if he revolts and not enough people join him. Thus, a person will revolt as long as he is sure that there is a sufficient number of people revolting. A person always gets utility 0 by staying at home. When he revolts, he gets utility 1 if the total number of people revolting (including himself) is at least θ_i.

The communication network is undirected and is represented by $G(N, L)$, where L denotes the set of links. The communication technology we consider here is "facebook-type" communication in which people write on each others' "walls." Let the pairwise relationship between two agents be represented by the binary variable $g_{ij} \in \{0, 1\}$. When $g_{ij} = 1$, two agents are linked; when $g_{ij} = 0$ there is no link between i and j. $g_{ij} = 1$ implies that person i writes his threshold, θ_i, (and action, a_i) on person j's *wall* (and vice versa). This post is observed by j's neighbors since they have access to the *wall* of j. Let $N_i(g) = \{j \in N \setminus \{i\} : g_{ij} = 1\}$ be the set of person i's neighbors and let $\eta_i(g) = \#N_i(g)$ be the cardinality of this set. The network structure is not observed by everyone, but person i knows about the thresholds and actions of the people in his '*ball*' which is denoted by $B_i = \{j \in N_i^2\}$, where N_i^2 is the set of all neighbors within distance-2 of i.

The timing of the game

$t = 0$. Nature determines people's preferences (thresholds to revolt), which are private information.

$t = 1$. People communicate their thresholds and action plans with their direct neighbors. They then choose their own actions, r or s, (decide whether to revolt or not).

$t > 1$. People who have chosen s in the previous periods observe the thresholds and past actions of the people in their ball and choose their own actions, r or s[1].

3.2 Common Knowledge

In game theory, information is represented by an *event* and an agent knows an event if he believes it with probability one. Informally, an event E is *mutual knowledge* among a set of agents if each agent knows that E. Mutual knowledge by itself implies nothing about what, if any, knowledge anyone attributes to anyone else. In interactive situations, not only individual and independent knowledge of "fundamentals" (first-order knowledge) is important, but also higher-order knowledge, i.e., knowledge about others' knowledge. A knowledge is common among a group of agents if everyone has it, everyone knows that everyone has it, everyone knows that everyone knows that everyone has it, and so on *ad infinitum* [2,14]. Common knowledge is shown to be a central concept and often a necessary condition for coordination.[2]

Although there are a number of ways in which the concept of common knowledge can be formalized, we adopt the set-theoretic approach described below.[3]

Events are subsets of a set Ω of possible worlds. A distinct actual world ω_α is an element of Ω. An event $E \subseteq \Omega$ is realized (or is true) if the actual world $\omega_\alpha \in E$. The event should be consistent with the actual state.

What an agent i knows about the set of possible worlds is stated formally in terms of a *knowledge operator* $\mathbf{K}_i(E)$. Given an event $E \subseteq \Omega$, $\mathbf{K}_i(E)$ denotes a new event, corresponding to the set of possible worlds in which agent i knows that E obtains. $\mathbf{K}_i(E)$ is read as 'i knows (that) E (is the case)'.

Definition. Agent i's *possibility set* $P_i(\omega)$ at $\omega \in \Omega$ is defined as

$$P_i(\omega) \equiv \bigcap \{E \mid \omega \in \mathbf{K}_i(E)\}.$$

The collection of sets $\mathcal{P}_i = \bigcup_{\omega \in \Omega} P_i(\omega)$ is i's *private information partition.*

$P_i(\omega)$ is the intersection of all events that i knows at ω, $P_i(\omega)$ is the smallest event in Ω that i knows at ω. In other words, $P_i(\omega)$ is the most specific information that i has about the possible world ω. The elements of i's information system represent what i knows immediately at a possible world. We can also write player i's knowledge function as:

$$\mathbf{K}_i(E) = \{\omega \in \Omega \mid P_i(\omega) \subseteq E\}.$$

We can now define mutual and common knowledge as follows:

Definition. Let a set Ω of possible worlds together with a set of agents N be given.

1. The event that E is (*first order*) *mutual knowledge* for the agents of N, $\mathbf{K}_N^1(E)$, is the set defined by

$$\mathbf{K}_N^1(E) \equiv \bigcap_{i \in N} \mathbf{K}_i(E).$$

2. The event that E is m^{th} *order mutual knowledge* among the agents of N, $\mathbf{K}_N^m(E)$, is defined recursively as the set

$$\mathbf{K}_N^m(E) \equiv \bigcap_{i \in N} \mathbf{K}_i(\mathbf{K}_N^{m-1}(E)).$$

3. The event that E is *common knowledge* among the agents of N, $\mathbf{K}_N^*(E)$, is defined as the set

$$\mathbf{K}_N^*(E) \equiv \bigcap_{m=1}^{\infty} \mathbf{K}_N^m(E).$$

It can be shown that:

[1]In this framework, we only allow for moving from s to r. One can think about introducing the possibility of changing the action from r to s.

[2]See survey [11].

[3]We follow the Stanford Encyclopedia of Philosophy [24] for definitions and notations regarding common knowledge. Related axioms and their properties can be found in [24].

(1) If $\omega \in \mathbf{K}^*_N(E)$ and $E \subseteq F$, then $\omega \in \mathbf{K}^*_N(F)$.

(2) $\omega \in \mathbf{K}^m_N(E)$ if and only if for all agents $i_1, i_2, ..., i_m \in N$, $\omega \in \mathbf{K}_{i_1}\mathbf{K}_{i_2}..\mathbf{K}_{i_m}(E)$.

Hence, $\omega \in \mathbf{K}^*_N(E)$ if and only if (2) is the case for each $m \geq 1$.[4]

In our framework, the set of possible states (and hence the events) are defined by the thresholds and the network structure since the links between agents are not observed by everyone in the network. The set of states is given by $\Omega = \Theta^n \times G$ with $G = \{0,1\}^{C^n_2}$ where $C^n_2 = n(n-1)/2$ is the number of 2-combinations of n nodes. A state can be formally written as $\omega = [(\theta_i)_{i\in N}, (g_{ij})_{i<j}]$. Thus, a state will be a $(n + C^n_2)$-*tuple*. For example, we can write one possible state for a network of 3 people with $g_{12} = 0, g_{13} = 0, g_{23} = 1$ as $\omega = (2, 2, 2, 0, 0, 1)$ where $\theta_i = 2$ for all i.

Common knowledge among a set of people implies that: *they know each others' thresholds (and actions) and they know that they know their thresholds (and actions). Therefore they can count on each other.* In other words, it is common knowledge that there is sufficient discontent. Figure 1 below highlights the importance of network structure and common knowledge in facilitating coordination.

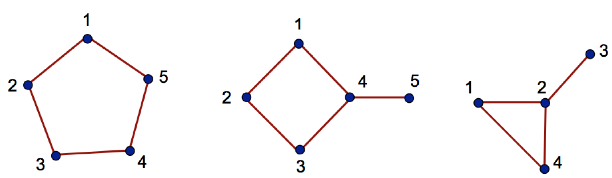

Figure 1: Pentagon, kite and half-kite with $\theta_i = 4$ $\forall i$.

In the pentagon, person 1 knows the thresholds of person 2 and person 5 directly since they write on each others' walls. He knows the thresholds of person 3 and person 4 through the walls of person 2 and person 5, respectively. Thus, he knows the thresholds of everyone and he knows that everyone knows his threshold (by symmetry). However, he has limited information on the network structure. Person 1 cannot observe the link between person 3 and person 4; thus he does not know whether they communicate. Moreover, he does not know that agents 2 and 4 (similarly agents 3 and 5) have a common neighbor, hence they can observe each other's thresholds. Although person 1 knows that all of them will benefit if they revolt, he does not know whether they know it. Therefore, he cannot count on them. Although everyone knows everyone's thresholds, and despite the fact that there is a sufficient number of people who would get positive payoffs if they revolted (i.e., $n > \theta_i \quad \forall i$), no one revolts in the pentagon.

In the kite, as with the pentagon, we observe that agents 1, 3 and 4 know the thresholds of everyone (everyone is within distance-2). Although person 2 and person 5 do not know about each other (since they do not have a common friend), there is common knowledge of thresholds among agents 1, 2, 3, 4 and among people 1, 3, 4 and 5. Since all agents have threshold 4, the people in each group who share common knowledge about each others' thresholds know that if they jointly revolt, each will gain. Therefore, everyone revolts.

Note that the difference between the pentagon and kite is only structural, since both have the same number of people with same thresholds, and the same number of links. In Chwe [6], if people know the thresholds of everyone, they can immediately observe the actual state of the world and revolt so long as there is sufficient discontent. If we assume the network structure is also common knowledge (as in [6]), everyone would revolt in the pentagon. In that case, person 1 would also know that there is common knowledge among everyone, which is more than enough for all people to revolt since they have threshold $\theta_i = 4$.

3.3 Characterization of Common Knowledge

We have assumed that the agents learn each others' thresholds if and only if they are within distance-2 of each other. Formally, we state this assumption as follows:

Axiom. Consider the actual state $\hat{\omega} \in \Omega$ and the event $\hat{E} = \{\omega \in \Omega : \omega = [(\theta_i)_{i\in N}, (g_{ij})_{i<j}]$ with $(\theta_i, \theta_j) = (\hat{\theta}_i, \hat{\theta}_j)\}$. We assume that $\hat{\omega} \in \mathbf{K}_i(\hat{E}) \cap \mathbf{K}_j(\hat{E})$ if and only if $i \in B_j$ and $j \in B_i$ for $i, j \in N$.

This suggests that at the actual state in which $(\theta_i, \theta_j) = (\hat{\theta}_i, \hat{\theta}_j)$, this event is mutually known by agents i and j if and only if they are within distance-2. Moreover, it allows us to prove that 3 agents i, j, l that form a star network share common knowledge of thresholds. We formally state this result as a lemma below:

Lemma. Given the actual state $\bar{\omega} \in \Omega$ and the event $\bar{E} = \{\omega \in \Omega : \omega = [(\theta_i)_{i\in N}, (g_{ij})_{i<j}]$ with $(\theta_i, \theta_j, \theta_l) = (\bar{\theta}_i, \bar{\theta}_j, \bar{\theta}_l)$ and $j, l \in N_i,\}$, $\bar{\omega} \in \mathbf{K}^*_M(\bar{E})$ where $M = \{i, j, l\}$.

We generalize this result and characterize the necessary and sufficient conditions for a subset of agents to have common knowledge about the thresholds of everyone in the subset. Therefore, we can state the following result:

Theorem 1. *Given the actual state $\hat{\omega} \in \Omega$ and the event $\hat{E} \equiv \{\omega \in \Omega : \omega = [(\theta_i)_{i\in N}, (g_{ij})_{i<j}]$ with $(\theta_i)_{i\in M\subseteq N} = (\hat{\theta}_i)_{i\in M\subseteq N}\}$, $\hat{\omega} \in \mathbf{K}_i\mathbf{K}_j...\mathbf{K}_m(\hat{E})$ for all $i, j, m \in M \subseteq N$ if and only if*

(1) $i \in B_j \quad \forall i, j \in M \subseteq N$,

(2) $\forall \{j, l\} \in L$, either (a) $\{i, j\} \in L \vee \{i, l\} \in L$ or (b) $\exists k : \{i, k\}, \{j, k\}, \{l, k\} \in L \quad \forall i, j, k, l \in M \subseteq N$.

Condition (1) implies that when $i \in B_j$ (and $j \in B_i$) $\forall i, j \in M$, people either know each others' thresholds directly, or indirectly through the wall of a common friend.

Condition (2) suggests that for every link $\{j, l\} \in L$ with $j, l \in M$, (2a) every other node $i \in M$ should be connected to one of the nodes of that link ($\{i, j\} \in L$ or $\{i, l\} \in L$) or (2b) there must be another node k that all three are connected to ($\{i, k\} \in L$, $\{j, k\} \in L$ and $\{l, k\} \in L$). In addition to knowing about the others' thresholds, this condition ensures that an agent also knows that the others know about each others' thresholds.

As we have discussed above, in the pentagon, the maximum distance between any two nodes is 2, hence everyone knows the thresholds of everyone else (Condition 1). However, in order for an agent to know whether others know about each other, the agent would need more information about the network structure. Condition (2a) suggests that

[4]The condition that $\omega \in \mathbf{K}_{i_1}\mathbf{K}_{i_2}...\mathbf{K}_{i_m}(E)$ for all $m \geq 1$ and all $i_1, i_2, ..., i_m \in N$ is Schiffer's definition of common knowledge [20], which is the one most often used in the literature.

person 1 should be connected to either person 3 or person 4, so that he can observe that they communicate and know about each other, which in turn allows person 1 to know that person 2 knows about person 4 and that person 5 knows about person 3. Similarly, person 2 should connect either to person 4 or to person 5, and so on. Thus, the event is common knowledge in the actual state among all agents.

Finally, we observe that condition (2a) is sufficient but not necessary to have common knowledge among a set of agents. We can have structures in which (2a) is not satisfied but the thresholds are still common knowledge. In the half-kite in Figure 1, we observe that person 3 is not linked to person 1 or person 4 directly (violating condition 2a), and thus, he cannot observe the link between them. However, person 3 knows that they know about each other through person 2, to whom all of them are linked. Condition (2b) captures these cases, in which agents do not need to know about all links but they need to know that the others' know about each other so that the thresholds are common knowledge among everyone. We now state the necessary and sufficient graph substructure that produces common knowledge.

Theorem 2. *Given a graph $G(N, L)$, if there exists a set of people $M \subseteq N$ such that* $\hat{\omega} \in \mathbf{K}^*_M(\hat{E}) = \mathbf{K}_i\mathbf{K}_j..\mathbf{K}_m(\hat{E}) \quad \forall i, j, m \in M$, *then there must exist a sub-graph formed by this set M that is a complete bipartite graph.*

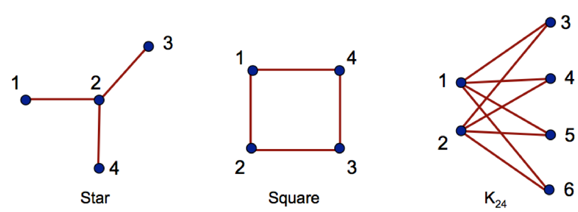

Figure 2: Complete bipartite graphs.

Figure 2 illustrates three examples of complete bipartite graphs, including star and square, in which the thresholds are commonly known among all agents. The star is the extreme case in which common knowledge is obtained only through the hub's wall.

3.4 Static Framework

Initially everyone is in state s, i.e., $a_i(t = 0) = s \quad \forall i \in N$.

Deviations: A subset of agents $M \subseteq N$ revolts at $t = 1$ if

(1) $\hat{\omega} \in \mathbf{K}^*_M(\hat{E}) = \mathbf{K}_i\mathbf{K}_j..\mathbf{K}_m(\hat{E})$ for all $i, j, m \in M$,

(2) $\theta_i \leq \#M \quad \forall i \in M$.

Therefore, the subset M deviates when the number of people in M who share the common knowledge of the thresholds is sufficiently high. In other words, they revolt when it is common knowledge that if all $i \in M$ deviate, then everyone gains.

Equilibrium: We say that an action profile $a^* = (a^*_1, a^*_2, ..., a^*_n)$ is an equilibrium if and only if

$\nexists M \subset N$ such that

(1) $\hat{\omega} \in \mathbf{K}^*_M(\hat{E}) = \mathbf{K}_i\mathbf{K}_j..\mathbf{K}_m(\hat{E})$ for all $i, j, m \in M$, and

(2) $\forall i \in M \quad U_i(r, a^*_{-M}) \geq U_i(a^*)$ with strict inequality for some $j \in M$.

Having defined deviations and the equilibrium for the static case, we can re-analyze some of the networks with $\theta_i = 4 \ \forall i$. In the pentagon, the thresholds are mutually known by all agents but not common. Thus, they do not revolt, i.e., $a^* = (s, s, s, s, s)$. In the kite, we have $M_1 = \{1, 2, 3, 4\}$ and $M_2 = \{1, 3, 4, 5\}$ such that the thresholds are common knowledge in each subset and $\theta_i \leq \#M \quad \forall i \in M$. Agents in the same subset know that if they jointly deviate and revolt, everyone in that set will gain. Therefore, everyone revolts in the kite, i.e., $a^* = (r, r, r, r, r)$.

Following Theorem 2, we know that in complete bipartite graphs, there is common knowledge of thresholds among all of the people. Note that in complete bipartite graphs, the number of people for whom the thresholds are common knowledge is not limited. As long as the structure remains the same, infinitely many agents can have common knowledge about the state of the world. In this case, whenever there exists a subset $M \subseteq N$ for which $\theta_i \leq \#M \quad \forall i \in M$, $a_i = r$ for all $i \in M$.

3.5 Dynamic Framework

We introduce dynamics to the model by assuming that people post also their actions on their direct neighbors' *walls*. In each period, agents obtain local information about the past actions of their neighbors in distance-2 in addition to the thresholds. We argue that once people learn the past actions of their neighbors, they can make inferences about the number of revolters. In other words, learning that someone with threshold $\bar{\theta}$ revolted at $t-1$ reveals to an agent that the number of revolters must be at least $\bar{\theta}$, since he knows that people do not revolt unless they know that there is sufficient number of revolters. We introduce more sophisticated behavior in the dynamic framework and we need to assume that it is commonly known that everyone is sophisticated enough to make these inferences.

The set of states is given by $\Omega(t) = \Theta^n \times A^n(t-1) \times G$ with $G = \{0, 1\}^{C^n_2}$ and $A(t-1) = \{r, s\}$. Formally, a state can be written as $\omega = [(\theta_i)_{i \in N}, (a_i)_{i \in N}, (g_{ij})_{i<j}]$. Thus, a state will be a $(2n + C^n_2)$*-tuple*.

The Law of Motion $(a_i(t))$

We can write the law of motion for $t > 1$ formally as:

$a_{i \in M}(t > 1) = r \iff \exists M \subseteq N :$

(1) $\hat{\omega} \in \mathbf{K}^*_M(\hat{E}) = \mathbf{K}_i\mathbf{K}_j..\mathbf{K}_m(\hat{E}) \ \forall i, j, m \in M$,

(2) $\exists K \subseteq N$ with $a_k(t-1) = r$ and $\hat{\omega} \in \mathbf{K}^*_M(\hat{\hat{E}}) \ \forall i, j, m \in M$ and $k \in K$: $\theta_i - \max\{\max\{\theta_k\}_{k \in K}, \#K\} \leq \#M \quad \forall i \in M$

where $\hat{\hat{E}} = \{\omega \in \Omega : \omega = [(\theta_i)_{i \in N}, (g_{ij})_{i<j}]$ with $(\theta_i)_{i \in M \subseteq N} = (\hat{\theta}_i)_{i \in M \subseteq N}$ and $(\theta_k)_{k \in K \subseteq N} = (\hat{\hat{\theta}}_k)_{k \in K \subseteq N}\}$.

This implies that the agents have common knowledge about (1) each others' thresholds, and (2) a set of agents who revolted at $t - 1$. In addition, either the number of agents in the latter set is sufficiently high, or at least one of them has a sufficiently high threshold.

Therefore, (i) an agent with $\theta_i = \bar{\theta}$ can unilaterally revolt once she observes that another agent with $\bar{\theta} - 1$ has revolted in the previous period, (ii) an agent with $\theta_i = \bar{\theta}$ can unilaterally revolt if she observes that the number of people that have revolted is at least $\bar{\theta} - 1$, (iii) agents can jointly revolt if they observe that another agent with a sufficiently high threshold has revolted and this information is common knowledge among these agents. Suppose that agents A and B have threshold 5 and they have common knowledge about C who has revolted. C's threshold must be at least 3 for A and B to revolt. If C's threshold is 3, A and B need to coordinate; if it is 4, A and B can unilaterally revolt. The next result is an example of how our common knowledge model differs from classic diffusion.

Proposition 1 *Given any connected social network, with d_{max} representing the maximum degree of any agent in the network. Let every agent have threshold $\theta^* = d_{max} + 1$. Then, the contagion will spread to every agent in the network. Moreover, the number of time steps required to reach this condition is no more than the diameter of the graph.*

This result shows a stark contrast between our model and complex contagion classic diffusion models. Under the classic diffusion model, the contagion will never spread beyond the seed set. The results on contagion spread also hold for the stochastic extension of our model considered in Section 4.

4. EXPERIMENTAL METHODOLOGY AND NETWORKS

4.1 Computational Challenges

Simulation of network dynamics for our model and Chwe's common knowledge model requires computations that are more challenging than those for typical influence models (e.g., [4, 18]), where information about nearest neighbors is the only information an agent needs to compute its next state. Each of the common knowledge models requires identification of subgraphs that are computationally intractable to find. We refer to our model as the CKF model (common knowledge through Facebook) in the remainder of the paper.

In Chwe's model, at time $t = 0$, common knowledge exists among all agents which form a clique substructure in a network. To computationally identify all sets of common knowledge, we need to find all maximal cliques in the network; this is the so-called Maximal Clique Enumeration (MCE) problem and is NP-hard [13]. Moreover, at each successive time $t > 0$, each agent v_i increases the number of agents for which it knows states and thresholds by increasing its "view distance" by 1. Hence, at each time, the subgraphs of interest for evaluating state transition of agent v_i are growing. The effect is that the number of agents that comprise these "growing" cliques increases rapidly. (We will see in the simulation section that this modeling choice has large implications for computations of contagion spread.) We use a combination of the Bron-Kerbosch algorithm [3] and a graph densification algorithm to approximately compute all maximal cliques in a graph up through time $t = 5$, which is a reasonable approximation considering graph diameters (see below).

For our CKF model, we require maximal (in terms of agents) complete bipartite graphs (bicliques). The problem of finding *all* maximal bicliques in a graph is NP-hard [1]. We use the algorithm [15] to compute maximal bicliques.

4.2 Simulation

To determine whether agents transition state $s \rightarrow r$, in the CKF model, we need to consider three steps:

1. State transitions owing to common knowledge sets. Consider each unique common knowledge set M of agents that must form a maximal complete bipartite subgraph B. Analyze the set M and, consecutively, subsets of M by removing, in turn, all agents with the current maximum threshold, until a subset of M meets the criterion for state transition, or until the remaining subset is the empty set. In the latter case, no transitions occur.

2. State transitions owing to distance-2 classical diffusion. Consider all of the agents v_j within distance-2 of v_i. Let $\eta_r(v_i)$ be the number of these agents in state r. Then v_i transitions to state r if $\eta_r(v_i) + 1 \geq \theta_i$, where the 1 is for v_i.

3. Agent participation. Each agent has a probability of participation p_p that describes whether it participates in the contagion process at each time t. If an agent participates, this means that it can change state, and it can contribute to the state change of other agents; if it is not participating, then it does neither. This probability reflects the fact that people are interacting through Facebook at different times; $p_p = 1$ corresponds to the deterministic model of Section 3.

These are the dynamics for our CKF model. Appropriate procedures are also implemented for the Chwe model [5], including loading-in progressively larger sets of cliques as simulation time progresses. In Chwe's model, the cliques at $t = 4$ are used for all $t \geq 5$.

4.3 Networks

Networks and characteristics are provided in Table 1, including a Facebook network (FB) [25], a social contact network that we generated from a high school in the New River Valley (NRV) in Southwest Virginia, and a mutual friendship network from Add-Health data (AH) [10]. All networks are considered as undirected. The headers in Table 1 refer to network name or type, number of nodes, number of edges, average degree, number of unique maximal bicliques, number of unique maximal cliques, graph diameter, and maximum degree.

Table 1: Networks and their characteristics.

Graph	n	m	d_{ave}	n_{bc}	n_c	**diam**	d_{max}
FB	43,953	182,384	8.30	258668	60294	18	223
NRV	769	4551	11.8	5752	1495	7	20
AH	2448	5277	4.31	1496	1140	10	10

4.4 Experimental Procedures

We simulate the spread of contagion through the three networks using Chwe, CKF, and classic threshold diffusion models. Inputs to the models are θ (or low and high thresholds for heterogeneous simulations), p_p, and the networks. For each set of conditions, we run either 20 or 50 diffusion instances. For the two common knowledge models, there are no seed nodes (i.e. no nodes are in state r initially) because these models can introduce a contagion organically. For the classic threshold model, we use random seeding. Each simulation starts at $t = 0$ and runs for a prescribed number of time steps. At each step, agents in state s are evaluated for transition to state r, and agent IDs with time and state are saved as output.

5. EXPERIMENTAL RESULTS

In this section, we use the term *affected* to indicate agents that have transitioned to state r, the revolting state.

Effect of participation probability and variance in diffusion instances. Figure 3 provides simulation results for the FB network using our CKF model, where all agents have $\theta = 9 \approx d_{ave}$. In the left plot, the fraction of network agents that transition to state r at each time step is plotted against simulation time (time units can be thought of as hours, days, or any other relevant unit). Each curve is the average of multiple simulations. For the deterministic model, where the participation probability $p_p = 1$, roughly 80% of agents transition to r at $t = 1$. This is a consequence of the large number of bicliques (n_{bc} in Table 1) in the FB network. As p_p decreases, the spike attenuates. There is no diffusion in the first 30 time units for $p_p = 0.01$. The corresponding cumulative fractions of affected agents are provided in the right plot. The curves with smaller rates of increase in fraction of affected agents (e.g., the light green curve for $p_p = 0.05$), result from a larger variance among diffusion instances in the time at which the contagion *initiates* via common knowledge. Although not shown, the variation in times at which contagion originates can be a factor of 100 across diffusion instances. Once the contagion organically appears, however, the rate of spreading is fast, with relatively little variance.

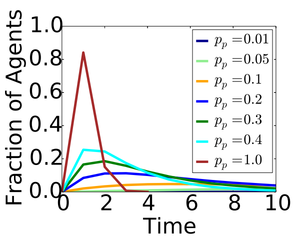

(a)

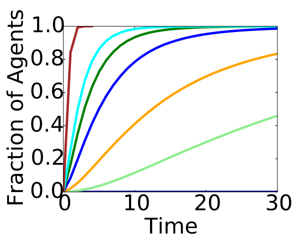

(b)

Figure 3: Our CKF model. (a) Average number of newly affected agents at each time and (b) average cumulative number of affected agents in time for the FB network. All agents have $\theta = 9 \approx d_{ave}$.

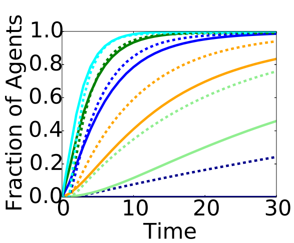

(a)

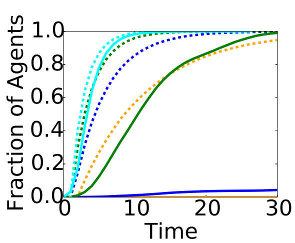

(b)

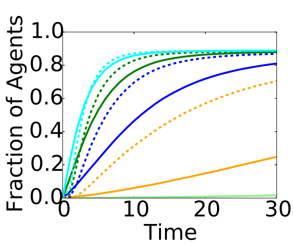

(c)

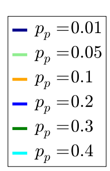

(d)

Figure 4: Comparison between CKF (solid) and Chwe (dashed) models. (a) FB with $\theta = 9 \approx d_{ave}$; (b) NRV with $\theta = 12 \approx d_{ave}$; (c) AH with $\theta = 4 \approx d_{ave}$; and (d) legend. Curves show the fraction of affected agents in time, and are averages over multiple simulation instances. Results show significant differences in contagion spread via common knowledge between models.

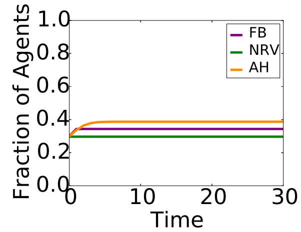

Figure 5: Contagion spread (in terms of fraction of agents affected) via classic threshold diffusion [4] in the three networks where $p_p = 1.0$, and for each network, 30% of the populations are *seeded* with the contagion. Thresholds for the FB, NRV, and AH networks are, respectively, 9, 12, and 4.

Comparisons of different dynamics models. Figure 4 provides data for the Chwe model (dashed curves) and our CKF model (solid curves) for a range of p_p for the three networks. Again we use uniform thresholds approximately equal to the average degrees of networks to make results easier to interpret. Differences are more pronounced at smaller p_p. For example, for FB with $p_p = 0.05$, the fraction of affected agents is 0.28 at $t = 20$ with our CKF model, versus 0.60 for the Chwe model. For NRV with $p_p = 0.2$ at $t = 20$, our CKF model produces a fraction of affected agents that is < 0.05, while the Chwe model predicts a fraction 0.97. The dominant factor at play is the assumption in the Chwe model that at each successive time, an agent learns the thresholds of agents at an increasing distance of 1. Thus, by time $t = 5$, for example, the Chwe model assumes that each agent v_i knows everything (including thresholds) about all agents within a geodesic distance of 6. This means that the number of agents with which v_i can form cliques (and hence common knowledge) grows very rapidly. For FB, for example, there is a two order of magnitude increase in the number of agents in cliques at $t = 2$ compared to that for $t = 0$. Moreover, as shown in Table 1, the diameters of these networks are not large, meaning that a distance of 3, 4, or 6 will encompass a large fraction of the agents: a local model at $t = 0$ quickly changes towards a global model even at small times. This leads to very high growth rates in numbers of affected agents. Our model, in contrast, remains a local model. As p_p increases, this difference between models diminishes because more agents participate and contagion spread is rapid in both models.

Figure 5 shows classic threshold diffusion for the three networks. We use the same thresholds as in Figure 4. For classic diffusion, seed agents are required. We aggressively assign 30% of agents to be in state r initially. Furthermore, $p_p = 1.0$. Yet, even under these conditions, there is minimal contagion spread, as illustrated by the fact that the time history curves remain essentially flat, indicating that very few agents are transitioning to state r. These results (and those in the previous figure) indicate the power of common knowledge to spread contagion, and suggest that Facebook and Twitter—which is typically modeled as a classic complex contagion mechanism [18]—might be complementary modes of social exchange, rather than redundant or interchangeable forums.

Sensitivity of contagion spread to thresholds. We perform experiments to determine how contagion spread changes with changes in both θ and p_p for the FB network. Simulations span 30 time units. We use a two-threshold system, where a given fraction $p(\theta_l)$ of agents possess a low threshold θ_l and $p(\theta_h) = 1 - p(\theta_l)$ fraction of agents possess a high threshold θ_h. We use threshold pairs $(p(\theta_l), p(\theta_h)) = (0.5., 0.5)$ and (0.8,0.2); the first set generates the same number of agents with each threshold (in expectation), while the second pair biases more agents to have the lesser threshold. In all cases $\theta_h = 225 = d_{max} + 2$, meaning that these agents cannot transition to state r via common knowledge when no agents are in state r. Low thresholds range from 5 to 210, and p_p is set to 0.2, 0.5, 0.8, and 1.0, in turn. The left plot in Figure 6 shows results for (0.5,0.5). As p_p increases, the lower threshold at which contagion can still propagate increases. These data are consolidated as the blue curve in the right plot, and the corresponding data for (0.8,0.2) are provided as the magenta curve. For a given threshold, p_p values that lay above the curves result in widespread diffusion. The bias in having more agents with the lower threshold decreases the critical value of p_p for generating widespread information transmission.

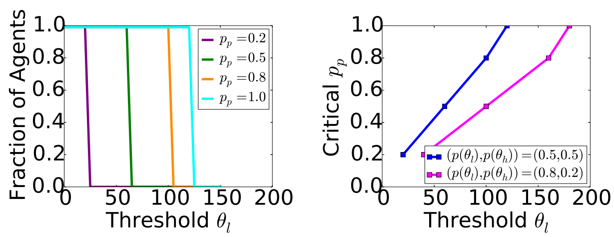

(a) (b)

Figure 6: Our CKF model and FB network. (a) Number of affected nodes versus θ_l for heterogeneous thresholds. Data exhibit sharp transitions from large to small contagion spread as the lower threshold increases, for different p_p. Here, $(p(\theta_l), p(\theta_h)) = (0.5, 0.5)$. (b) Critical participation probability as a function of θ_l for different $(p(\theta_l), p(\theta_h))$ pairs.

6. CONCLUSIONS

We introduce a contagion dynamics model that, to our knowledge, is the first of its kind. The model employs unique classic threshold diffusion and common knowledge mechanisms. Both mechanisms are based on the Facebook method of information transfer via users' wall postings. We provide a rigorous formal definition and theoretical results describing various aspects of system dynamics, and also provide insights from experiments on real networks.

7. ACKNOWLEDGMENTS

The authors thank the reviewers for their helpful comments and suggestions. This work has been partially supported by NSF Netse CNS-1011769, NSF ICES CCF-1216000, DTRA, R&D Grant HDTRA1-11-1-0016 and DTRA CNIMS Grant HDTRA1-11-D-0016-0001.

8. REFERENCES

[1] G. Alexe, S. Alexe, Y. Crama, S. Foldes, P. L. Hammer, and B. Simeone. Consensus algorithms for the generation of all maximal bicliques. *DAM*, 145:11–21, 2004.

[2] R. Aumann. Agreeing to disagree. *Ann. of Stat.*, 4:1236–1239, 1976.

[3] C. Bron and J. Kerbosch. Finding all cliques of an undirected graph. *CACM*, 16:575–577, 1973.

[4] D. Centola and M. Macy. Complex Contagions and the Weakness of Long Ties. *Am. J. of Soc.*, 113(3):702–734, 2007.

[5] M. S.-Y. Chwe. Structure and strategy in collective action. *Am. J. of Soc.*, 105:128–156, 1999.

[6] M. S.-Y. Chwe. Communication and coordination in social networks. *Rev. of Econ. Stud.*, 67:1–16, 2000.

[7] P. S. Dodds and D. J. Watts. A Generalized Model of Social and Biological Contagion. *J. of Theo. Bio.*, 232(4):587–604, 2005.

[8] S. Gonzalez-Bailon, J. Borge-Holthoefer, A. Rivero, and Y. Moreno. The Dynamics of Protest Recruitment Through an Online Network. *Nature*, pages 1–7, 2011.

[9] M. Granovetter. Threshold Models of Collective Behavior. *Am. J. of Soc.*, 83(6):1420–1443, 1978.

[10] K. M. Harris. The National Longitudinal Study of Adolescent Health (Add Health), Waves III, 2008. Carolina Population Center, UNC.

[11] F. Koessler. Common knowledge and interactive behaviors: A survey. *Euro. J. of Econ. and Soc. Sys*, 14:271–308, 2000.

[12] M. Lahiri and M. Cebrian. Genetic algorithm as general diffusion model for social networks. *(AAAI)*, 2010.

[13] E. L. Lawler, J. K. Lenstra, , and A. H. G. R. Kan. Generating all maximal independent sets: Np-hardness and polynomial-time algorithms. *SIAM J. on Comp.*, 9:558–565, 1980.

[14] R. Lewis. *Convention: A Philosophical Study*. Harvard U. Press, 1969.

[15] G. Liu, K. Sim, and J. Li. Efficient mining of large maximal bicliques. In *LNCS 4081, Conf. DaWak*, pages 437–448, 2006.

[16] P. Oliver. Formal models of collective actions. *Annual Review of Soc.*, 19:271–300, 1993.

[17] M. Olson. *The Logic of Collective Action*. Harvard U. Press, 1965.

[18] D. Romero, B. Meeder, and J. Kleinberg. Differences in the Mechanics of Information Diffusion. In *(WWW)*, 2011.

[19] T. Schelling. *Micromotives and Macrobehavior*. W. W. Norton and Company, 1978.

[20] S. Schiffer. *Meaning*. Oxford U. Press, 1972.

[21] D. Siegel. Social networks and collective action. *Am. J. of Pol. Sc.*, 53:122–138, 2009.

[22] E. Sun, I. Rosenn, C. A. Marlow, and T. M. Lento. Gesundheit! modeling contagion through facebook news feed. In *(ICWSM)*, 2009.

[23] J. Tsai, E. Bowring, S. Marsella, and M. Tambe. Empirical evaluation of computational fear contagion models in crowd dispersion. In *(JAAMAS)*, 2013.

[24] P. Vanderschraaf and G. Sillari. Common Knowledge. In *The Stanford Encyclopedia of Philosophy*. 2009.

[25] B. Viswanath, A. Mislove, M. Cha, and K. P. Gummadi. On the evolution of user interaction in facebook. In *(WOSN)*, 2009.