Despotic Regimes Instilling Fear in Citizens to Suppress Protests

Karen Kuhlman Amos
School of Social Work
Aurora University
Aurora, Illinois
karenkuhlman6@gmail.com

Chris J. Kuhlman
Biocomplexity Institute
University of Virginia
Charlottesville, Virginia
hugo3751@gmail.com

S. S. Ravi
Biocomplexity Institute
University of Virginia
Charlottesville, Virginia
ssravi0@gmail.com

***Abstract*—Fear of reprisals such as violence and punishment can inhibit citizens from speaking out, or make them more reluctant to act, in opposition to a repressive regime. Protests are one form of opposition, and their growth has been successfully modeled as an influence-based contagion process within a social network (representing a population). In these models, an individual joins a protest if a sufficient number of her neighbors has already joined. This required number of neighbors is often called a "threshold." In this study, we model a regime's ability to suppress protests by instilling fear in a subset of a population, and this fear is manifested by an increase in a person's threshold. We consider different social networks, numbers of seed nodes, and amounts of fear. Through simulations, we present several results. For example, we demonstrate that, for the objective of reducing the size of a protest, inducing fear can be more advantageous than removing nodes from a network.**

***Index Terms*—fear modeling, contagion blocking, threshold models, social networks**

I. Introduction

A. Background and Motivation

Despots and authoritarian regimes have been studied at least back to the time of Aristotle [3]. These regimes may impose their will through coercion, intimidation, violence, and threats against its citizens [3], [7], [8]. Such tactics are designed to control dissent by instilling fear in a population [19]. This means, for example, that citizens are less likely to protest against a tyrant for concern of harmful repercussions. Our goal is to study the effects of fear in inhibiting the process of individuals joining a protest in a population. Following [9], joining a protest is modeled as a contagion process where individuals respond to social influence; they are more likely to participate if their friends and neighbors do.

Our model represents a population as a social network. The nodes are individuals, and undirected edges denote interactions (e.g., friendship relations) between nodes. Nodes may be **inactive**, meaning that they are not involved in a protest, or **active**, meaning that they are protesting. We use a **threshold model** [4], [10], [14] to quantify whether a node transitions from the inactive state to the active state in a contagion process. An inactive node v_i transitions to the active state if at least a threshold θ_i number of its neighbors are already in the active state. Otherwise, v_i remains inactive. Once activated, a node remains active because the individual is assumed to be subjected to reprisals—she has reached "a point of no return."

Threshold models are used in many contexts, from modeling protests [5], [9], [13] to joining online platforms [17]. Other applications where threshold models are used in decision making are given in [10], [14], [18].

B. Our Approach to Fear Modeling and Analysis

Threshold models are a natural approach for modeling fear because θ_i represents v_i's intrinsic resistance to become active. According to [6], fear is a deterrent, and a deterrent, in turn, is a reluctance to participate. This reluctance can be quantified as a threshold increase.

By increasing v_i's threshold θ_i by some amount $\Delta\theta_i > 0$, v_i becomes more resistant to contagion adoption, which is our goal in modeling fear. That is, we use $\Delta\theta_i$ as a proxy for fear instilled in a person. We refer to these nodes as **blocking nodes** because they may inhibit protest adoption for themselves and their neighbors in a social network. We model successive increases in fear as successive increases in $\Delta\theta_i$. In addition, we evaluate different human contact networks G, different numbers of seed nodes for the contagion process (**seed nodes** are nodes that are protesting at the start of the protest contagion spreading process), and different amounts of fear. We quantify the effects of these input variables on the contagion process of protest participation. In particular, we compute the effects of these inputs on the fraction of nodes in a population that participates in a protest. We use agent-based modeling and simulation (ABMS) to quantify these effects. Section III provides the formal model.

C. Novelty and Significance of Our Work

To our knowledge, there is only one work that studies thwarting of contagion process by increasing node thresholds, i.e., increasing fear among population members to make them more reluctant to join a protest. Our work, and our goals, are significantly different from theirs. See Section II. Our approach of blocking protest contagions by fear (modeled as threshold increases) rather than by extreme violence/killing (modeled as removing nodes from a social network)—and comparing the two approaches—is significant because despots, in reality, make these calculations: to control by subtlety (inducing fear) or brutality (killing) [19]. Our work helps to *explain* and *understand* these decisions.

D. Contributions

1. Methodology to model the effects of fear in inhibiting participation in protests against a regime. Modeling fear through a principled approach, i.e., increasing thresholds of selected nodes of a social network, enables one to quantify the effects of fear on the spread of a dissent contagion. Similarly, changes in the numbers of seed nodes and in the amount of fear can be quantified in terms of the resulting fraction σ of nodes that join a protest. Ours is a global methodology in the sense that a set of blocking nodes is designed to halt a contagion that starts *anywhere* in a social network.

2. Effectiveness of threshold increases in thwarting the contagion of joining a protest. We show the following results. (i) Threshold increases in roughly 40% of nodes are required to completely halt contagion spreading (exact values depend on simulation conditions). (ii) We demonstrate that the more subtle approach of instilling fear in a population, as opposed to the more draconian approaches of *removing* dissenters from society by killing or incapacitating them, can be highly effective. Removing a node v_i from a graph is equivalent, from a network perspective, to setting its threshold to $\theta_i = d(v_i) + 1$, where $d(v_i)$ is the degree of v_i. For the same number k of blocking nodes, we demonstrate that for some conditions, increasing fear can produce the same reduced protest contagion spread fraction σ as node removal, but with a total cost (i.e., threshold increase) over all nodes that is only 38% or less of the total cost required to remove these nodes. (iii) These two conditions combined indicate that although a despot may have to engender fear in a significant fraction of a population, the *amount* of required total fear—the *cost* of fear—is significantly less than taking more overt actions such as abducting or killing people.

II. Related Work

Seigel [15], [16] has studied repression; this work is closest to our own. However, there are several differences between his work and ours. First, the networks studied in [15], [16] have 1000 nodes. In [16], the network is presumably a clique while those studied in [15] are 1000-node stylized small world, clustered, scale-free, and hierarchical networks. Our networks have been developed from sophisticated synthetic population procedures [2] and are much larger in size. Second, we use a Granovetter type threshold model (see Section III for a formal description) which is different from the ones used in [15], [16]. Third, and central to our entire work, is that when we apply fear in our model, we do so to a targeted subset of k people, and we vary the amount of fear over the k nodes. The selected set of blocking nodes (in size and composition) can make a big difference on the blocking efficacy of nodes and on the cost of the blocking nodes (see Section V). Seigel's work applies the fear model to all the nodes of a network.

III. Contagion Model

A. Graph Dynamical System (GDS)

Here, we define the discrete dynamical systems framework, called **graph dynamical system**, that we use to model and simulate the Granovetter, Schelling, and Centola and Macy [4], [10], [14] type of complex contagion.

A **graph dynamical system** (GDS) [1], [12], denoted S, is a four-tuple (G, K, F, R) where G is a (social) network $G(V, E)$ with vertex (or node or agent) set V and edge set E, with $n = |V|$ and $m = |E|$. An undirected edge between $v_i \in V$ and $v_j \in V$, denoted $\{v_i, v_j\} \in E$, means that v_i and v_j interact and therefore can influence each other. The set K of **vertex states** is the set of states of a node; an agent v_i is in exactly one state $s_i \in K$ at each time t. A **sequence of local functions** F, with $|F| = |V|$, provides a local function f_i for each $v_i \in V$. The function f_i specifies how agents update their states, and will be explained below. The **update scheme** R is the manner in which the f_i are invoked at each time. For this work, we assume a **synchronous** or **parallel** scheme whereby all agents update their states in parallel.

The **system state**, also called a **configuration**, $s = (s_1, s_2, \ldots, s_n)$ is the n-vector of all vertex states. Let N_i denote the closed neighborhood of v_i (i.e., N_i contains v_i and all of v_i's distance-1 neighbors in G). Let $s[v_i]$ be the vector of states of the vertices in N_i; the sequence is of length $d(v_i)+1$, where $d(v_i)$ is the degree of v_i in G. Then the next state s_i' of v_i is computed with the local function f_i on $s[v_i]$, the current states of the vertices in N_i; i.e., $s_i' = f_i(s[v_i])$. Introducing time t as a superscript, the state $s_i^{(t+1)}$ at time $t+1$ is computed by f_i and is given by $s_i^{(t+1)} = f_i(s^t[v_i])$.

B. Contagion Threshold Model

For the GDS model, we have specified G and R. Here, we specify K and F in the GDS in Section III-A. For the threshold contagion model, the node state set $K = \{0, 1\}$, where state 0 means that a node is *inactive* (i.e., does not possess the contagion, does not participate in a protest), and state 1 means that a node is *active* (i.e., does possess the contagion, does participate in a protest). Once a node reaches state 1, it remains in state 1; it never transitions back to state 0. This model is appropriate when state 1 involves a commitment that is not easily rescinded, such as acting against government wishes. This is a progressive threshold model [11].

We now specify the local functions $f_i \in F$ $(i \in \{1, 2, \ldots, n\})$. Let $n_{1,i}$ represent the number of v_i's distance-1 neighbors in G that are already in state 1 at time t. Also, each node v_i is assigned an integer $\theta_i \geq 0$ that represents its resistance or inertia to adopting a contagion (e.g., joining a protest). The next state $s_i^{(t+1)}$ of v_i at time $(t+1)$ is 1 if $n_{1,i} \geq \theta_i$ and 0 otherwise. Hence, if the number of neighbors of v_i in state 1 is at least θ_i, then v_i will transition 0 to 1. If each node v_i has its threshold $\theta_i = 1$, then the model is called a **simple contagion**. If at least one node v_i has $\theta_i > 1$, then the model is a **complex contagion**. Thus, complex contagions require social reinforcement [4].

C. Examples of GDS State Transitions and Blocking Nodes

Figure 1 shows a 7-node social network G. The forms of threshold T and state $s(t)$ vectors for the seven nodes are provided in the figure below the graph. Two cases of dynamics are given, and each case corresponds to a separate threshold vector T. The node thresholds for these two cases are the same, except that θ_4 increases from 2 (on the left) to 3 (on the right). That is, $\Delta\theta_4 = 1$ between cases 1 and 2; all other $\Delta\theta_i = 0$ between the two cases for $i \in \{1, 2, 3, 5, 6, 7\}$. These two cases use the initial system state, namely $s^0 = s(0) = (0, 0, 1, 0, 0, 1, 0)$; that is, only nodes v_3 and v_6 are in state 1. Since nodes v_3 and v_6 are initially in state 1, they are called **seed nodes**.

In both cases, v_7 has $\theta_7 = 1$. Since $n_{1,7} = 1$ (because v_7's neighbor v_6 is in state 1) and $\theta_7 = 1$, v_7 transitions from 0 to 1. It can be verified by iterating through the nodes that this is the only node that transitions at $t = 1$, in both cases. In a similar manner, one can verify the other successive system states shown in Figure 1. The **spread fraction** σ is the fraction of nodes that are active at the end of a simulation. Here, for each of cases 1 and 2, $\sigma = 5/7$.

If θ_4 increases further to 4 in a case 3 (not shown), so that $\Delta\theta_4 = 2$ between cases 1 and 3, then v_4 will not transition from 0 to 1, for the *specific initial conditions of this example*. This is because only three neighbors of v_4 can be active based on the seeding; they are v_3, v_5, and v_7.

To ensure that v_4 can *never* transition to state 1 (i.e., can never participate), regardless of the choice of seed nodes, we set $\theta_4 = d(v_4) + 1 = 6$ in a new case 4 (compared to $\theta_4 = 2$ in case 1). Hence, $\Delta\theta_4 = 6 - 2 = 4$ ensures that this node never transitions to state 1. Thus, increasing its threshold to $d(v_4) + 1$ ***effectively removes*** v_4 ***from the network***.

This is one of the behaviors that we seek to study in realistic social networks in Section V: can we reduce contagion spread by increasing thresholds rather than by removing nodes from G (i.e., can we reduce spread by increasing thresholds to a lesser extent than to $\theta = d(v) + 1$)? When a node threshold is increased between two simulations (i.e., cases) of contagion, as done here for node v_4, we say that v_4 is a **blocking node**. That is, the node may inhibit contagion spreading (our goal is to have blocking nodes reduce σ, but some or all of them may be ineffective). The set of all blocking nodes in a simulation is the **blocking set** B.

The **threshold increase ratio** is a comparison between two simulations (cases), each relative to the same base case. The numerator is the sum of threshold increases *due to fear* over all k blocking nodes. The denominator is the sum of threshold increases *due to node removal* over all k blocking nodes. Here, the base case is case 1, with no threshold increases, and $k = 1$ because $B = \{v_4\}$. Case 2 has a threshold increase *due to fear* of $\Sigma^{fear} = \Sigma_{v_i \in B}(\Delta\theta_i) = 1$. From case 4, we know that $\Sigma^{removal} = \Sigma_{v_i \in B}(\Delta\theta_i) = 4$. So the threshold increase ratio is $\Sigma^{fear}/\Sigma^{removal} = 1/4 = 0.25$. This is a way to measure the **cost** of instilling fear, relative to removing nodes.

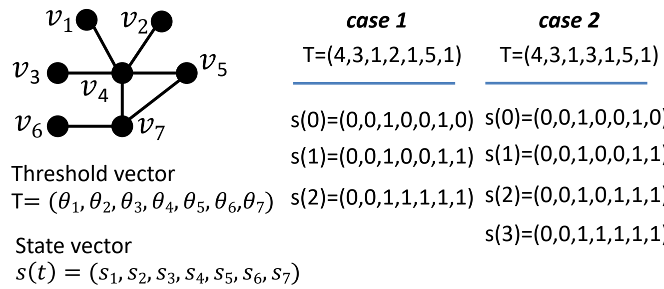

Fig. 1: Schematic of a GDS, showing a 7-node graph G, threshold vector T, and state vectors $s(t)$ as they update over time. There are two cases, each involving a threshold vector T where the only difference between them is that θ_4 increases from 2 to 3 for v_4. These cases are discussed in the text.

IV. Networks and Results

Networks evaluated in this work are listed in Table I. Both are human proximity (contact) networks. The human contact networks were generated using synthetic populations of cities, and the procedures in [2]. For this work, we make the assumption that issues that pertain to protest and fear involve only people between the ages of 15 to 70, inclusive, so that the networks are subgraphs of the full city social networks that include only those people in the specified age range.

TABLE I: Human proximity contact networks. The city-based human contact networks were made with the procedures in [2]. Each network is the giant component of the network, since we run dynamics on these networks.

Network	Type	Num. Nodes	Num. Edges	Ave. Deg.	Max. Deg.
Danville, VA	human contact	12961	44393	6.85	93
Newport News, VA	human contact	64425	418879	13.00	344

V. Simulations and Results

A. Simulation Process

Simulations are of the following scenario. A population wants to protest/revolt against a repressive regime. But protest has risk in that if it is unsuccessful, there will be adverse consequences, perhaps violence and/or death, for those individuals revolting. The regime seeks to thwart the protest by inducing fear of reprisal in a subset of its citizens, which is manifested through threshold increases $\Delta\theta$ for these fearful people.

A **simulation** is composed of 100 runs. A **run** is one contagion spread instance. The following applies to one run. All nodes are assigned a base threshold θ_{base}, and only the fearful nodes have their thresholds increased by $\Delta\theta$ to θ_{blk}. Fearful nodes—which are the blocking nodes—are identified by (i) ordering the nodes of a network by their degrees, from greatest to least (so that the first node has the greatest degree; break ties arbitrarily); (ii) specifying a number k ($\leq n$); and (iii) selecting the top k nodes that have the greatest degrees as the fearful (blocking) nodes. Given a base threshold θ_{base} and threshold increase $\Delta\theta$ for fearful (blocking) nodes, a node v_i's threshold is *never* more than $d(v_i) + 1$, as explained in Section III-C. For example, suppose $\theta_{base} = 1$, $\Delta\theta = 6$, and $d(v_i) = 3$. Then the actual $\Delta\theta$ applied to v_i is $\Delta\theta = d(v_i) + 1 - \theta_{base} = 3$. This ensures that the threshold increases are not spurious.

Each run has a set of seed nodes that are in state 1 at time $t = 0$. The seed nodes are selected using the Centola/Macy method [4] in that a node is chosen at random, and it and all of its distance-1 neighbors are made seed nodes. Because the number of seed nodes can vary depending on the degree of the chosen node, we set a minimum number $n_{s,min}$ of seed nodes per run, and this procedure is repeated for a run, as needed, until at least $n_{s,min}$ total nodes are seeds. Given the above description, the only variations among runs in a simulation are the specific number $n_s \geq n_{s,min}$ of selected seed nodes, the particular nodes that are the seeds, and the k selected blocking nodes. The k blocking nodes only change if one or more of the desired blocking nodes is also a seed node, as described next. Because we have 100 seed sets (100 runs in a simulation), we are testing the blocking method against 100 different seed sets, and thus are testing our blocking method against seeding "anywhere" in a network.

There are many features of a simulation that make the goal of blocking contagions more onerous. First, seed nodes are selected before the blocking nodes. If a desired blocking node is already a seed node, then this node is not chosen for blocking; a lesser ranked blocking node takes its place so that the blocking set size is k. Second, the number of seed nodes for a run can be significantly greater than the minimum $n_{s,min}$. Third, the diffusion process is deterministic. This drives contagion through the graph faster, making it more difficult to block. Fourth, we use small values for base threshold θ_{base}. This makes it easier for non-fearful nodes to become active. Fifth, each social network is the giant component of the graph: this prevents seed nodes being specified in a small component and eliminates the possibility that contagion cannot reach some disconnected nodes.

B. Simulation Results

Each subsection below addresses the effect of a parameter on the spread fraction σ (the fraction of nodes that join a protest movement). Each box in a boxplot represents 100 values of σ from the 100 runs of a simulation.

Effect of threshold increase $\Delta\theta$. We evaluate the Danville network with $\theta_{base} = 2$, and we choose the k nodes using $(\mathrm{p}, k) = (dg, 5000)$, where p means we are selecting blocking nodes based on greatest degrees (dg). Figure 2 shows the results. In each plot, as $\Delta\theta$ increases from 1 to 6, σ decreases.

Figure 2a shows that in Danville, if the 5000 people with the greatest degrees (i.e., those with the most connections to other people) out of a population of 12961 have their thresholds increased by three ($\Delta\theta = 3$), then seeding nominally 50 ($n_{s,min} = 50$) people anywhere in the graph will result in a zero median spread fraction ($\sigma = 0$). Moreover, this achieves the same result as *removing* those 5000 people from the graph, which is the case denoted by the symbol ∞ on the x-axis (more on this below). Additionally, since the base threshold $\theta_{base} = 2$ in these simulations, even the 5000 nodes have threshold $\theta_{blk} = \theta_{base} + \Delta\theta = 5 < d_{ave} = 6.85$. That is, the thresholds θ_{blk} are not excessively large: they are not even equal to the average degree in the network. Thus, these results indicate that inducing fear in a population, resulting in threshold increases that are relatively modest, can be effective in deterring protests. This same conclusion can be drawn from the other plot in Figure 2.

Effect of seeding. We continue with Figure 2, and note that the only condition that changes across the plots is $n_{s,min}$ increasing from 50 to 500. The trend holds for $n_{s,min} = 10$, 50, 100, and 500. As the minimum number of seeds $n_{s,min}$ increases, from 50 in Figure 2a to 500 in Figure 2b, the blocking nodes have a more difficult time thwarting the contagion spread, so σ increases across the two plots for a fixed value of $\Delta\theta$.

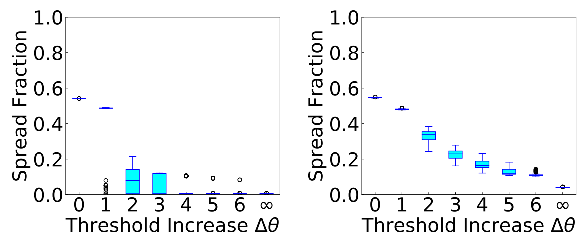

(a) (b)

Fig. 2: Fractional spread size σ as a function of threshold increase $\Delta\theta$ for the Danville network with $\theta_{base} = 2$, with blocking nodes $(\mathrm{p}, k) = (dg, 5000)$. The number of seed nodes per run increases in plots as (a) $n_{s,min} = 50$, and (b) $n_{s,min} = 500$. As $\Delta\theta$ increases, σ decreases and as the minimum number of seeds $n_{s,min}$ increases, σ increases. "Base" case is threshold increase $\Delta\theta = 0$. $\Delta\theta = \infty$ means that nodes are *removed* from the network.

Effectiveness of smaller threshold increases versus complete node removal. Figure 2 has on the x-axis $\Delta\theta = \infty$. This is the case that the $k = 5000$ blocking nodes are *removed* from the network (rather than having their thresholds increased by a smaller amount). For $n_{s,min} = 50$ in Figure 2a, $\Delta\theta = 6$ produces the same median spread fraction σ as removing those nodes. However, in Figure 2b for $n_{s,min} = 500$, removing blocking nodes results in less spread than does increasing the thresholds by six for the $k = 5000$ nodes with the greatest degrees. Node removal can be represented as each of the $k = 5000$ nodes v_i having their thresholds increased to $d(v_i) + 1$ so that they never transition to state 1. Then, we can compare the effectiveness of threshold increases due to fear, to node removal, by computing the ratio of the sum of units of threshold increase for a given $\Delta\theta$ (from fear) to the sum of units of threshold increase for node removal (the sums are over all blocking nodes). This is the **threshold increase ratio**. Figure 4a shows the results of these computations for the data in Figure 2. Specifically, this plot shows that the good performance of reducing σ in Figure 2 is achieved by total threshold increases that are only 0.38 fraction of the total threshold increase that would be required for node removal. The 0.38 value is for $\Delta\theta = 6$; even smaller ratios are required for smaller $\Delta\theta$. Figures 2 and 4a, taken together, show that instilling fear in a population can be just as effective at stymieing contagion spreading—at far less cost, if cost is measured by total threshold increases over k nodes—as removing people from society.

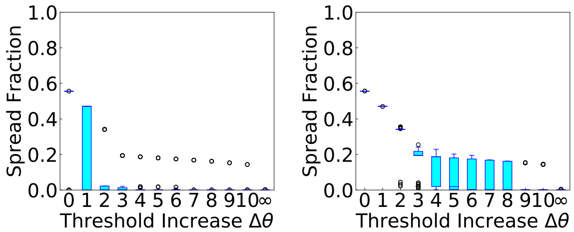

(a) (b)

Fig. 3: Fractional spread size σ as a function of threshold increase $\Delta\theta$ for the Newport News network. The base threshold $\theta_{base} = 2$, blocking nodes $(\mathrm{p}, k) = (dg, 30000)$, and number of seed nodes per run increases across two plots as (a) $n_{s,min} = 10$ and (b) $n_{s,min} = 100$, "Base" case is threshold increase $\Delta\theta = 0$. $\Delta\theta = \infty$ means that nodes are *removed* from the network. In (a), the median spread $\sigma = 0$ for $\Delta\theta \geq 2$. In (b), where the $n_{s,min}$ increases by $10\times$ to 100, the median spread $\sigma = 0$ for $\Delta\theta \geq 6$. In both cases, σ can be reduced to zero, as is the case when nodes are *removed* from the network. However, for the cases of threshold increases, there are individual runs among the 100 runs where the spread fraction $\sigma > 0$.

These results are not specific to one network. Data analogous to those in Figure 2 for Danville are shown for the Newport News network in Figure 3, for $k = 30000$ blocking nodes in this much larger graph. The same trends are observed: spread fraction decreases as $\Delta\theta$ increases; and the threshold increases, representing increasing fear, can be as effective as removing nodes (the latter represented by the ∞ value on the x-axis). Moreover, Figure 4b shows that threshold increases are effective compared to node removal.

VI. Summary

This work demonstrates that instilling fear in a subset of the people in a population (modeled by threshold increases in selected blocking nodes) can significantly diminish, and in many cases halt, a contagion of citizens joining a protest. In addition, this can be done by a repressive regime at a cost, as measured by total threshold increase of blocking nodes, that is small compared to taking more overt, drastic actions.

Acknowledgments: We thank the referees for their comments. We thank Michele Co and other system staff (University of Virginia) for compute cluster support. This work is partially supported by NSF Grants: ACI-1443054 (DIBBS), IIS-1633028 (BIG DATA), CMMI-1745207 (EAGER), OAC-1916805 (CINES), CCF-1918656 (Expeditions), CRISP 2.0-1832587 and IIS-1908530.

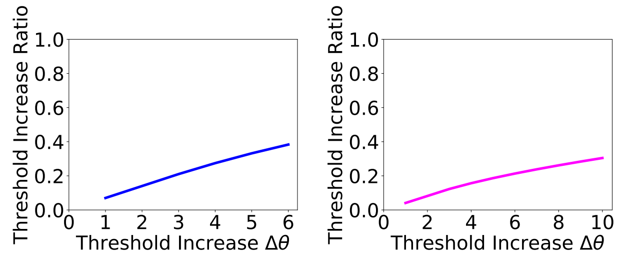

(a) (b)

Fig. 4: The plot in (a) is related to the data in Figure 2 for Danville and the plot in (b) is related to the data in Figure 3 for Newport News. Threshold increase ratio is the sum of threshold increases over the $k = 5000$ blocking nodes, in (a), as a result of fear, divided by the sum of threshold increases for *removal* of these same k nodes in the Danville network. Similarly, the threshold increase ratio in (b) is for the threshold increases for all $k = 30000$ blocking nodes in the Newport News network. The data show that for all $\Delta\theta$ values, the cumulative threshold increase over all k nodes (for fear) is at most 0.38 of the total threshold increase required for node removal. Hence, blocking by fear is a cost-effective way (cost in terms of threshold increases) to control a population.

References

[1] A. Adiga *et al.*, "Graphical dynamical systems and their applications to bio-social systems," *International Journal of Advances in Engineering Sciences and Applied Mathematics*, pp. 1–19, 2018.

[2] C. L. Barrett, R. J. Beckman, M. Khan *et al.*, "Generation and analysis of large synthetic social contact networks," in *WSC*, 2009, pp. 1003–1014.

[3] R. Boesche, "Fearing monarchs and merchants: Montesquieu's two theories of despotism," *The Western Political Quarterly*, vol. 4, pp. 741–761, 1990.

[4] D. Centola and M. Macy, "Complex contagions and the weakness of long ties," *American Journal of Sociology*, vol. 113, pp. 702–734, 2007.

[5] M. S.-Y. Chwe, "Structure and strategy in collective action," *American Journal of Sociology*, vol. 105, no. 1, pp. 128–156, 1999.

[6] C. R. Conrad, "Divergent incentives for dictators: Domestic institutions and (international promises not to) torture," *Journal of Conflict Resolution*, vol. 58, pp. 34–67, 2014.

[7] C. Davenport, "State repression and political order," *Annu. Rev. Polit. Sci.*, vol. 10, pp. 1–23, 2007.

[8] H. Fjelde, "Generals, dectators, and kings: Authoritarian regimes and civil conflicts, 1973–2004," *Conflict Management and Peace Science*, vol. 27, pp. 195–218, 2010.

[9] S. Gonzalez-Bailon *et al.*, "The dynamics of protest recruitment through an online network," *Scientific Reports*, pp. 1–7, 2011.

[10] M. Granovetter, "Threshold models of collective behavior," *American Journal of Sociology*, vol. 83, no. 6, pp. 1420–1443, 1978.

[11] D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the Spread of Influence Through a Social Network," in *KDD*, 2003, pp. 137–146.

[12] H. Mortveit and C. Reidys, *An Introduction to Sequential Dynamical Systems*. New York, NY: Springer, 2007.

[13] D. Romero, B. Meeder, and J. Kleinberg, "Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on twitter," in *WWW*, 2011.

[14] T. Schelling, *Micromotives and Macrobehavior*, 1978.

[15] D. Siegel, "When does repression work?: Collective action and social networks," *Journal of Politics*, vol. 73, pp. 993–1010, 2011.

[16] D. A. Siegel, "Non-disruptive tactics of suppression are superior in countering terrorism, insurgency, and financial panics," *PLoS ONE*, vol. 6, pp. e18 545–1–e18 545–6, 2011.

[17] J. Ugander, L. Backstrom, C. Marlow, and J. Kleinberg, "Structural Diversity in Social Contagion," *PNAS*, vol. 109, pp. 5962–5966, 2012.

[18] D. Watts, "A simple model of global cascades on random networks," *PNAS*, vol. 99, pp. 5766–5771, 2002.

[19] L. E. Young, "The psychology of state repression: Fear and dissent decisions in Zimbabwe," *Am. Pol. Sci. Rev.*, vol. 113, pp. 140–155, 2019.