Proc. NATO RTO HFM-201 Specialists Meeting,
Tallinn, Estonia, April 16-18, 2012, pp. 32-1 -- 32-15

On the Influence of Committed Minorities on Social Consensus

Boleslaw Szymanski[1,2], Sameet Sreenivasan[1,2,3], Jierui Xie[1,2], Weituo Zhang[2,4], Jeff Emenheiser[2,3], Matt Kirby[2,3], Chjan Lim[2,4] and G. Korniss[2,3]

1. Dept. of Computer Science, RPI, Troy, NY, USA
2. Social Cognitive Networks Academic Research Center, RPI, Troy, NY, USA
3. Dept. of Physics, RPI, Troy, NY, USA
4. Dept. of Mathematics, RPI, Troy, NY, USA

szymansk@cs.rpi.edu, {sreens, xiej2, zhangw10, emenhj, kirbym, limc, korniss} @rpi.edu

ABSTRACT

Human behavior is profoundly affected by the influenceability of individuals and the social networks that link them together. In the sociological context, spread of ideas, ideologies and innovations is often studied to understand how individuals adopt new states in behavior, opinion, ideology or consumption through the influence of their neighbors. In this paper, we study the evolution of opinions and the dynamics of its spread. We use the binary agreement model starting from an initial state where all individuals adopt a given opinion B, except for a fraction p<1 of the total number of individuals who are committed to opinion A. In a generalization of this model, we consider also the initial state in which some of the individuals holding opinion B are also committed to it. Committed individuals are defined as those who are immune to influence but they can influence others to alter their opinion through the usual prescribed rules for opinion change. The question that we specifically ask is: how does the consensus time vary with the size of the committed fraction? More generally, our work addresses the conditions under which an inflexible set of minority opinion holders can win over the rest of the population. We show that the prevailing majority opinion in a population can be rapidly reversed by a small fraction p of randomly distributed committed agents who are immune to influence. Specifically, we show that when the committed fraction grows beyond a critical value p_c*=9.79%, there is a dramatic decrease in the time,* T_c*, taken for the entire population to adopt the committed opinion. Below this value, the consensus time is proportional to the exponential function of the network size, while above this value this time is proportional to the logarithm of the network size. This has enormous impact on stability/instability of the society opinions. We also discuss conditions under which the committed minority can rapidly reverse the influenceable majority even if the latter is supported by a small fraction of individuals committed to their opinion. The results are relevant in understanding and influencing the social perceptions of international missions operating in various countries and attitudes to the mission goals.*

1.0 INTRODUCTION

The propagation of social influence through the social networks that link individuals together has a profound effect on human behavior. Prior to the proliferation of online social networking, offline or interpersonal social networks have been known to play a major role in determining how societies move towards consensus in the adoption of ideologies, traditions and attitudes [1] [2]. As a consequence, the dynamics of social influence has been heavily studied in sociological, physics and computer science literature [3] [4] [5] [6]. In the sociological context, work on diffusion of innovations has emphasized how individuals adopt new states in behavior, opinion or consumption through the influence of their neighbors. Commonly used models for this process include the threshold model [7] and the Bass model [8]. A key feature in both these models is that once an individual adopts the new state, his state remains unchanged at all subsequent times. While appropriate for modeling the diffusion of innovation where investment in a new idea comes at a cost, these models are less suited to studying the dynamics of competing opinions where switching one's state has little overhead.

Here we address the latter case. From among the vast repertoire of models in statistical physics and mathematical sociology, we focus on one which is a 2-opinion variant [9] [10] of the Naming Game (NG) [11] [12] [13] [14] and that we refer to as the *binary agreement model*. The evolution of the system in this model takes place through the usual NG dynamics, wherein at each simulation time step, a randomly chosen speaker voices a random opinion from his list to a randomly chosen neighbor, designated the listener. If the listener has the spoken opinion in his list, both speaker and listener retain only that opinion; else the listener adds the spoken opinion to his list. The order of selecting speakers and listeners is known to influence the dynamics, and we stick to choosing the speaker first, followed by the listener.

It serves to point out that an important difference between the binary agreement model and the predominantly used opinion dynamics models [4] is that an agent is allowed to possess both opinions simultaneously in the former, and this significantly alters the time required to attain consensus starting from uniform initial conditions. Numerical studies in [10] have shown that for the binary agreement model on a complete graph, starting from an initial condition where each agent randomly adopts one of the two opinions with equal probability, the system reaches consensus in time $T_c \sim \ln N$ (in contrast, for example, to $T_c \sim N$ for the voter model [4]). Here N is the number of nodes in the network, and unit time consists of N speaker-listener interactions. The binary agreement model is well suited to understanding how opinions, perceptions or behaviors of individuals are altered through social interactions specifically in situations where the cost associated with changing one's opinion is low [15], or where changes in state are not deliberate or calculated, but unconscious [16]. Furthermore, by its very definition, the binary agreement model may be applicable to situations where agents while trying to influence others, simultaneously have a desire to reach global consensus [17].

Another merit of the binary agreement dynamics in modeling social opinion change seems worth mentioning. Two state epidemic-like models of social "contagion" (examples in [18]) suffer from the drawback that the rules governing the conversion of a node from a given state to the other are not symmetric for the two states. In contrast, in the binary agreement model, both singular opinion states are treated symmetrically in their susceptibility to change.

Here we study the evolution of opinions in the binary agreement model in a scenario where fractions of the network can be *committed* [19] [20] to one of the two opinions. Committed nodes are defined as those that can influence other nodes to alter their state through the usual prescribed rules, but which themselves are stubbornly devoted to one opinion and thus, immune to influence.

The effect of having un-influencable agents has been considered to some extent in prior studies. Biswas et al. [21] considered for two-state opinion dynamics models in one dimension, the case where some individuals are "rigid" in both segments of the population, and studied the time evolution of the magnetization and the fraction of domain walls in the system. Mobilia et al. [22] considered the case of the voter model with some fraction of spins representing "zealots" who never change their state, and studied the magnetization distribution of the system on the complete graph, and in one and two dimensions. Our study differs from these not only in the particular model of opinion dynamics considered, but also in its explicit consideration of different network topologies and of finite sized networks, specifically in its derivation of how consensus times scale with network size for the case of the complete graph. Furthermore, the above mentioned studies do not explicitly consider the initial state that we care about - one where the entire minority set is un-influencable. A notable exception to the latter is the study by Galam and Jacobs [19] in which the authors considered the case of "inflexibles" in a two state model of opinion dynamics with opinion updates obeying a majority rule. While that study provides several useful insights and is certainly the seminal quantitative attempt at understanding the effect of committed minorities, its analysis is restricted to the mean-field case, and has no explicit consideration of consensus times for finite systems.

Specifically, we study two different situations that involve the effect of committed agents on the eventual steady state opinion in the network. In Sec. 2 we study the situation where there exists only one kind of committed set, i.e. all committed nodes are devoted to the same opinion. The question that we specifically address in this section is: how does the consensus time vary with the size of the committed fraction p? More generally, our work addresses the conditions under which an inflexible set of minority opinion holders can win over the rest of the population. In Sec. 3 we extend our study to the case where competing committed groups are present in the population. Here the relevant question is similar to that in Sec. 2 but more formally posed as follows. Suppose the majority of individuals on a social network subscribe to a particular opinion on a given issue, and additionally some fraction of this majority consists of unshakeable in their commitment to the opinion. Then, what should be the minimal fractional size of a competing committed group in order to effect a fast reversal in the majority opinion?

2.0 EFFECT OF A SINGLE GROUP OF COMMITTED AGENTS IN COMPLETE GRAPHS

We start by considering the case where a single group of committed agents is present in the population, and the connectivity of nodes within the population is defined by a complete graph. Additionally we are interested here in the situation where initially all uncommitted nodes adopt opinion B while the committed nodes are always committed to opinion A.

2.1 Infinite Network Size Limit

We start along similar lines as [19] by considering the case where the social network connecting agents is a complete graph with the size of the network $N\rightarrow\infty$. We designate the densities of uncommitted nodes in states A, B as n_A, n_B respectively. Consequently, the density of nodes in the mixed state AB is $n_{AB}=1-p-n_A-n_B$, where p is the fraction of the total number of nodes that are committed. Neglecting correlations between nodes, and fluctuations, one can write the following rate equations for the evolution of densities:

$$dn_A/dt = -n_A n_B + n_{AB}^{\;2} + n_{AB} n_A + 1.5\, p n_{AB}$$

(1)

$$dn_B/dt = -n_A n_B + n_{AB}^{\;2} + n_{AB} n_B - p n_B$$

The terms in these equations are obtained by considering all interactions which increase (decrease) the density of agents in a particular state and computing the probability of that interaction occurring. Table 1 lists all possible interactions. As an example, the probability of the interaction listed in row eight is equal to the probability that a node in state AB is chosen as speaker and a node in state B is chosen as listener ($n_{AB}\, n_A$) times the probability that the speaker voices opinion A (1/2).

The fixed-point and stability analyses [23] of these *mean-field* equations show that for any value of p, the consensus state in the committed opinion ($n_A=1-p$, $n_B=0$) is a stable fixed point of the mean-field dynamics.

However, below $p=p_c=\frac{5}{2}-\frac{3}{2}\left(\sqrt[3]{5+\sqrt{24}}-1\right)^2-\frac{3}{2}\left(\sqrt[3]{5-\sqrt{24}}-1\right)^2\approx 0.09789$, two additional fixed points appear: one of these is an unstable fixed point (saddle point), whereas the second is stable and represents an *active* steady state where n_A, n_B and n_{AB} are all non-zero (except in the trivial case where p=0). Fig. 1 shows (asterisks) the steady state density of nodes in state *B* obtained by numerically integrating the mean-field equations at different values of the committed fraction p and with initial condition $n_A=0$, $n_B=1-p$. As p is increased, the stable density of *B* nodes n_B abruptly jumps from ≈0.6504 to zero at the critical committed fraction p_c. A similar abrupt jump also occurs for the stable density of *A* nodes from a value very close to zero below p_c, to a value of 1, indicating consensus in the *A* state (not shown). In the study of phase transitions, an "order parameter" is a suitable quantity changing (either continuously or discontinuously) from zero to a non-zero value at the critical point. Following this convention, we use n_B - the density of uncommitted nodes in state *B* - as the order parameter appropriate for our case, characterizing the transition from an active steady state to the absorbing consensus state.

Table 1: Interactions in Naming Game

Before interaction	After interaction
$A \xrightarrow{A} A$	A - A
$A \xrightarrow{A} B$	A - AB
$A \xrightarrow{A} AB$	A - A
$B \xrightarrow{B} A$	B - AB
$B \xrightarrow{B} B$	B - B
$B \xrightarrow{B} AB$	B - B
$AB \xrightarrow{A} A$	A - A
$AB \xrightarrow{A} B$	AB - AB
$AB \xrightarrow{A} AB$	A - A
$AB \xrightarrow{B} A$	AB - AB
$AB \xrightarrow{B} B$	B - B
$AB \xrightarrow{B} AB$	B - B

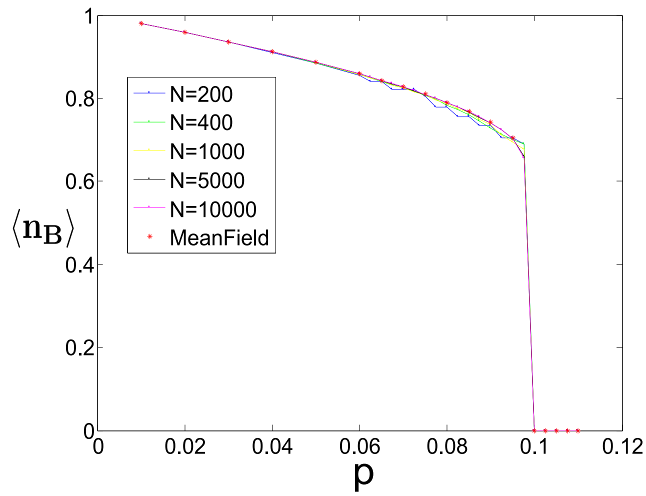

Figure 1. Steady-state density of nodes in state B as a function of the committed fraction *p*. Comparison between mean-field and simulation results.

2.2 Simulation Results for Finite Complete Graphs

In practice, for a complete graph of any finite size, consensus is always reached. However, we can still probe how the system evolves, conditioned on the system not having reached consensus. Fig. 1 shows the results of simulating the binary agreement model on a complete graph for different system sizes (solid lines). For $p<p_c$, in each realization of agreement dynamics, neglecting the initial transient, the density of nodes in state *B*, n_B, fluctuates around a non-zero steady state value, until a large fluctuation causes the system to escape from this active steady state to the consensus state. Fig. 1 shows these steady state values of n_B conditioned on survival, for several values of p. As expected, agreement of simulation results with the mean-field curve improves with increasing system size, since Eq. 1 represents the true evolution of the system in the asymptotic large network-size limit. Accordingly the critical value of the committed fraction obtained from the mean-field equations is designated as $p_c(\infty)$, although, for brevity, we refer to it simply as p_c throughout this paper.

The existence of the transition as p is varied and when the initial condition for densities is ($n_A=0$, $n_B=1-p$) can be further understood by observing the motion of the fixed points in phase space. As p is varied from 0 to p_c, the active steady state moves downward and right while the saddle point moves upwards and left. At the critical value p_c the two meet and the only remaining stable fixed point is the consensus fixed point. A similar observation was made in the model studied in [19]. The fact that as the stable fixed point and the saddle point approach each other, the value of n_B converges to 0.65 rather than smoothly approaching zero. Thus the transition is first-order in nature. Fig. 2 shows representative trajectories obtained by integrating the mean-field equations for the cases where p=0.05 ($<p_c$) and p=0.1($>p_c$).

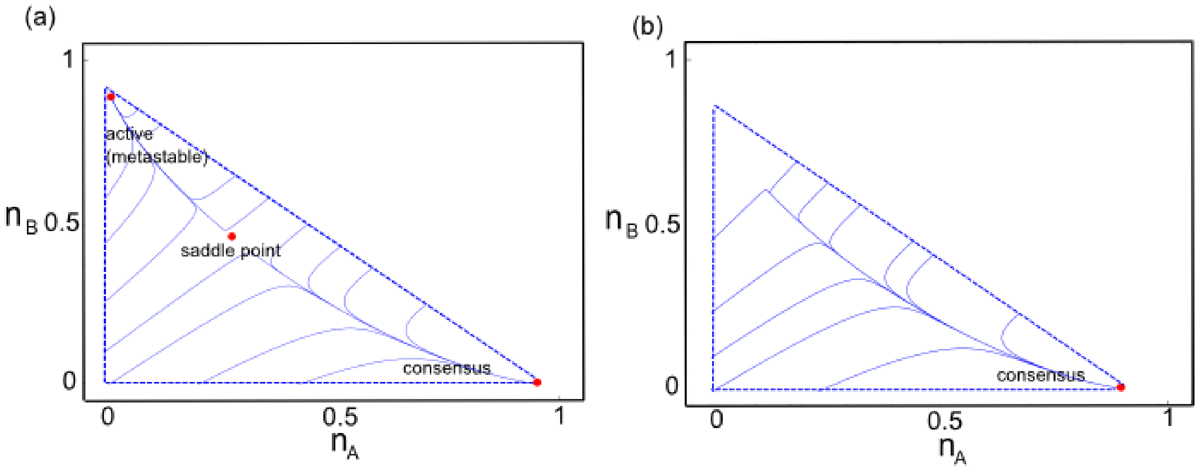

Figure 2. Phase plane picture of the dynamics. (a) The situation when $p < p_c$. Two stable states exists, one of which is the all A consensus state and the other is a mixed state with non-zero densities of all three classes of opinion holders. These stable states are separated by an unstable state or saddle point (b) Above p_c, the mixed stable state and the saddle point vanish, leaving the consensus state as the only stable state remaining.

We know from the mean field equations that in the asymptotic limit, and below a critical fraction of committed agents, there exists a stable fixed point. For finite stochastic systems, escape from this fixed point is always possible, and therefore it is termed *metastable*. Even though consensus is always reached for finite N, limits on computation time prohibit the investigation of the consensus time, T_c, for values of p below or very close to p_c. This can be overcome analytically using a quasi-stationary approximation [24]. This approximation assumes that after a short transient, the occupation probability (of different states in the configuration space of the system), conditioned on survival of allowed states excluding the consensus state, is stationary. For brevity we skip the details here which are given in [23]. Following the above method, we obtain the QS distribution, and consequently the mean consensus times T_c, for different values of committed fraction p and system size N. Fig. 3 (a) shows how the consensus time grows as p is decreased beyond the asymptotic critical point p_c for finite N. For $p<p_c$, the growth of T_c is exponential in N (Fig. 3 (b)), consistent with what is known regarding escape times from metastable states. For $p>p_c$, the QS approximation does not reliably provide information on mean consensus times, since consensus times themselves are small and comparable to transient times required to establish a QS state. However, simulation results show that above p_c the scaling of mean consensus time with N is logarithmic (Fig. 3 (c)). The precise dependence of consensus times on p can also be obtained for $p<p_c$ by considering the rate of exponential growth of T_c with N. In other words, assuming T_c~ exp($\alpha(p)N$), we can obtain $\alpha(p)$ as a function of p. Thus we find that below p_c:

$$T_c(p < p_c) \sim \exp((p_c - p)^{\nu} N) \qquad (2)$$

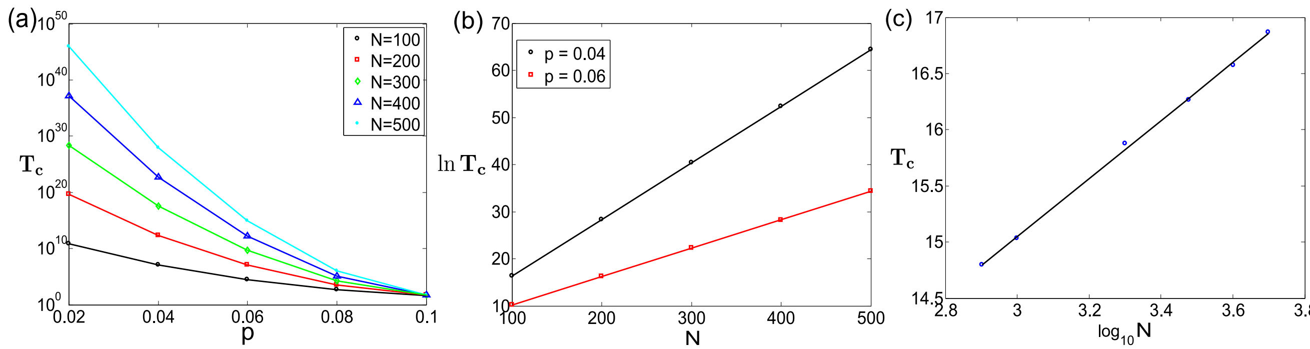

Figure 3. Consensus time as a function of committed fraction size p. (a) Shows the dramatic increase of consensus times below p_c = *0.0989* for different network sizes. Consensus times here are obtained using the quasistationary approximation. (b) The exponential scaling of consensus times with network size for two values of $p < p_c$. (c) Simulation results showing the logarithmic scaling of consensus times with network size for $p > p_c$.

2.3 Recursive Approach for Obtaining Consensus Time Distributions

Suppose the numbers of uncommitted nodes in states A, B and AB are denoted by N_A, N_B and N_{AB} respectively, the number of committed nodes is denoted by $N_p=pN$. We establish a random walk model over the space containing all possible macrostates (N_A, N_B). The transition matrix of this random walk is Q, each element of which, $Q(N_A', N_B' \mid N_A, N_B)$, is the probability for the system jumping from the current macrostate (N_A, N_B) to the macrostate (N_A', N_B') in one time step. All the elements in Q are zero except for the following cases:

Table 2: Non-zero elements in the transition matrix Q

Old State	New State	Events	Probability
(N_A, N_B)	(N_A-1,N_B)	A→AB	$N_A(N_B+N_{AB}/2)/[N(N-1)]$
	(N_A+1, N_B)	AB→A	$N_{AB}(N_A+N_p)*2/3/[N(N-1)]$
	(N_A, N_B-1)	B→AB	$N_B(N_A+N_p+N_{AB}/2)/[N(N-1)]$
	(N_A, N_B+1)	AB→B	$N_BN_{AB}*3/2/[N(N-1)]$
	(N_A+2, N_B)	AB-AB→A-A	$N_{AB}(N_{AB}-1)/2/[N(N-1)]$
	(N_A, N_B+2)	AB-AB→B-B	$N_{AB}(N_{AB}-1)/2/[N(N-1)]$
	(N_A, N_B)	no change	1-sum of the probabilities above

The analysis here refers to Table 2, and the detailed calculations are in [25]. Let $P(N_A, N_B, T)$ be the probability for the dynamics starting from the state (N_A, N_B) to achieve consensus at time T. Then we have the following recursive relationship through the first step analysis:

$$P(N_A, N_B, T+1)=Q(N_A+1, N_B \mid N_A, N_B)\, P(N_A+1, N_B, T)+Q(N_A-1, N_B \mid N_A, n_B)\, P(N_A-1, N_B, T)$$
$$+Q(N_A, N_B+1 \mid N_A, N_B)\, P(N_A, N_B+1, T)+ Q(N_A, N_B-1 \mid N_A, N_B)\, P(N_A, N_B-1, T)$$
$$+Q(N_A+2, N_B \mid N_A, N_B)\, P(N_A+2, N_B, T)+ Q(N_A, N_B+2 \mid N_A, N_B)\, P(N_A, N_B+2, T)$$
$$+Q(N_A, N_B \mid N_A, N_B)\, P(N_A, N_B, T)$$

where $Q(\cdot \mid \cdot)$'s are elements of the transition matrix Q. Representing probabilities corresponding to a given time T and the different possible (N_A, N_B) states into a vector $P(T)$ we get :

$$P(T+1) = Q * P(T) \tag{3}$$

with initial conditions: $P(N\text{-}N_p,0,0)=1$. The sequence of vectors corresponding to different T generated by Eq.3 can be thought of as a vector-valued function of T. Along a given coordinate direction, this sequence is a real-valued function of T which represents the consensus time distribution starting from the macrostate (N_A,N_B) corresponding to this coordinate direction: Prob$(N_A,N_B,T_c=T)=P(N_A,N_B,T)$. Note that we calculate the consensus time distributions starting from all (N_A,N_B) at the same time. Figure 4 shows the consensus time distribution starting from macrostate $(0,N\text{-}N_p)$ obtained using Eq. 3 compared with simulation results.

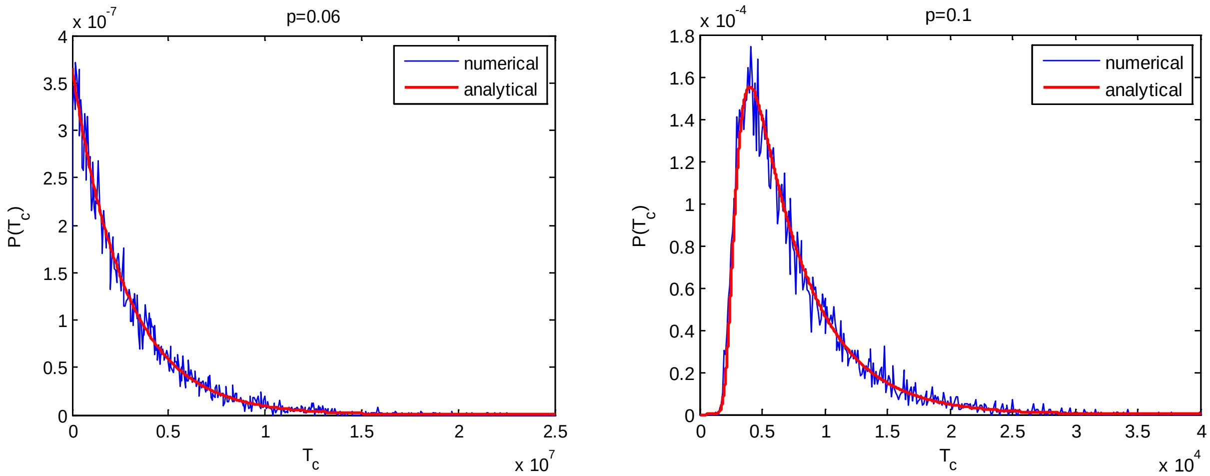

Figure 4. Compare the consensus time distributions from the analysis via the recursive approach (red) with those from statistics of numerical simulations (blue) . $p(T_c)$ stands for the probability of achieving consensus at T_c. The left plot is for $p=0.06<p_c$, the right plot is for $p=0.1>p_c$.

A direct consequence of Eq. 3 is that, $P(T)$ has an eigenvector representation:

$$P(T) = \Sigma_k C_k \exp(\log(\lambda_k)\, T) \quad (k=1,\ldots, M)$$

where k is the index of summation, M is the total number of macrostates, C_k's are constants, λ_k's are the eigenvalues of Q in decreasing order. As $\lambda_k<1$, the first term $C_1 \exp(\log(\lambda_1)T)$ dominates when T is big enough. Therefore the consensus time distribution Prob$(N_A,N_B, T_c=T)$ always has an exponentially decaying tail with a decay constant -log(λ_1), where λ_1 is the largest eigenvalue of Q.

Using Eq. 3, we check the asymptotic behavior of the consensus time distribution as the network size grows. Figure 5 shows the normalized consensus time distribution for different p and N. When $p<p_c$, the consensus time distribution tends to exponential distribution. When $p>p_c$, the central region of the consensus-time distribution tends to a Gaussian distribution, but always exhibits an exponentially decaying tail according to the expression shown in the previous paragraph. We also observe the concentration of the consensus time distribution for $p>p_c$ as N grows which validates the mean field assumption in calculating the consensus time. Rigorous analysis of the concentration is in [26].

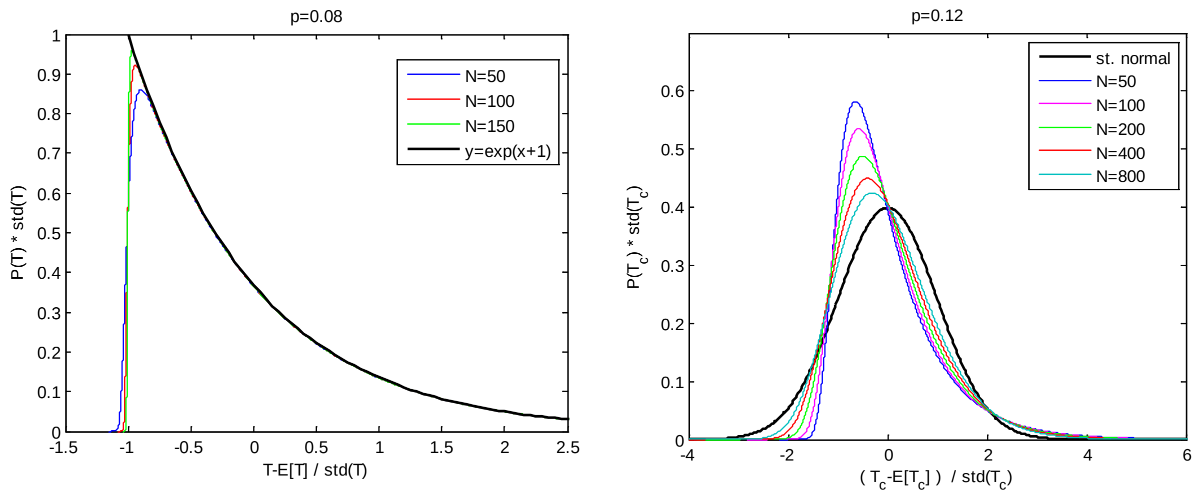

Figure 5. Asymptotic behavior of consensus time distribution as network size grows. When $p=0.08<p_c$, the consensus time distribution tends to exponential distribution. When $p=0.12>p_c$, the consensus time distribution tends to Gaussian distribution. The *E* and std denote the mean and the standard deviation repectively.

3.0 EFFECT OF TWO COMPETING GROUPS OF COMMITTED AGENTS IN COMPLETE GRAPHS

As earlier we begin with the case of the mean-field or infinite network size limit.

3.1 Infinite Network Size Limit

We designate the densities of uncommitted agents in the states *A*, *B* and *AB* by n_A, n_B, and n_{AB}. We also designate the fraction of nodes committed to state *A*, *B* by p_A, p_B respectively. These quantities naturally obey the condition: $n_A+n_B+n_{AB}+p_A+p_B=1$. In the asymptotic limit of network size, and neglecting fluctuations and correlations, the system can be described by the following mean-field equations [27], for given values of the parameters p_A and p_B:

$$dn_A/dt = -n_A n_B + n_{AB}^2 + n_{AB} n_A + 1.5 p n_{AB} - p_B n_A$$

$$dn_B/dt = -n_A n_B + n_{AB}^2 + n_{AB} n_B + 1.5 p n_{AB} - p_A n_B \qquad (4)$$

The evolution of n_{AB} follows from the constraint on densities defined above.

We now characterize the behavior of the system governed by Eqs. (4) for p_A, $p_B>0$, we systematically explore the parameter space (p_A, p_B) by dividing it into a grid with a resolution of 0.000125 along each dimension. We then numerically integrate Eqs. (4) for each (p_A, p_B) pair on this grid, assuming two distinct initial conditions, $n_A=1-p_A-p_B$, $n_B=n_{AB}=0$ and $n_B=1-p_A-p_B$, $n_A=n_{AB}=0$, representing diagonally opposite extremes in phase space. The results of this procedure reveal the picture shown in Fig. 6 in different regions of parameter space. As is obvious, with non-zero values for both p_A, p_B, consensus on a single opinion can never be reached, and therefore all fixed points (steady-states) are non-absorbing. With (p_A, p_B) values within the region denoted as I which we refer to as the "beak" (borrowing terminology used in [28]) the phase space contains two stable fixed points, separated by a saddle point, while outside the beak, in region II, only a single stable fixed point exists in phase space. In region I, one fixed point corresponds to a state where opinion *A* is the majority opinion (*A*-dominant) while the other fixed point corresponds to a state where opinion *B* constitutes the majority opinion (*B*-dominant). Figure 6 also shows representative trajectories and fixed points in phase space, in different regions of parameter space.

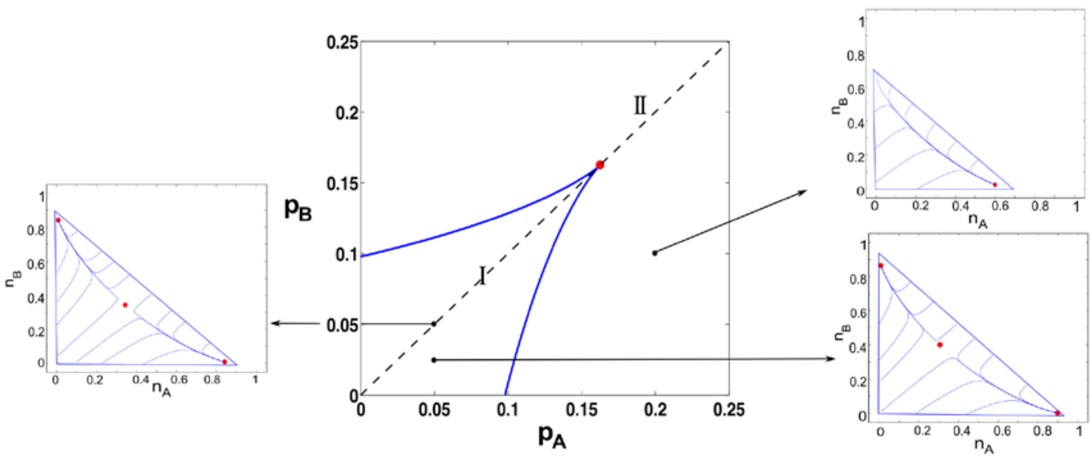

Figure 6. Representative trajectories (smaller plots) in different regions of parameter space (central plot). In Region I, the phase-plane has two co-existing stable steady-states separated by a saddle point. In Region II, only a single stable steady-state exists. The two regions are separated by bifurcation lines (blue) which meet tangentially at a cusp point (red).

In order to study the nature of the transitions that occur when we cross the boundaries of the beak, we parameterize the system by denoting $p_B=cp_A$ where c is a real number. Then, we systematically analyze the transitions occurring in two cases: (i) c=1 and (ii) $c\neq1$. It can be shown that along the diagonal line c=1 the system undergoes a cusp bifurcation at $p_A=p_B=0.1623$. Henceforth, we denote the value of p_A and p_B at the cusp as p_c. As is well known, at the cusp bifurcation two branches of a saddle-node (or fold) bifurcation meet tangentially [29]. These two bifurcation curves form the boundary of the beak shown in Fig. 6. The cusp bifurcation point is analogous to a second-order (or continuous) critical point seen in equilibrium systems, while bifurcation curves are analogous to spinodal transition lines.

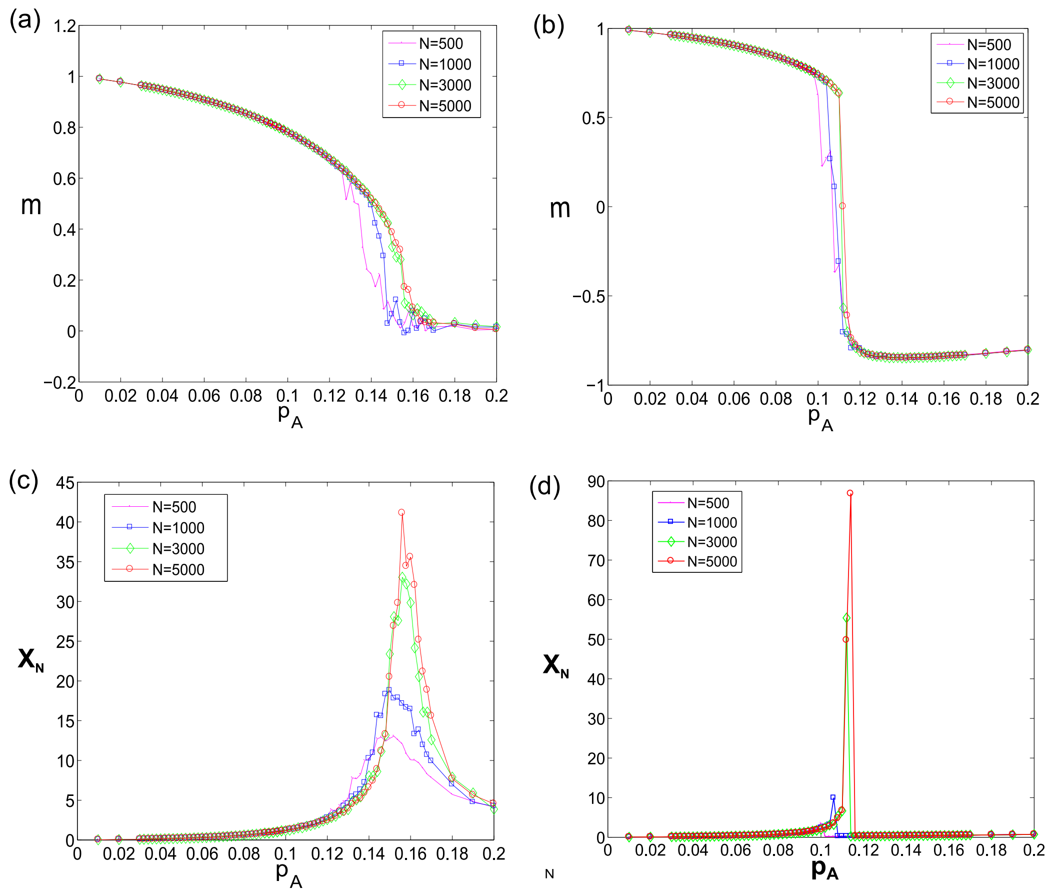

Figure 7. Behavior of order parameters as a function of linear trajectories. (a) ,(b). Magnetization as a function of fraction p_A of committed nodes in state A for a linear trajectories with slope $c = 1$ (diagonal) and $c = 0.5$ (off-diagonal) respectively. (c), (d). Scaled variance of the magnetization as a function of p_A along trajectories with $c = 1$ and $c = 0.5$ respectively.

3.2 Simulation Results for Finite Complete Graphs

Next, we show results for the stochastic evolution of opinions on finite-sized complete graphs through simulations [27]. Here, we systematically vary *c* from 1 to 0 to obtain the right bifurcation curve, and therefore by virtue of the *A*-*B* symmetry in the system, also obtain the left bifurcation curve. In particular for a given value of *c* we obtain the transition point by varying p_A (with $p_B = cp_A$) and measuring the quantity:

$$m = (n_B - n_A)/(1 - p_A - p_B) \tag{5}$$

which we utilize as an order parameter. The above order parameter is analogous to the "magnetization" in a spin system as it captures the degree of dominance of opinion *B* over opinion *A* and is conventionally used to characterize the nature of phase transitions exhibited by such a system. The behavior of this order parameter for trajectories crossing the cusp point and the bifurcation curves are shown in Fig. 7 (a), (b), respectively.

The behavior of the transitions in magnetization demonstrate quantitatively that the cusp bifurcation point and the saddle-node bifurcation curves observed for mean-field system are manifested in finite networks as a second-order transition point and two first-order transition (spinodal) lines respectively.

The fluctuations of the quantity *m* can be used to identify a transition point, particularly for the case of the second-order transition. In particular, in formal analogy with methods employed in the study of equilibrium spin systems, the scaled variance:

$$X_N = N\left\langle (|m| - \langle |m| \rangle)^2 \right\rangle \tag{6}$$

serves as an excellent estimate for the second-order transition point p_c for a finite network. As shown in Fig. 7 (c), X_N peaks at a particular value of p_A, with the size of the peak growing with N (and expected to diverge as $N\rightarrow\infty$). In the case of the spinodal transition, one studies fluctuations of m ($X_N=N<(m-<m>)^2>$) restricted to the metastable state [30] [31] until the spinodal point at which the metastable state disappears, and fluctuations of *m* in the unique stable state beyond the spinodal point (Fig. 7 (d)).

Figure 8 shows the bifurcation (spinodal) lines obtained via simulations of finite complete graphs, and demonstrating that its agreement with the mean-field curves improves as *N* grows.

In the region within the beak, the switching time between the co-existing steady-states represents the longest time-scale of relevance in the system. The switching time is defined as the time the system takes to escape to a distinct co-existing steady-state, after having been trapped in one of the steady-states. In stochastic systems exhibiting multistability or metastability, it is well known that switching times increase exponentially with *N* for large *N* (the weak-noise limit) [28] [32] [33] [34]. Furthermore, the exponential growth rate of the switching time in such cases can be determined using the *eikonal* approximation [28] [35]. The basic idea in the approximation involves (i) assuming an *eikonal* form for the probability of occupying a state far from the steady-state and (ii) smoothness of transition probabilities in the master equation of the system. This allows the interpretation of fluctuational trajectories as paths conforming to an auxiliary Hamilton-Jacobi system. This in turn enables us to calculate the probability of escape allowing an *optimal fluctuational path* that takes the system from the vicinity of the steady-state to the vicinity of the saddle point of the deterministic system. The switching time is simply the inverse of the probability of escape along this optimal fluctuational path. For the sake of brevity we skip the details of this procedure here. Using this approach we find that for the symmetric case, $p_A=p_B=p<p_c$, the exponential growth rate of the switching time $s\sim(p_c-p)^\nu$ with $\nu\approx1.3$ (Fig. 9). Thus, along the portion of the diagonal within the beak:

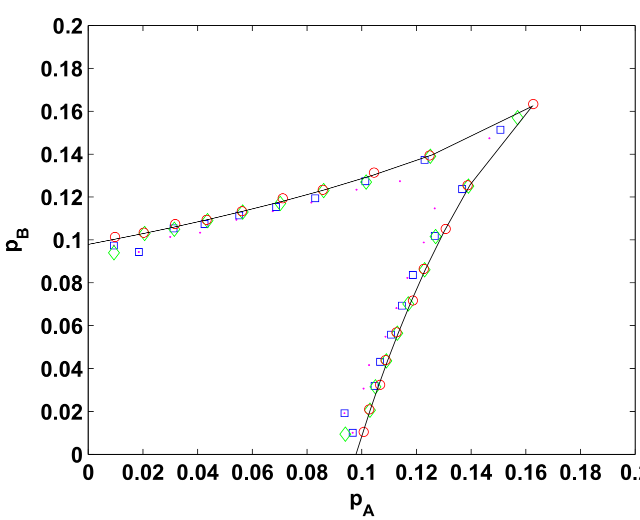

Figure 8. Bifurcation curves using the mean-field equations (solid line) and simulations (symbols). The network sizes used are *N*=500 (dots), 1000 (squares), 3000 (diamonds) and 5000 (circles).

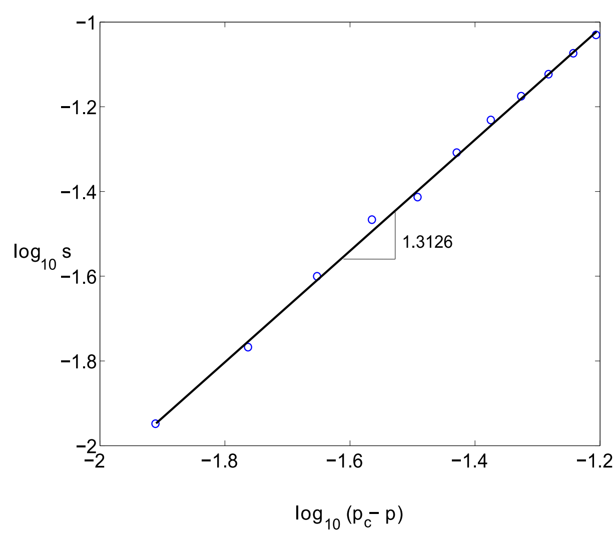

Figure 9. Scaling of the exponential rate constant in $T_c \sim \exp(sN)$ as a function of distance from the critical point p_c along the diagonal trajectory.

$$T_{switching} \sim \exp\left[(p_c - p)^{\nu} N\right] \quad (7)$$

Outside the beak, the time to get arbitrarily close to the sole steady-state value grows logarithmically with N (not shown).

The results presented in Sec. 3 show that there exists a transition in the time needed by a committed minority to influence the entire population to adopt its opinion, even in the presence of a committed opposition (i.e. in the case where both p_A, $p_B>0$), as long as p_A, $p_B<p_c$ (the case $p_A>0$, $p_B=0$ was considered in Sec 2.). For example, assume that initially all the uncommitted nodes adopt opinion B, and that $p_A=p_B<p_c$. Then, the steady-state that the system reaches in $\ln N$ time is the one in which the majority of nodes hold opinion B. Despite the fact that there exist committed agents in state A continuously proselytizing their state, it takes an exponentially long time before a large (spontaneous) fluctuation switches the system to the A-dominant steady-state. For identical initial conditions, the picture is qualitatively the same if we increase p_A keeping p_B fixed, as long as (p_A, p_B) lies within the beak. However, when (p_A, p_B) lies on the bifurcation curve or beyond, the B-dominant steady-state vanishes, and with the same initial conditions - where B is the initial majority - it takes the system only $\ln N$ time to reach the A-dominant state (the only existing steady-state). Thus, for every value of an existing committed fraction p_B ($<p_c$) of B nodes, there exists a corresponding critical fraction of A nodes beyond which it is guaranteed that the system will reach an A dominant state in $\ln N$ time, irrespective of the initial conditions. However, for any trajectory in parameter space in a region where either p_A or p_B is (or both are) greater than p_c, no abrupt changes in dominance or consensus times are observed. Instead, the dominance of A or B at the single fixed point smoothly varies as the associated committed fractions are varied. Moreover, the system always reaches this single fixed point in $\ln N$ time.

In conclusion, we have demonstrated how opinions within a social network evolve in when groups of committed nodes that are immune to influence exist within the network. Our results may be applicable to a wide variety of situations including social revolutions, debates on global warming, and the competition between different industrial standards.

Acknowledgments. This work was supported in part by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053, by the Army Research Office Grant No. W911NF-09-1-0254, and by the Office of Naval Research Grant No. N00014-09-1-0607. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government.

4.0 REFERENCES

[1] Harary, F. (1959). *A criterion for unanimity* in French's theory of social power.

[2] Friedkin, N. E. & Johnson, E. C. (1990). *Social influence and opinions*. Journal of Mathematical Sociology 15, 193.

[3] Schelling, T. C. (1978). *Micromotives and macrobehavior*. W. W. Norton.

[4] Castellano, C., Fortunato, S., & Loreto, V. (2009). *Statistical physics of social dynamics*. Rev. Mod. Phys. 81, 591.

[5] Kempe, D., Kleinberg, J., & Tardos, E. (2003). *Maximizing the spread of influence through a social network*. Proc. of 9th ACM SIGKDD international conference on Knowledge discovery and data

mining, 137.

[6] Galam, S. (2008). *SOCIOPHYSICS: A review of Galam models*. Internat. J. Modern Phys. C 19, 409.

[7] Granovetter, M. (1978). *Threshold models of diffusion and collective behavior*. Am. J. Sociol. 83, 1420.

[8] Bass, F. M. (1969). *A new product growth for model consumer durables*. Management Science 15, 215.

[9] Baronchelli, A., Dall'Asta, L., Barrat, A., & Loreto, V. (2007). *Nonequilibrium phase transition in negotiation dynamics*. Phys. Rev. E 76, 051102.

[10] Castelló, X., Baronchelli, A., & Loreto, V. (2009). *Consensus and ordering in language dynamics*. Eur. Phys. J. B 71, 557.

[11] Steels, L. (1995). *A self-organizing spatial vocabulary*. Artificial Life 2, 319.

[12] Baronchelli, A., Felici, M., Caglioti, E., Loreto, V., & Steels, L. (2006). *Sharp transition towards shared vocabularies in multi-agent systems*. J. Stat. Mech.:Theory Exp. P06014, 0509075.

[13] Dall'Asta, L. & Baronchelli, A. (2006). *Microscopic activity patterns in the naming game*. J. Phys. A: Math. Gen 39, 14851.

[14] Lu, Q., Korniss, G., & Szymanski, B. K. (2008). *Naming games in two-dimensional and small-world-connected random geometric networks*. Phys. Rev. E 77, 016111.

[15] Uzzi, B., Soderstrom, S., & Diermeier, D. (2011). *Buzz and the contagion of cultural products: The case of Hollywood movies*.

[16] Christakis, N. A. & Fowler, J. H. (2007). *The spread of obesity in a large social network over 32 years*. New Engl. J. Med. 357, 370.

[17] Kearns, M., Judd, S., Tan, J., & Wortman, J. (2009). *Behavioral experiments on biased voting in networks*. Proc. Natl. Acad. Sci. U. S. A 106, 1347.

[18] Watts, D. J. & Dodds, P. S. (2007). *Influentials, networks, and public opinion formation*. J. Cons. Res. 34, 441-458.

[19] Galam, S. & Jacobs, F. (2007). *The role of inflexible minorities in the breaking of democratic opinion dynamics*. Physica A 381, 366 - 376.

[20] Lu, Q., Korniss, G. & Szymanski, B. K. (2009). *The Naming Game in social networks: community formation and consensus engineering*. J. Econ. Interact. Coord. 4, 221.

[21] Biswas, S. & Sen, P. (2009). *Model of binary opinion dynamics: Coarsening and effect of disorder*. Phys. Rev. E 80, 027101.

[22] Mobilia, M., Petersen, A., & Redner, S. (2007). *On the role of zealotry in the voter model*. J. Stat. Mech.:Theory Exp. 2007, P08029.

[23] Xie, J., Sreenivasan, S., Korniss, G., Zhang, W., Lim, C., & Szymanski, B. K. (2011). *Social consensus through the influence of committed minorities*. Phys. Rev. E 84, 011130.

[24] Dickman, R. & Vidigal, R. (2002). *Quasi-stationary distributions for stochastic processes with an absorbing state*. J. Phys. A 35.

[25] Zhang, W., Lim, C., Sreenivasan, S., Xie, J., Szymanski, B. K. & Korniss, G. (2011). *Social influencing and associated random walk models: Asymptotic consensus times on the complete graph*. Chaos 21, 025115

[26] Zhang, W. & Lim, C. (manuscript). *Fluctuation of consensus time in the Naming Game model*.

[27] Xie, J., Emenheiser, J., Kirby, M., Sreenivasan, S., Szymanski, B. K., & Korniss, G. (2012). *Evolution of opinions on social networks in the presence of competing committed groups*. PLoS ONE (in press).

[28] Dykman, M. I., Mori, E., Ross, J., & Hunt, P. M. (1994). *Large fluctuations and optimal paths in chemical kinetics*. J. Chem. Phys. 100, 5735-5750.

[29] Arnold, V. I. (1988). *Geometrical methods in the theory of ordinary differential equations*. Springer.

[30] Herrmann, D. W., Klein, W., & Stauffer, D. (1982). *Spinodals in a long-range interaction system*. Phys. Rev. Lett. 49, 1262--1264.

[31] Ray, T. S. (1991). *Evidence for spinodal singularities in high-dimensional nearest-neighbor ising models*. Journal of Statistical Physics 62, 463--472.

[32] Maier, R. S. (1992). *Large fluctuations in stochastically perturbed nonlinear systems: applications in computing*. Addison-Wesley.

[33] Graham, R. & Tél, T. (1984-06-01). *On the weak-noise limit of Fokker-Planck models*. Journal of Statistical Physics 35, 729--748.

[34] Gang, H. (1987). S*tationary solution of master equations in the large-system-size limit*. Phys. Rev. A 36, 5782--5790.

[35] Luchinsky, D. G., McClintock, P. V. E., & Dykman, M. I. (1998). *Analogue studies of nonlinear systems*. Rep. Prog. Phys. 61, 889.