Opinion Formation in Ising Networks *

M. H. Afrasiabi, R. Guérin and S. S. Venkatesh
Department of Electrical and Systems Engineering
University of Pennsylvania
{afram@seas, guerin@ee, venkates@ee}.upenn.edu

Abstract—**We study a network of connected nodes where each node holds an *opinion* — a binary state that may update over time under the influence of a node's neighbors. Nodes have biased affinities, which logically partition the network into distinct *parties*. Nodes in the same party tend to have a positive influence on each other, but the extent to which this holds varies across nodes and depends on the chosen affinity model. This paper considers two variations on an Ising spin-glass network model that investigate opinion formation in such biased affinity systems. These models differ in how they determine the pairwise influence between nodes. The first of these in what we dub the *random interactions model* randomly selects the influence two nodes exert on each other based on their respective party affiliation. The second, a *profile-based model*, relies on a profile, a κ-bit vector of ± 1 entries based on the node's known positions regarding each of κ independent topics. In this model the similarity of the profiles of two nodes determines whether they have a positive or negative influence on each other's opinions. We investigate the formation of opinions under both models and characterize their equilibria. We show that while these systems always converge to an equilibrium, they differ in their number and types of equilibria. These differences manifest themselves in the level of influence of initial opinions, and in the likelihood of polarized outcomes across party lines.**

I. Preview of Attractions

Ising spin-glass-inspired models of random interactions borrowed from statistical physics have been adapted for use in economics, models of neural computation, and in social network settings [2], [3], [6], [7]. In this paper we describe a variation on this theme with a view to understanding essential features of opinion formation in social networks in a (highly) sanitized setting.

Consider a fully connected network of n nodes, where node i holds a binary (for or against) *opinion* $x_i \in \{-1, 1\}$ about an issue under consideration. A node's opinion evolves over time as a function of its own opinion and that of its network neighbors. A neighbor's opinion is weighed based on its *influence.* Neighbors have a symmetric influence on each other which depends on their level of affinity. Affinity is biased positively or negatively based on nodes' *party* affiliations. Nodes from the same party are more likely to exhibit a positive affinity bias. This paper investigates the extent to which such party-based influence biases affect opinion formation, and in particular which opinions emerge in each party at equilibrium.

The two models considered in the paper differ in how the inter-nodal influences are specified. In the *random interactions model*, neighborhood influences are selected randomly to be positive or negative with a bias based on party affiliation. Two principals from the same party are more likely to have a positive affinity, while the influence on each other of two principals from different parties is more likely to be negative. The second of the models we consider, the *profile-based model*, specifies affinities in a more nuanced fashion. In this model each node comes equipped with a *profile*—a vector of positions on a fixed set of prior issues with nodes in the same party more likely to take similar positions on those issues. Inter-nodal influences in this setting are determined based on profile similarity.

The two models share common properties, and in particular, opinions in both models always converge to stable equilibria. However, they also exhibit significant differences. Of most interest is the fact that with high probability the random interactions model gives rise to a polarized outcome with nodes in each party converging to a common opinion opposed to that of nodes in the other party. In contrast, the profile-based model permits a more diverse set of distinct equilibria, with the specification of opinion equilibrium driven by the initial distribution of opinions in each party.

II. Embedding a Party Structure in an Ising Network Model of Interaction

The basic Ising model is a stochastic system which specifies a dynamics on the vertices $\{-1, +1\}^n$ of the n-dimensional cube. The system is characterized by a symmetric stochastic matrix $[w_{ij}]$ of interaction weights. At each epoch the state of the system is represented by a state vector $\mathbf{x} = (x_1, \dots, x_n) \in \{-1, +1\}^n$ which represents a collection of spins or, in the current context, opinions in a community of n interconnected principals. Updates to the state are performed asynchronously according to some update schedule: a node i is selected and a state update $x_i \mapsto x_i'$ performed according to the sign of a linear form of the node's current inputs,

$$x_i' = \operatorname{sgn} S_i = \operatorname{sgn}\left(\sum_{j=1}^{n} w_{ij} x_j\right), \qquad (1)$$

* This work was supported by NSF grant NSF CCF-1137519.

all other nodal states are kept unchanged.[1] We refer to $S_i = S_i(\mathbf{x})$ as the *update sum* for node i. In our context x_j is the (current) opinion of node j, w_{ij} denotes the weight of the influence node j has on node i, and x'_i represents the updated opinion of node i. The updates determine a dynamics $\mathbf{x} \mapsto \mathbf{x}' \mapsto \mathbf{x}'' \mapsto \cdots$ in the state space of vertices where, at any update epoch, the new state differs from the previous state in at most one coordinate (if the node update actually resulted in a change in sign). The specific update schedule is not critical for our purposes; it suffices if each state is updated infinitely often with probability one. A simple deterministic update schedule with this property is a round-robin schedule of state updates; a stochastic example is provided by a random update schedule where the node whose state is to be updated is selected randomly and independently at each update epoch. Call any such update schedule *honest*.

We assume throughout that the matrix $[w_{ij}]$ of interaction weights is symmetric, $w_{ij} = w_{ji}$, and w_{ii} is non-negative. In the context of interacting principals in a social network this models a situation where inter-agent influences are bilateral and symmetric, each agent having a positive self-reinforcement. A key classical result in this setting says that *under any honest update schedule the system dynamics converges to a fixed point* (see [5], [2]). Any fixed point (or *equilibrium*) $\mathbf{x}^*$ of the system satisfies the stationary system of update equations

$$x_i^* = \operatorname{sgn}\left(\sum_{j=1}^{n} w_{ij} x_j^*\right) \qquad (1 \le i \le n)$$

and it is now naturally of interest to characterize the number and nature of such equilibria.

In the classical Ising paradigm, the weights $\{\, w_{ij}, 1 \le i < j \le n \,\}$ are independent, standard normal random variables. We consider a variation on this theme where there is an embedded party structure with individuals within a party more likely to have a positive influence on each other, while individuals across party lines tend to have a neutral or negative influence on each other.

The general setting is as follows. The nodes $\{1, \dots, n\}$ are partitioned into, say, m groups $G_1, \dots, G_m$ which determine party memberships in a multi-party system. To keep algebra transparent, we may suppose that the nodes have been so ordered that the indices increase monotonically with party number. For each group G_k, the intra-group interaction weights $\{\, w_{ij}, i < j, (i, j) \in G_k \times G_k \,\}$ form a system of (positively biased) exchangeable random variables. Likewise, inter-group interaction weights $\{\, w_{ij}, (i, j) \in G_k \times G_l \,\}$ form systems of (negatively biased) exchangeable random variables. The nature of the dynamics is now, of course, determined by the specifics of the interaction distributions.

[1]For definiteness, set $\operatorname{sgn}(0) = -1$.

We restrict attention in this paper to a symmetric two-party system leaving extensions for elsewhere. In the setting at hand, $\{G_1, G_2\}$ is a partition of the nodes into memberships in two parties. We suppose further that the intra-party distributions are the same for both parties, that is to say, $\{\, w_{i,j}, i < j, (i, j) \in G_1 \times G_1 \text{ or } (i, j) \in G_2 \times G_2 \,\}$ is a system of (positively biased) exchangeable random variables, and the inter-party distributions are complementary, that is to say, the negatives of the inter-party interaction weights have the same distributions as the intra-party interaction weights. The symmetries inherent in the situation permit us to simplify exposition and consider an equivalent single party system (though it should be borne in mind that these algebraic simplifications will not be available when there are more than two parties). We briefly sketch the argument.

Begin with a two party partition $\{G_1, G_2\}$ and an associated symmetric stochastic system of weights $[w_{ij}]$. Form a new membership partition $\{G'_1, G'_2\}$ by moving one member, say, k from G_2 into G_1, $G'_1 = G_1 \cup \{k\}$ and $G'_2 = G_2 \setminus \{k\}$, associating with the new partition a new symmetric system of weights $[w'_{ij}]$ where $w'_{ij} = w_{ij}$ if $i \neq k$ and $j \neq k$, $w'_{kk} = w_{kk}$, and $w'_{jk} = -w_{jk}$ for $j < k$, by simply negating the weights of all the interconnections incident on k. We may now establish an isomorphism between dynamics in the $\{G_1, G_2\}$ and $\{G'_1, G'_2\}$ systems by putting states $\mathbf{x} = (x_1, \dots, x_n)$ in the $\{G_1, G_2\}$ system into one-to-one correspondence with states $\mathbf{x}' = (x'_1, \dots, x'_n)$ in the $\{G'_1, G'_2\}$ system by setting $x'_j = x_j$ if $j \neq k$ and $x'_k = -x_k$. It is now easy to verify that the update sums in the two systems starting with states $\mathbf{x}$ and $\mathbf{x}'$, respectively, satisfy $S'_j(\mathbf{x}') = S_j(\mathbf{x})$ if $j \neq k$ and $S'_k(\mathbf{x}') = -S_k(\mathbf{x})$, and so a dynamics $\mathbf{x}(0) \mapsto \mathbf{x}(1) \mapsto \mathbf{x}(2) \mapsto \cdots$ in the $\{G_1, G_2\}$ system is exactly mirrored by the dynamics $\mathbf{x}'(0) \mapsto \mathbf{x}'(1) \mapsto \mathbf{x}'(2) \mapsto \cdots$ in the $\{G'_1, G'_2\}$ system, the symmetry of the distributions ensuring that all probabilities are preserved. By iterating the process we end up with a single party system with weights $\{\, w_{ij}, i < j \,\}$ forming an exchangeable system of random variables with a positive bias. The new system is stochastically equivalent to the original two party system, the dynamics in the two systems being isomorphic. The single party formulation provides the greatest transparency in the statement of the results and the proofs and *we assume without comment henceforth that we are dealing with an equivalent single party system of nodes where the weights $\{\, w_{ij}, i < j \,\}$ form a system of exchangeable random variables.*

III. Random Interactions Model

The interaction model most closely related to the classical Ising model of Gaussian interactions is to consider a system of independent variables with a drift. And the simplest of these arises when we have signed Bernoulli interactions. Suppose $1/2 < p \le 1$ and let $\{\, w_{ij}, i < j \,\}$ be a system of signed Bernoulli trials with success parameter p: $\Pr\{w_{ij} = 1\} = p$ and $\Pr\{w_{ij} = -1\} = 1 - p$. [In the equivalent two party system, the intra-party weights are $+1$ with probability p (and -1 with probability $1 - p$) for both parties, while the inter-party weights are $+1$ with probability $1 - p$ (and -1 with probability p). In other words, two nodes are likely to positively influence each other if they belong to the same party; they are likely to negatively influence each other if they belong to different parties.] The algebra is simplest if there are no self-interactions, $w_{ii} = 0$, and we so assume. Extensions of this simple model to varying distributions and self-reinforcement may be handled by tweaking this basic framework and are outlined in the discussions.

We begin by a characterization of the dominant equilibria in this setting.

We begin by the observation that a state $\mathbf{x}$ is a fixed point if, and only if, each partial sum $S_i(\mathbf{x}) = \sum_j w_{ij} x_j$ has the same sign as x_i.[2] This leads to the following *criterion*: a state $\mathbf{x}$ is a fixed point of the system if, and only if, $x_i S_i(\mathbf{x}) > 0$ for each i. As $S_i(-\mathbf{x}) = -S_i(\mathbf{x})$, it follows that if $\mathbf{x}$ is a fixed point then so is $-\mathbf{x}$, and vice versa. Thus, fixed points appear in pairs.

In view of our criterion, we see that, for any given state $\mathbf{x} \in \{-1, +1\}^n$,

$$\Pr\{\mathbf{x} \text{ is a fixed point}\} = 1 - \Pr\left(\bigcup_{j=1}^{n} x_j S_j(\mathbf{x}) \leq 0\right).$$

Write $\mathbf{x}^+ = (1, 1, \ldots, 1)$ for the vector all of whose components are $+1$ and in increasing compaction of notation, write $S_i(\mathbf{x}^+) = S_i^+$ for the partial sums corresponding to state $\mathbf{x}^+$. Naturally enough, we expect $\mathbf{x}^+$ and $\mathbf{x}^-$ to be fixed points. And this is indeed the case (in a suitable probabilistic interpretation). Identifying the dependence on n explicitly, write $P^+ = P_n^+$ for the probability that $\mathbf{x}^+$ is a fixed point.

Theorem 1. *Fix any $0 < \delta < 1$ and suppose*

$$p \geq \frac{1}{2} + \sqrt{\frac{\log(n/\delta)}{2(n-1)}}.$$

Then $P_n^+ \geq 1 - \delta$. In particular, if $p > 1/2$ is bounded away from $1/2$, then $P_n^+ \to 1$ as $n \to \infty$.

Proof: If the system is in state $\mathbf{x}^+$ the partial sums are given by

$$x_i^+ S_i^+ = \sum_{j \neq i} w_{ij} x_j^+ = \sum_{j \neq} w_{ij}.$$

The sum on the right represents a random walk with a positive drift. As the signed Bernoulli variables w_{ij} have expectation $2p - 1$, everything sets up nicely for an application of Hoeffding's inequality [4, Theorem 2] (see [8, Section XVI.1] for the particular version considered here). We hence obtain

$$\begin{aligned}
&\Pr(x_i^+ S_i^+ \leq 0) \\
&\quad = \Pr\left(\sum_{j \neq i} (w_{ij} - (2p - 1)) \leq -(n-1)(2p-1)\right) \\
&\leq \exp\left(-\frac{(n-1)^2 (2p-1)^2}{2(n-1)}\right) = \exp\left(-2(n-1)\left(p - \frac{1}{2}\right)^2\right).
\end{aligned}$$

By Boole's inequality, it follows that

$$\begin{aligned}
\Pr\left(\bigcup_{i=1}^{n} \left(x_i^+ S_i^+ \leq 0\right)\right) &\leq \sum_{i=1}^{n} \Pr\left(x_i^+ S_i^+ \leq 0\right) \\
&\leq n. \exp\left(-2(n-1)\left(p - \frac{1}{2}\right)^2\right) \leq \delta
\end{aligned}$$

for the given selection of p. ■

Our proof shows that the probability that $\mathbf{x}^+$, hence also $\mathbf{x}^-$, is a fixed point converges very fast indeed, exponentially in n, to one. In the rest of this paper we refer to these equilibria as *polarized equilibria*. [The corresponding equilibria in the equivalent two party model are states where all nodes in a party have an identical opinion which is opposed to the common opinion of the nodes in the other party.]

We can do a little better, the mechanism of proof permitting a characterization of the *region of attraction* around the equilibria $\mathbf{x}^+$ and $\mathbf{x}^-$. The term region of attraction is a little vague; more precisely, we would like to estimate the probability that a given initial state $\mathbf{x}$ is mapped, eventually, over possibly many asynchronous steps, into, say, the equilibrium $\mathbf{x}^+$, and determine for what range of Hamming distances $d(\mathbf{x}, \mathbf{x}^+)$ we obtain a high probability convergence to the equilibrium.

The situation with respect to $\mathbf{x}^+$ and $\mathbf{x}^-$ is symmetric. Suppose, for definitenes, that an initial state vector $\mathbf{x}$ is at Hamming distance $0 < m < n/2$ from $\mathbf{x}^+$. Write B_m for the set of all $\binom{n}{m}$ such states $\mathbf{x}$,

$$B_m = \{\mathbf{x} : d(x, \mathbf{x}^+) = m\}.$$

For any $\mathbf{x}$ in B_m, let $M(\mathbf{x})$ be the set of m nodes that have different opinions under $\mathbf{x}$ and $\mathbf{x}^+$, that is to say,

$$M(\mathbf{x}) = \{i : x_i \neq x_i^+\}.$$

In a vivid nonce terminology, we call these nodes *non-conforming*. Now, let $S_i(\mathbf{x})$ be the update sum of node i under state $\mathbf{x} \in B_m$,

$$S_i(\mathbf{x}) = -\sum_{j \in M(\mathbf{x})} w_{ij} + \sum_{j \notin M(\mathbf{x})} w_{ij}.$$

A preliminary estimate of the probability that all nodes have positive update sum under state $\mathbf{x} \in B_m$ sets the stage.

[2] There is an irritating possibility of the sum being zero in which case the adopted convention of the sign function becomes important. But this has an exponentially small probability and we will ignore this nuisance. Alternatively, assume n is even.

Lemma 1. *Fix any* $m < n/2$ *and suppose* $\mathbf{x} \in B_m$. *Then*

$$\Pr\left(\bigcap_{i=1}^{n}(S_i(\mathbf{x}) > 0)\right) \geq 1 - n \cdot \exp\left(-\frac{(n-1-2m)^2\ (2p-1)^2}{2(n-1)}\right).$$

Proof: By Boole's inequality, we see that

$$\Pr\left(\bigcap_{i=1}^{n}(S_i(\mathbf{x}) > 0)\right) = 1 - \Pr\left(\bigcup_{i=1}^{n} S_i(\mathbf{x}) \leq 0\right)$$
$$\geq 1 - \sum_{i=1}^{n} \Pr\left(S_i(\mathbf{x}) \leq 0\right). \quad (2)$$

The event $\{S_i(\mathbf{x}) \leq 0\}$ occurs if, and only if,

$$-\sum_{j \in M(x)} w_{ij} + \sum_{j \notin M(x)} w_{ij} \leq 0. \quad (3)$$

(Bear in mind that $w_{ii} = 0$.) Now first consider a node $i \in M(x)$. Then the left side of Eq. (3) is the sum of $n-1$ random variables with mean $(n+1-2m)(2p-1)$. Another application of Hoeffding's inequality shows then that

$$\Pr\left(S_i(\mathbf{x}) \leq 0\right) \leq \exp\left(-\frac{(n+1-2m)^2\ (2p-1)^2}{2(n-1)}\right)\ [i \in M(\mathbf{x})]. \quad (4)$$

Next, consider a node $i \notin M(x)$. Then the left side of Eq. (3) is the sum of $n-1$ random variables with mean $(n-1-2m)(2p-1)$. Hoeffding's inequality hence shows that

$$\Pr\left(S_i(\mathbf{x}) \leq 0\right) \leq \exp\left(-\frac{(n-1-2m)^2\ (2p-1)^2}{2(n-1)}\right)\ [i \notin M(\mathbf{x})]. \quad (5)$$

Comparing Eq. (4) and Eq. (5) we see that

$$\Pr\left(S_i(\mathbf{x}) \leq 0\right) \leq \exp\left(-\frac{(n-1-2m)^2\ (2p-1)^2}{2(n-1)}\right) \quad (6)$$

for all i. Substituting the bound on the right into Eq. (2) completes the proof. ■

The introduction of a little notation and terminology helps streamline the results. Fix $0 < p < 1$ and on the interval $[0, 1/2]$ define the function

$$f(\alpha) = 2(1-2\alpha)^2(p-1/2)^2 - h(\alpha)$$

where, with logarithms to base e, $h(\cdot)$ is the binary entropy function (in nats) defined by

$$h(\alpha) = -\alpha \log(\alpha) - (1-\alpha)\log(1-\alpha).$$

The function $2(1-2\alpha)^2(p-1/2)^2$ is decreasing in the interval $[0, 1/2]$ while $h(\alpha)$ is increasing in this interval. It follows that $f(\alpha)$ decreases monotonically from a value of $f(0) = 2(p-1/2)^2 > 0$ at $\alpha = 0$ to a value of $f(1/2) = -\log 2 < 0$ at $\alpha = 1/2$. By the intermediate value theorem of calculus, it follows that f has a unique root $\alpha_0 = \alpha_0(p)$ in the interior of the interval $(0, 1/2)$ at which $f(\alpha_0) = 0$. Fig. 1 shows the dependence of α_0 on p.

Say that a state $\mathbf{x}$ lies in the *attraction region* of the polarized equilibrium $\mathbf{x}^+$ if, starting with $\mathbf{x}$ as the initial state, the state updates converge, eventually, to the fixed point $\mathbf{x}^+$. Write $A^+(\mathbf{x})$ for the event that $\mathbf{x}$ is in the attraction region of $\mathbf{x}^+$.

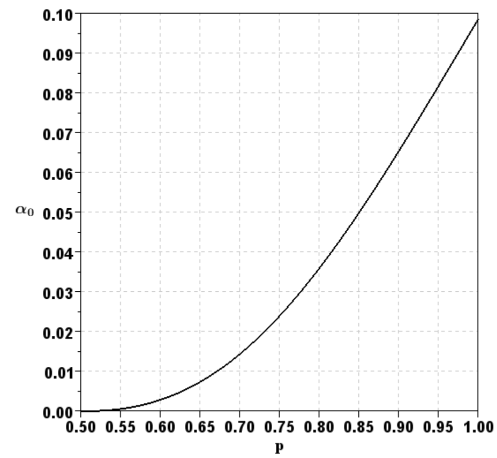

Fig. 1: Lower-bound on the radius of attraction region for the polarized equilibria as a function of p.

Theorem 2. *Select any tiny, positive* ϵ. *Fix any* $1/2 < p < 1$ *and any value* $0 < \alpha < \alpha_0(p)$. *If* $\mathbf{x}$ *is any state with* $d(\mathbf{x}, \mathbf{x}^+) \leq \alpha n$, *then* $\Pr(A_x^+) > 1 - \epsilon$ *whenever* n *is sufficiently large.*

Proof: Suppose that the system is at a state $\mathbf{x} \in B_m$ for some $m \leq \alpha n$ and the event $\bigcap_{i=1}^{n}\{S_i(\mathbf{x}) > 0\}$ occurs. Now consider an arbitrary sequence of (asynchronous) opinion updates according to the update rule specified in Eq. (1). Since all of the update sums are positive, the first node to change its opinion is a non-conforming node that becomes conforming, the update moving the system to a state $\mathbf{x}' \in B_{m-1}$. At this point we say that *one round* of updates has happened, and the system has *shrunk* one step towards the (closer) polarized equilibrium. We denote by $R_\mathbf{x}$ the event that the system moves from a particular state $\mathbf{x} \in B_m$ to any state $\mathbf{x}' \in B_{m-1}$ in *one round* of updates. It is now clear that the occurrence of the event $\bigcap_{i=1}^{n}\{S_i(\mathbf{x}) > 0\}$ implies the occurrence of $R_\mathbf{x}$ and so, by lemma 1, we obtain

$$\Pr(R_\mathbf{x}) \geq \Pr\left(\bigcap_{i=1}^{n}\{S_i(\mathbf{x}) > 0\}\right)$$
$$\geq 1 - n \cdot \exp\left(-2n\left(1 - 2\alpha - \frac{1}{n}\right)^2 (p-1/2)^2\right),$$

or, what is the same thing,

$$\Pr(R_\mathbf{x}^c) \leq n \cdot \exp\left(-2n\left(1 - 2\alpha - \frac{1}{n}\right)^2 (p-1/2)^2\right).$$

Another deployment of Boole's inequality shows now that

$$\begin{aligned}
&\Pr\left(\bigcup_{x\in B_m} R_x^c\right)\\
&\le \binom{n}{m} n\cdot\exp\left(-2n\left(1-2\alpha-\frac{1}{n}\right)^2 (p-1/2)^2\right)\\
&\le \frac{\exp(h(\alpha)\cdot n)}{\sqrt{n}} n\cdot\exp\left(-2n\left(1-2\alpha-\frac{1}{n}\right)^2 (p-1/2)^2\right)\\
&= \sqrt{n}\cdot\exp\left[n\left(h(\alpha)-2\left(1-2\alpha-\frac{1}{n}\right)^2 (p-1/2)^2\right)\right].
\end{aligned} \tag{7}$$

The expression on the right bounds from above the probability that at least one of the states $x \in B_m$ fails to demonstrate the shrink property R_x. One more application of Boole's inequality now gives

$$\begin{aligned}
\Pr\left(\bigcup_{m\le\alpha n}\bigcup_{x\in B_m} R_x^c\right) &\le \sum_{m\le\alpha n}\Pr\left(\bigcup_{x\in B_m} R_x^c\right)\\
&\le n\cdot\Pr\left(\bigcup_{x\in B_m} R_x^c\right) \le n\cdot n^{\frac{1}{2}}.\exp(-n\cdot f_n(\alpha)),
\end{aligned} \tag{8}$$

where $f_n(\alpha) = 2(1-2\alpha-1/n)^2\ (p-1/2)^2 - h(\alpha)$.

For every choice of $0 < \epsilon \le 1$, the right side of Eq. (8) is strictly less than ϵ if

$$\frac{\frac{3}{2}\log(n) - \log(\epsilon)}{n} < f_n(\alpha).$$

Since the left side of this inequality goes to 0 as n grows large, if $f_n(\alpha)$ is bounded away from zero then there exists $n(\epsilon)$ such that for all $n \ge n(\epsilon)$,

$$\Pr\left(\bigcup_{m\le\alpha n}\bigcup_{x\in B_m} R_x^c\right) < \epsilon.$$

It now remains to be shown that if $\alpha < \alpha_0$ then $f_n(\alpha) > 0$ is bounded away from zero for n sufficiently large. Arguing as for $f(\alpha$, we see that $f_n(\alpha)$ is strictly decreasing in the interval $0 \le \alpha \le \frac{1}{2} - \frac{1}{2n}$ and goes from $f_n(0) = 2(1-1/n)\ (p-1/2)^2 > 0$ at $\alpha = 0$ to $f_n\left(\frac{1}{2}-\frac{1}{2n}\right) < 0$ at $\alpha = \frac{1}{2}-\frac{1}{2n}$. By the intermediate value theorem again, it follows that f_n has exactly one root $\alpha_{0,n}$ in the interior of the interval $\left[0, \frac{1}{2}-\frac{1}{2n}\right]$. But $f_n(\alpha)$ differs from $f(\alpha)$ only in a term of order $1/n$ and, by examination, it is clear that $\alpha_{0,n} = \alpha_0 + O(n^{-1})$. It follows that if $\alpha < \alpha_0$ then $\alpha < \alpha_{0,n}$ for sufficiently large n and this concludes the proof. ■

To summarise, the states x^+ and x^- are equilibria (with high probability) for large n, each having a large (linear in n) region of attraction. The estimates can be improved but the results are intuitive (and limiting) and we won't expend any further effort in this direction.

A. Special Cases and Extensions

An examination of the proofs of the two theorems of this section shows that the basic method of proof extends to settings where there is self-reinforcement, $w_{ii} > 0$, and also to settings where the signed Bernoulli variables w_{ij} arise from different Bernoulli processes. We leave these details for elsewhere [1].

B. Simulation results

The theorems of this section show that as n becomes large, the polarized equilibria become persistent, their regions of attraction becoming larger as p increases. As a result, we should see more and more nodes of the same opinion, the population of nodes holding the majority opinion growing to potentially span the whole space. Fig. 2 illustrates this behavior through simulation for networks of different sizes. The figure illustrates the presence of a "threshold effect," where, for moderate population sizes, a polarized outcome rapidly arises when p exceeds $1/2$.

IV. Profile-Based model

The random interactions model is tempting in its simplicity and elegance but, as suggested by the results of the previous section, results in opinions of a bland conformity along party lines. Somewhat more nuanced and varied opinion formations arise in the profile-based model that we describe next.

A. Model formulation

In the profile-based model each node i has a *profile* $\pi_i = (\pi_{i1}, \ldots, \pi_{i\kappa}) \in \{-1, +1\}^\kappa$ where each entry in the profile takes a positive $(+1)$ or negative (-1) value based on the node's known position regarding one of κ independent topics. The influence weight w_{ij} between two nodes is specified as the dot product of their profiles: $w_{ij} = \pi_i \cdot \pi_j = \sum_{k=1}^{\kappa} \pi_{ik}\pi_{jk}$. We suppose that the profile bits are randomly chosen for each node and independently across nodes. Reusing notation for success probabilities, suppose $1/2 < p < 1$. The sequence of profile bits $\{\pi_{ij}, 1 \le j \le \kappa, 1 \le i \le n\}$ is then supposed to constitute a family of independent, signed Bernoulli variables with success parameter p: $\Pr\{\pi_{ij} = +1\} = p$, $\Pr\{\pi_{ij} = -1\} = 1 - p$. It is now clear that the collection of weights $\{w_{ij}, 1 \le i < j \le n\}$ forms an exchangeable system of random variables in the de Finetti sense, each weight w_{ij} having a positive bias. [Our description is for the equivalent single party formulation. In the two party formulation, the profile bits for members of one party are $\pm$Bernoulli(p) while the profile bits for members of the other parts are $\pm$Bernoulli($1-p$). Thus, if i and j are nodes in the same party then w_{ij} is likely to be positive; if they are from different parties then w_{ij} is likely to be negative. The equivalence between the two formulations is now as outlined in Section II.]

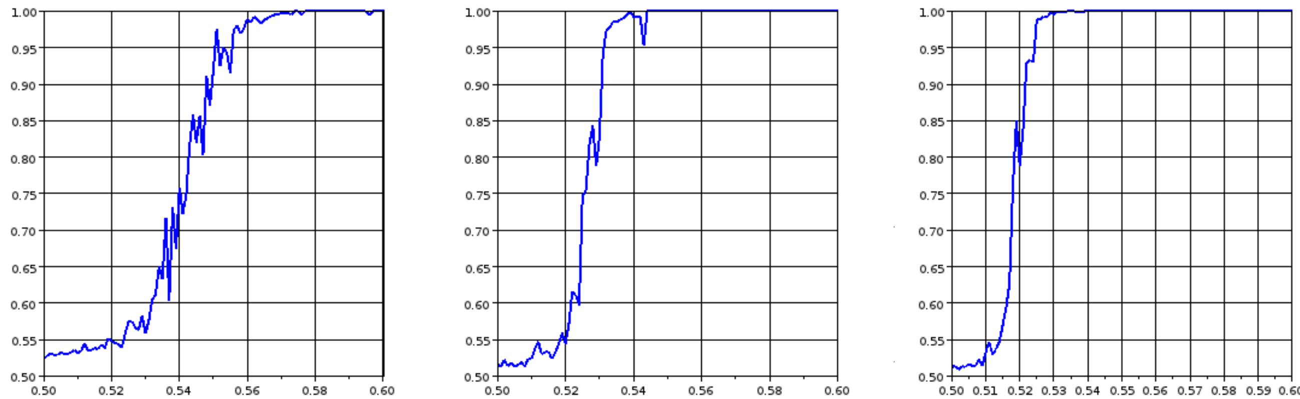

(a) $n = 400$ (b) $n = 1000$ (c) $n = 2000$

Fig. 2: The relative size of the majority opinion at equilibrium in the random interactions model.

B. Solution tools and results

For a profile size of κ, there are 2^κ distinct profile "types," with Π_ν, $0 \le \nu < 2^\kappa$, denoting the possible profile types. Nodes of the same profile type can be grouped into what we term a *cluster*. Let *cluster* C_ν be the set of all nodes whose profile is equal to Π_ν. In what follows we show that nodes in the same cluster always converge to the same opinion. *Cluster opinions*, therefore, fully describe the state of the system, so that equilibria only need to be characterized at the cluster level, *i.e.*, what opinion prevails in each cluster.

1) Cluster-based characterization of equilibria: In this section, we first establish that all nodes in a cluster share the same opinion at equilibrium, and then proceed to demonstrate that with probability approaching one, cluster opinions only depend on expected cluster sizes.

Proposition 1. *At equilibrium, nodes in a cluster all have the same opinion.*

Proof: Let i and j be any two nodes in C_ν. Per our criterion, at equilibrium, the update sums for i and j satisfy $x_i S_i > 0$ and $x_j S_j > 0$. To verify, we have

$$\begin{aligned} x_i S_i &= x_i \sum_k (\Pi_\nu \cdot \Pi_k)\, x_k \\ &= x_i \sum_{k \neq i, j} (\Pi_\nu \cdot \Pi_k)\, x_k + \left(x_i^2 + x_i x_j\right)(\Pi_\nu \cdot \Pi_\nu) \\ &= x_i \sum_{k \neq i, j} (\Pi_\nu \cdot \Pi_k)\, x_k + \kappa \left(x_i^2 + x_i x_j\right) \\ &\triangleq x_i S + \kappa (x_i x_j + 1) > 0, \end{aligned}$$

and likewise,

$$\begin{aligned} x_j S_j &= x_j \sum_{k \neq j, i} (\Pi_\nu \cdot \Pi_k)\, x_k + \kappa \left(x_j^2 + x_j x_i\right) \\ &= x_j S + \kappa (x_j x_i + 1) > 0. \end{aligned}$$

Now suppose the proposition is not true, *i.e.*, $x_i x_j = -1$. Then we replace $x_i x_j = -1$ and $x_j = -x_i$ in the above equations to get $x_i S > 0$ and $-x_i S > 0$. Contradiction. ■

In light of this result, we only need to consider *cluster level* opinions; write X_ν for the common opinion of cluster C_ν. The state of the system can, therefore, be described by the vector of cluster opinions $\mathbf{X} = (X_0, \ldots, X_{2^\kappa - 1})$. State $\mathbf{X}$ is stable if all clusters are stable. As with nodes, the updated opinion of cluster C_ν is the sign of its update sum, and can be written as

$$\begin{aligned} X_\nu^{\text{new}} &= x_i^{\text{new}} \qquad \text{for } i \in C_\nu \\ &= \operatorname{sgn}\left(\sum_{\nu'=0}^{2^\kappa - 1} (\Pi_\nu \cdot \Pi_{\nu'})\, c_{\nu'} X_{\nu'}\right) \\ &= \operatorname{sgn}\left(\sum_{\nu'=0}^{2^\kappa - 1} (\Pi_\nu \cdot \Pi_{\nu'})\, \frac{c_{\nu'}}{n} X_{\nu'}\right) \\ &\triangleq \operatorname{sgn}(S_\nu)\ , \end{aligned} \tag{9}$$

where c_ν denotes the size (number of nodes) of cluster C_ν, S_ν is a function of the cluster opinions $\mathbf{X}$, and the realized cluster sizes. The stability of a state $\mathbf{X}$ is, therefore, determined by the vector of realized cluster sizes $c = (c_0, \ldots, c_{2^\kappa - 1})$.

Lemma 2. *For a particular vector of cluster sizes,* c*, a state* $\mathbf{X}$ *is an equilibrium if, and only if,*

$$X_\nu = \operatorname{sgn}\left(\sum_{\nu'=0}^{2^\kappa - 1} (\Pi_\nu \cdot \Pi_{\nu'})\, c_{\nu'} X_{\nu'}\right) \qquad (0 \le \nu \le 2^\kappa - 1).$$

By the pigeon-hole principle, this result proffers a crude upper bound of 2^{2^κ}, the maximum number of cluster level states, for the number of possible equilibria. While the bound is not particularly sharp, it is already informative: *the number of equilibria in the profile-based model* *is bounded*. As we shall see, the number of equilibria is not trivially small nor are they so large (growing with n) that analysis is fruitless. The number of equilibria falls in the Goldilocks zone of not too many and not too few.

The following proposition shows that the pigeonhole bound can be improved by considering symmetries across clusters. For that purpose, we order clusters "lexicographically" according to their profile value, with Π_0 the all "-1" cluster and $\Pi_{2^\kappa-1}$ the all "1" cluster. Under this ordering, $\Pi_\nu = 2\overrightarrow{b}_\nu - \mathbb{1}_{1\times\kappa}$, where $\overrightarrow{b}_\nu$ is the binary vector[3] representation of decimal ν.

Proposition 2. *At equilibrium, clusters C_ν and $C_{2^\kappa-1-\nu}$ have opposite opinions.*

Proof: For the respective profiles Π_ν and $\Pi_{2^\kappa-1-\nu}$ of these clusters, we have

$$\begin{aligned}\Pi_\nu + \Pi_{2^\kappa-1-\nu} &= 2\overrightarrow{b}_\nu - \mathbb{1}_{1\times\kappa} + 2\overrightarrow{b}_{2^\kappa-1-\nu} - \mathbb{1}_{1\times\kappa}\\ &= 2\left(\overrightarrow{b}_\nu + \overrightarrow{b}_{2^\kappa-1-\nu} - \mathbb{1}_{1\times\kappa}\right)\\ &= 2\left(\mathbb{1}_{1\times\kappa} - \mathbb{1}_{1\times\kappa}\right) = 0,\end{aligned}$$

and as a result

$$\Pi_\nu = -\Pi_{2^\kappa-1-\nu}.$$

On the other hand, since the system is at equilibrium, updated and current opinions are identical. Therefore using Eq. (9) we have

$$\begin{aligned}X_\nu = X_\nu^{\text{new}} &= \operatorname{sgn}\left(\sum_{\nu'=0}^{2^\kappa-1}\left(\Pi_\nu\cdot\Pi_{\nu'}\right)c_{\nu'}X_{\nu'}\right)\\ &= -\operatorname{sgn}\left(\sum_{\nu'=0}^{2^\kappa-1}\left(\Pi_{2^\kappa-1-\nu}\cdot\Pi_{\nu'}\right)c_{\nu'}X_{\nu'}\right)\\ &= -X_{2^\kappa-1-\nu}^{\text{new}} = -X_{2^\kappa-1-\nu}.\end{aligned}$$

This concludes the proof. ■

Proposition 2 establishes that any equilibrium of the system is of the form $\mathbf{X} = [\mathbf{X}^L, -\mathbf{X}^L]$, where $\mathbf{X}^L$ is a vector of size $2^{\kappa-1}$ consisting of -1 and $+1$ entries only. This yields the tighter upper bound of $2^{2^{\kappa-1}}$ on the number of possible equilibria. (As we shall see in the simulations, the bound is not particularly tight; several of these states may not be feasible equilibria.) In the case of $\kappa = 3$ this upper bound is $2^{2^{\kappa-1}} = 16$. As per Lemma 2, this will depend on the relative sizes of the clusters. While it is straight-forward to compute the *expected value* of cluster sizes, actual values may differ. In the following, we establish that, with probability approaching one, expected cluster sizes are typically sufficient to characterize possible equilibria.

Note that since the probability distribution of every profile entry is known, the probability μ_ν that a node profile will be of a certain type Π_ν can be readily computed. For $\kappa = 3$, for example, $\mu_0 = q^3$ and $\mu_1 = q^2p$, where $q = 1-p$. Expected cluster sizes can now be easily computed as follows.

Let $I_j(\Pi)$ be an indicator random variable which takes value $+1$ if $\pi_j = \Pi$, and 0 if $\pi_j \neq \Pi$. It is cleat that $\mathbb{E}\left(I_j(\Pi_\nu)\right) = \mu_\nu$. As the profiles of different nodes are independent of one another, $I_j(\Pi)$ and $I_{j'}(\Pi)$ are also independent for $j \neq j'$. The size c_ν of cluster C_ν can now be written in the form

$$c_\nu = |C_\nu| = \sum_{j=1}^{n} I_j(\Pi_\nu). \qquad (10)$$

By additivity, the expected cluster size is hence given by

$$\mathbb{E}(c_\nu) = \sum_{j=1}^{n}\mathbb{E}\left(I_j(\Pi_\nu)\right) = \sum_{j=1}^{n}\mu_\nu = n\mu_\nu.$$

In any realization of the profile vectors, cluster sizes will vary around these expected values, and these variations can conceivably affect the set of possible equilibria. Fig. 3 illustrates this by plotting the distribution of the number of equilibria obtained across a set of $10,000$ realizations for $\kappa = 3$ and different combinations of n and p. We see that while, as expected, the upper bound of 16 holds, the number of observed equilibria varies as a function of p. More interesting though is the fact that as n increases for a constant p (Fig. 3a and Fig. 3b), the number of observed equilibria appears to concentrate on fewer values. We formalize this insight next, starting with a lemma that bounds the probability of cluster size variations around their mean value.

Lemma 3. *Fix any $\epsilon > 0$. Then*

$$Pr\left(\left|\frac{c_\nu}{n} - \mu_\nu\right| \le \epsilon\right) \ge 1 - 2\exp\left(-2n\epsilon^2\right).$$

In words: the fractional cluster size c_ν/n is concentrated at its mean value μ_ν.

Proof: We start from Eq. (10) that expresses c_ν as the sum of n independent random variables, each of which satisfies $0 \le I_j(\Pi) \le 1$. An application of Hoeffding's inequality (Theorem 1 of [4]) shows then that

$$\Pr\left(\frac{c_\nu}{n} - \mu_\nu \ge \epsilon\right) \le \exp\left(-2n\epsilon^2\right)$$

and

$$\Pr\left(\frac{c_\nu}{n} - \mu_\nu \le -\epsilon\right) \le \exp\left(-2n\epsilon^2\right).$$

An application of Boole's inequality shows hence that

$$\Pr\left(\left|\frac{c_\nu}{n} - \mu_\nu\right| \ge \epsilon\right) \le 2\exp\left(-2n\epsilon^2\right).$$

Taking the complement of both sides finishes the proof. ■

Now that we know c_ν/n is close to μ_ν, we will show that for identifying the possible outcomes it suffices to only investigate the *expected value* of cluster sizes, $n\mu_\nu$, or the expected fractional size, μ_ν.

[3] If needed, with zeros appended to make length κ.

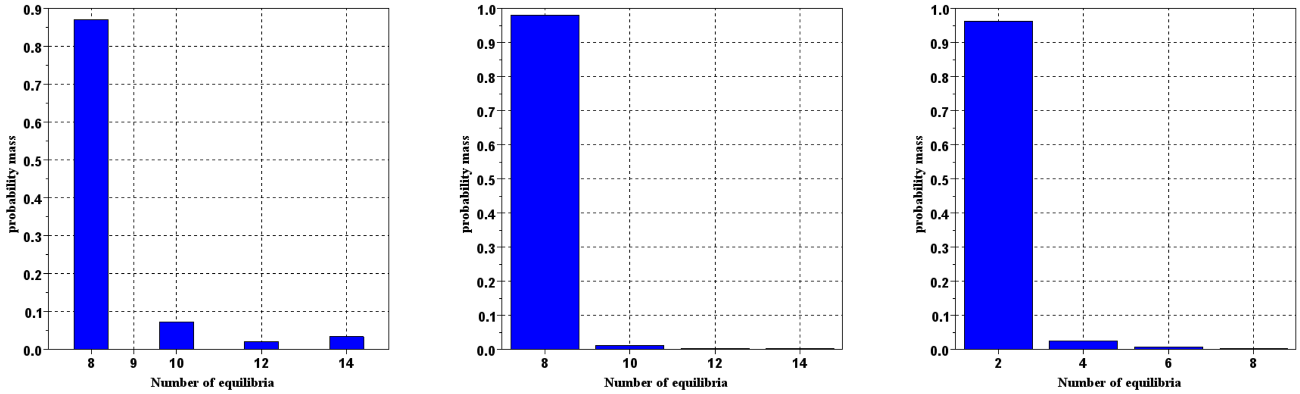

(a) $n = 400$ and $p = 0.8$ (b) $n = 2000$ and $p = 0.8$ (c) $n = 100$ and $p = 0.9$

Fig. 3: Probability mass function of the number of realized equilibria for $\kappa = 3$.

Before we proceed, note that μ_ν is a function of p, the probability that a profile entry is positive. As a result, the expected update sum $\overline{S}_\nu$ for cluster ν is a function of state $\boldsymbol{x}$ as also of p, and it will be useful to take explicit note of this by writing $\overline{S}_\nu = \overline{S}_\nu(\boldsymbol{x}, p)$. It will also be convenient to define

$$\sigma(p) \triangleq \min_{\substack{0 \le \nu \le 2^\kappa - 1 \\ \boldsymbol{x} \in \{-1,+1\}^n}} \left| \overline{S}_\nu(\boldsymbol{x}, p) \right|.$$

Theorem 3. *For any p with $\sigma(p) > 0$, the probability Q_n that the set of equilibria of the system under actual cluster sizes c_ν is the same as that under expected cluster sizes $n\mu_\nu$ is bounded from below by*

$$Q_n \ge 1 - 2^{\kappa+1} \exp\left(-n \left(\frac{\sigma(p)}{2^\kappa} \right)^2 \right).$$

In particular, $Q_n \to 1$ as $n \to \infty$.

Proof: In the proof of Lemma 3 we saw that the probability that any of the c_ν values lies outside $[\mu_\nu - \epsilon, \mu_\nu + \epsilon]$ is bounded by $2 \exp\left(-2n\epsilon^2\right)$. By Boole's inequality we see that the probability that at least one of the c_ν values lies outside that interval is bounded by $2^\kappa \times 2 \exp\left(-2n\epsilon^2\right)$. Taking complements we obtain

$$\Pr\left(\bigcap_\nu \left\{ \left| \frac{c_\nu}{n} - \mu_\nu \right| \le \epsilon \right\} \right) \ge 1 - 2^{\kappa+1} \exp\left(-2n\epsilon^2\right).$$

We recall that

$$S_\nu = \sum_{\nu'=0}^{2^\kappa - 1} \frac{(\Pi_\nu \cdot \Pi_{\nu'})}{\kappa} \frac{c_{\nu'}}{n} x_{\nu'}, \tag{11}$$

and so, replacing $\frac{c_{\nu'}}{n}$ with $\mu_{\nu'}$ in Eq. (11) can introduce an error of at most ϵ to each term in the sum and at most $\epsilon 2^\kappa$ to the total sum. Therefore

$$\Pr\left(\bigcap_\nu \left\{ \left| S_\nu - \overline{S_\nu} \right| \le \epsilon 2^\kappa \right\} \right) \ge 1 - 2^{\kappa+1} \exp\left(-2n\epsilon^2\right). \tag{12}$$

If $\sigma(p) \neq 0$, then we can pick $\epsilon < \frac{\sigma(p)}{2^\kappa}$, *e.g.*, $\epsilon = \frac{\sigma(p)}{2^{\kappa+1/2}}$ and Eq. (12) becomes

$$\Pr\left(\bigcap_\nu \left\{ \left| S_\nu - \overline{S_\nu} \right| \le \epsilon 2^\kappa \right\} \right) \ge 1 - 2^{\kappa+1} \exp\left(-n \left(\frac{\sigma(p)}{2^\kappa} \right)^2 \right).$$

But $\left| S_\nu - \overline{S_\nu} \right| \le \sigma(p)$ means that S_ν and $\overline{S_\nu}$ have the same sign. Since the sign of the update sums determines the dynamics of the system, this concludes the proof. ■

The zeros of $\sigma(p)$ are determined by the zeros of the constituent polynomials comprising the update sums and so are finite in number. It follows that $\sigma(p)$ is strictly positive except at a finite set of p values.

Corollary. *Except for a finite set of p values, the set of equilibria in a profile-based network with random cluster sizes is, with probability approaching one as $n \to \infty$, the same as that for a deterministic profile-based network with cluster sizes fixed at their expected values.*

As a result of Theorem 3 we can describe the equilibria of the profile-based model in a compact way. In the following we formulate an explicit matrix equation for these equilibria. As before, let Π_i, $i = 0, \ldots, 2^\kappa - 1$, be the i^{th} profile type of length κ. Then we can represent the system by a graph of 2^κ nodes, where each node corresponds to a cluster. The combined effect of the expected size μ_j of cluster j together with the dot product $\Pi_j \cdot \Pi_i$ determines the influence of cluster j on cluster i. These effects are lumped into the weighted edge from j to i which we represent as

$$a_{ij} = \mu_j (\Pi_j \cdot \Pi_i), \; i, j = 0, \ldots, 2^\kappa - 1.$$

The adjacency matrix $A = [a_{ij}]$ for the graph of clusters can now be written in the form

$$\begin{bmatrix} \kappa p^{\kappa} & p^{(\kappa-1)} q \Pi_1 \cdot \Pi_0 & p^{(\kappa-2)} q^2 \Pi_2 \cdot \Pi_0 & \dots \\ p^{\kappa} \Pi_0 \cdot \Pi_1 & \kappa p^{(\kappa-1)} q & & \dots \\ \vdots & \vdots & \ddots & \\ p^{\kappa} \Pi_0 \cdot \Pi_{2^{\kappa}-1} & \dots & & \end{bmatrix}$$

whence the fixed points of the system are the solution to the system of simultaneous equations specified by

$$\mathrm{sgn}(A\mathbf{x}) = \mathbf{x},$$

where the signum operation applied to a vector is to be interpreted component-wise.

In the next subsection we compute the number of equilibria for p values close to 1.

2) Dominance of opinion for large p: In this subsection we show that for p values sufficiently close to 1, the cluster sizes become so disproportionate that the diversity in opinions gets eliminated. In particular, when p is sufficiently large, the opinions of most clusters at equilibrium will be determined be the opinion of one *dominant* cluster.

Let $C^{(\lambda)}$ denote a cluster whose profile has exactly λ entries that are $+1$ (and, hence, $\kappa - \lambda$ entries that are -1).

Lemma 4. *There exists $p^* < 1$ such that for all p in the interval $p^* \leq p < 1$ the opinion of any cluster $C^{(\lambda)}$ agrees with that of $C^{(\kappa)}$ if $\lambda > \kappa/2$ and disagrees with that of $C^{(\kappa)}$ if $\lambda < \kappa/2$.*

Proof: The expected size of (the only) $C^{(\kappa)}$ is p^{κ} which eventually outgrows that of all other clusters combined, as p approaches 1. Since any cluster with $\lambda < \kappa/2$ or $\lambda > \kappa/2$ has a non-zero edge-weight to $C^{(\kappa)}$, this cluster contributes a dominant effect and determines the opinion of $C^{(\lambda)}$, either positively or negatively, based on the sign of their edge-weight. ■

We conclude with a collection of refinements whose proofs we defer in the interests of space (see [1]).

While Lemma 4 determines the fate of clusters with $\lambda > \kappa/2$ and $\lambda < \kappa/2$ entries taking value $+1$ in their profiles, It does not say anything about the case where $\lambda = \kappa/2$. When κ is an even number, there are *centric* clusters that have exactly $\kappa/2$ entries of $+1$ in their profiles. The next proposition shows that these centric clusters are not affected by the dominance effect of Lemma 4. This is because they incur an overall zero influence from the outside world if p is large enough. Consequently, those clusters can decide independently of the rest of the network.

Proposition 3. *There exists $p^* < 1$ such that, if $p^* \leq p < 1$, then the clusters of type $C^{(\kappa/2)}$ are unaffected by the ensemble of clusters corresponding to profiles that are not exactly balanced.*

C. Special cases

In this subsection, we analyze the system for some special cases of profile length.

1) Profile of size 1: The following theorem says that for $\kappa = 1$, the system has only two equilibria. Note that 'when $\kappa = 1$, the profile vector π_i is just a scalar.

Proposition 4. *For any network size n, the profile-based model with profile size $\kappa = 1$ has exactly two stochastic equilibria.*

2) Profile of size ∞: The case $\kappa = \infty$ presents another extreme of profile size.

Proposition 5. *Suppose $\kappa = \kappa_n$ grows sufficiently rapidly with n so that*

$$\kappa_n = \frac{2n^2}{(4p^2 - 4p + 1)^2}(\log(n) + \log(n+1)).$$

Then with probability tending to one as $n \to \infty$, the only equilibria are the pair of polarized equilibria $\mathbf{x}^+$ and $\mathbf{x}^-$.

A very large historical record captured in the profile has the effect of smoothing out all the wrinkles in the dynamics and collapsing the equilibria to two.

Note that while both Proposition 4 and Proposition 5 predict only two equilibria, the compositions of these equilibria are very different—stochastic in the one case, deterministic in the other.

3) Profile of size 2: If $\kappa = 2$, any profile is one of these types:

$$\Pi_a = (+1, +1)$$
$$\Pi_b = (-1, -1)$$
$$\Pi_c = (+1, -1)$$
$$\Pi_d = (-1, +1).$$

Two nodes with types a and b behave regardless of nodes with types c and d, and vice versa. This is because those profiles have an inner product of zero which results in zero edge-weight. As a result, nodes with types a and b will behave similar to the $\kappa = 1$ case, demonstrating two different outcomes. Similarly, nodes with types c and d will demonstrate two different outcomes. Overall, the results are independent of the value of p; we obtain four different equilibria which are determined *stochastically*, *i.e.*, based on the specific initialization of profile values.

D. Simulation results

The concentration argument that was presented in Theorem 3 facilitates the numerical investigation of possible equilibria for systems with a finite profile size. Assuming the expected cluster sizes, the number of possible equilibria as a function of p is plotted in Fig. 4 for different κ values. This number is irrespective of the network size n. As expected from Lemma 4 for $\kappa = 3$ and $\kappa = 5$, when p gets sufficiently close to 1 only two

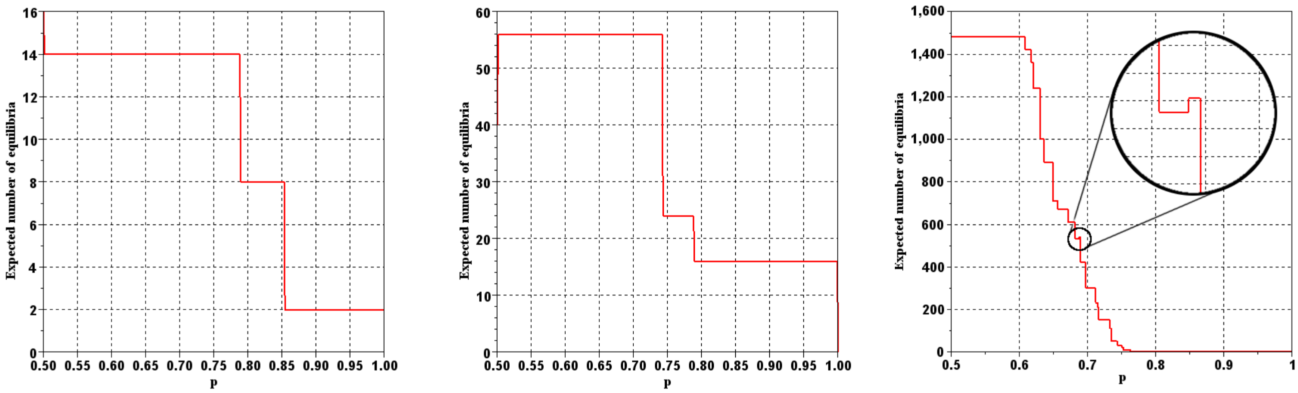

(a) $\kappa = 3$ (b) $\kappa = 4$ (c) $\kappa = 5$

Fig. 4: The expected number of equilibria in the profile-based model. As seen in (c), this number is not monotone in p.

equilibria remain feasible. In the case of the even number $\kappa = 4$, however, the number of equilibria never drops to 2. This is because, per Proposition 3, there are some clusters that are not forced to change their opinions, and therefore they maintain some level of diversity.

While the behaviors in figures 4a and 4b show a monotone decrease of number of equilibria as p increases, Fig. 4c gives evidence for a counter example where there is a rise in the number of equilibria.

V. Conclusion

References

[1] M. H. Afrasiabi, R. Guérin, and S. Venkatesh. Opinion formation in ising networks. Technical report, U. Pennsylvania, 2013. Available at http://www.seas.upenn.edu/~afram/tr/is2013.pdf.

[2] P. Baldi and S. Venkatesh. Number of stable points for spin-glasses and neural networks of higher orders. *Physical Review Letters*, 58:913–916, 1987.

[3] J.-P. Bouchaud. The (unfortunate) complexity of the economy. *ArXiv e-prints*, Apr. 2009.

[4] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

[5] J. J. Hopfield. NNeural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the USA*, 79:2554–âĂŞ2558, 1982.

[6] J. Nadal, D. Phan, M. Gordon, and J. Vannimenus. Multiple equilibria in a monopoly market with heterogeneous agents and externalities. *Quantitative Finance*, 5(6):557–568, 2005.

[7] M. Ostilli, E. Yoneki, I. Leung, J. Mendes, P. Lio, and J. Crowcroft. Ising model of rumour spreading in interacting communities. Technical Report UCAM-CL-TR-767, University of Cambridge, 2010.

[8] S. S. Venkatesh. *The Theory of Probability, Explorations and Applications*. Cambridge University Press, 2013.