Available online at www.sciencedirect.com
SCIENCE DIRECT
ELSEVIER
Simulation Modelling Practice and Theory 14 (2006) 407–422
SIMULATION MODELLING PRACTICE AND THEORY
www.elsevier.com/locate/simpat

Social influence and the emergence of norms amid ties of amity and enmity

James A. Kitts

Department of Sociology, University of Washington, 202 Savery Hall, Box 353340, Seattle, WA 98195, United States

Received 7 June 2004; accepted 5 May 2005
Available online 4 November 2005

Abstract

This paper explores the coevolution of social networks and behavioral norms. Previous research has investigated the long-term behavior of feedback systems of attraction and influence, particularly the tendency toward homogenization in arbitrary cultural fields. This paper extends those models by allowing that norms diffuse not only by simple contagion but through intentional sanctioning behavior among peers. Further, the model allows for negative relations, where actors differentiate themselves from enemies while seeking to align themselves with friends. Sociometric maps reveal non-trivial system dynamics—structural bifurcation, discrimination between factions, and cycles of deviance and solidarity—emerging from a few elementary agent-level assumptions.

Keywords: Social networks; Group dynamics; Social influence; Norms; Organizations

1. Introduction

Several basic theories of interaction have grown from one of the starkest regularities in the social world: the tendency of social ties to connect individuals who are similar in attributes, attitudes, or behaviors. This observed lawlike regularity of relative homogeneity in social relations—or *homophily* [39,48]—has inspired prominent "first principles" for models of emergent structure.

One such principle is the preference among actors to choose interaction partners who are similar to themselves [30,31,41,50–52]. Psychologists have proposed the "Law of Attraction" [7,8], according to which actors hold a positive affective bias toward similar others, which leads them to choose these similar alters as interaction partners. Other work has posited the same preference for homophilous sociometric choices, but has emphasized shared knowledge and reduced costs of communication [9,18,44,55] rather than an affective bias. Both research programs suppose that similarities drive network change. This sociometric choice explanation is

E-mail address: kitts@u.washington.edu

1569-190X/$ - see front matter
doi:10.1016/j.simpat.2005.09.006

most compelling when the attributes underlying homophily are fixed (e.g. race, gender) or change slowly (e.g. age, language) relative to changes in social ties.

Classic work in group dynamics [12,19,23,59] and recent work in social networks [25,46] posit a converse explanation for homophily, arguing that interpersonal influence operates through social ties, engendering common attitudes or behaviors among friends and other close relations. This social influence explanation is most compelling when social relations are fixed or change slowly relative to the attributes underlying homophilous choices (e.g. opinions, behaviors).

When individuals' attributes and social ties are both subject to change, sociometric choice and influence processes may operate simultaneously to generate homophily. This combination of differential attraction and influence creates a self-reinforcing dynamic in which similarity increases the likelihood or intensity of dyadic influence and influence reduces differences between the interaction partners. This feedback of course leads to inferential problems for natural observation and statistical analysis, and so empirical research does not give us much leverage on the relative importance of sociometric choice and social influence when the two processes operate together. For this important and challenging case, exploring formal models in "thought experiments" allows us to observe the qualitative implications of our assumptions and to design critical tests for empirical research.

Recent work has used simulation to study these dynamics of attraction and influence, primarily their effects on cultural diffusion and convergence. For example, several prominent projects [2,9,44] have assumed that agents who share more cultural traits have a higher probability of interacting, while interaction promotes further cultural similarity. Applications to complete networks such as organizational task groups [26,37] assume that all members interact on some level, but relative similarity determines the strength of dyadic influence. The positive feedback of homophilous sociometric choices and social influence generates a local homogenization that some have presented as an explanation for the emergence of "cultural norms" [38]. In fact, the tendency toward homogeneity is so robust that scholars are left with the opposite puzzle of explaining social differentiation [43–45]. They ask under what conditions social differences can emerge and survive in populations governed by generic processes of homophilous attraction and social influence.

In this paper, we contribute two elements of model design. First, we supplement the conventional models of arbitrary cultural diffusion by allowing that actors have vested interests in their behaviors. In pursuing these interests, agents learn from experience (adjust their behavior to seek rewards and avoid punishments) and also aim to influence their peers' behavior by applying social sanctions (rewards and punishments). As a second innovation, we allow that social ties may have a negative valence, so that actors differentiate themselves from enemies and seek to align themselves with friends.

Most previous research has used computational experiments to investigate stability conditions of system-level equilibria (especially cultural uniformity). We instead explore system dynamics out of equilibrium, examining sequential sociometric maps of group structure from a single simulation run. This qualitative analysis of model behavior will reveal intriguing dynamics at the system level that emerge from a very basic set of agent-level propositions and that are not observable from the model's equilibrium response surface.

1.1. *Toward consideration of inductive influence: peer pressure and regulatory interest*

The simulation literature reviewed above has examined the case of arbitrary influence, where the behavioral states are empty symbols, inconsequential to the actors who carry them. We refer to this process of arbitrary transmission or imitation as *mimetic influence* [37].

In this paper we incorporate a core property of social norms, the notion that agents may intentionally influence peers to behave in ways that those peers would otherwise not behave. Following Heckathorn [28], we consider the case where each agent's preferences over its own behavior ("inclinations") differ from its preferences over others' behavior ("regulatory interests"). This mismatch of inclinations and regulatory interests implies that agents have a reason to pressure one another, such as through informal social sanctions. Specifically, we model a group of interdependent agents facing a problem of collective goods production, where all agents receive a benefit when peers' "work" toward a collective good, but a net cost of working implies that each agent faces a perverse temptation to "shirk", or allow others to bear this burden. If actors use sanctions to affect peers' behavior, we call the resulting social pressure *inductive influence*.[1]

1.2. Toward consideration of valued ties: xenophobia and disinfluence

While the notion of homophilous attraction has received much attention in the sociological and psychological literatures [14,36,66], only a few scholars have investigated the obverse process, by which individuals may be repulsed by others who are very different. This dynamic, which we call "xenophobia", leads to the same cross-sectional pattern as homophilous attraction: disproportionate homogeneity in dyads connected by social ties. In fact, many empirical findings of homophily may have reflected some mixture of differential attraction and repulsion dynamics [56,62].

We know that disagreements may lead to negative interpersonal evaluations and thus local homogeneity through sociometric choice. Negativity is also relevant to the two forms of social influence we have mentioned. First consider mimetic influence: Classic research in group dynamics showed that disliked and dissimilar others may serve as "negative referents" [24,61] who inspire contrary behavior [40,63]. Second consider inductive influence: Research has shown that persuasive communication from similar others tends to positively influence attitudes [4,5,47] and behavior [60], while information from dissimilar others causes inverted change, even when the dimension of similarity is involuntary (e.g. race). The possibility for tie value to reverse the valence of influence has also received much empirical attention [1,57]. For example, Sampson and Insko [58] show that individuals tend to make judgments both in accord with liked others and contrary to disliked others.

In a standard network influence model without negativity, agents may imitate their friends and ignore other agents that are not friends. An extension to inductive influence suggests that agents may also be influenced by sanctions (such as social approval or disapproval) from their friends while ignoring sanctions from strangers. In either case, variation in social ties may be interpreted as social distance, density, or probability of interaction. Considering negativity allows that agents try to differentiate themselves from negative referents (a process opposite to imitation) and that they are paradoxically reinforced by sanctions from negative referents. Negative ties cannot then be interpreted as a great social distance or a very low density or probability of interaction because negative values violate the definitions of distance, density, and probability. The strength of influence may vary with distance and the density of interaction, but positive or negative valence must be evaluative.

2. Model design

This project borrows technical insights from the study of neural networks in artificial intelligence [3,13,42]. We build on an attractor neural network model originally developed and investigated by Hopfield [32–34] and extended to the study of group dynamics by Nowak and Vallacher [53]. This model includes a set of nodes each of which is characterized by a vector of *states* and connected to other nodes by a vector of *weights*. These networks "learn" stable configurations as nodes iteratively adjust their states and weights to minimize "energy" (also "stress" or "dissonance") across all relations with other nodes.

In this model, weights change over time through a Hebbian learning rule [27]: The weight between node i and node j—denoted w_{ij}—changes as a function of the concordance of nodes' states over time. To the extent that i and j tend to occupy the same state at the same time, the tie between them will grow more positive. To the extent that i and j occupy discrepant states, the tie between them will grow more negative. We refer to these self-organizing network dynamics as *structural learning*. This process operates in tandem with mimetic influence, where nodes adjust their states to minimize local dissonance, given a configuration of weights to other nodes.

To illustrate the operation of mimetic influence and structural learning, Fig. 1 shows a simple triad in which nodes i, j, and k must take a position (+1 or −1) on a single issue. If we additionally assume that these nodes

[1] Mimetic influence and inductive influence correspond to French and Raven's [24] *referent power* and *reward* (or *coercive*) power, respectively, but we focus on the dyad rather than the sender of influence.

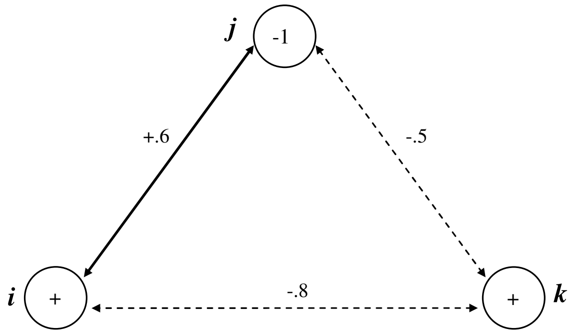

Fig. 1. A simple attractor neural network.

have a varying attraction for one another, as indicated by the weights on the ties between them, we can find the sources of dissonance in the graph.

In this illustration, nodes i and k each begin with a $+1$ state, while node j begins with a -1 state. Because i and j are connected by a positive weight, their contradictory states imply dissonance. Similarly, nodes i and k are frustrated by their agreement, given their negative relationship. Nodes j and k are satisfied despite their contradictory states, due to their negative path. Nodes iteratively adjust their weights (structural learning) and states (mimetic influence) to resolve local dissonance.[2]

There are obvious parallels here between the architecture of these neural networks and classic theories of group dynamics: dissonance theory [20,22], social comparison theory [21,54], and balance theory [29,52]. These theories predict instability where one individual values a cognitive object positively and a friend or ally values that object negatively. The friends may relieve this dissonance by cooling their ties toward one another or by resolving their disagreement through social influence. Structural Balance Theory [11,15] addresses the special case where the cognitive object is a third actor, predicting balance where the sign product in a triad is positive and imbalance where it is negative. Note that our model does not assume a direct tendency toward structural balance in triads, as actors have no knowledge of other actors' weights.

We depart substantially from the neural network architecture by allowing that each node in the network is an "agent" that has preferences over its own state and over the states of other agents. If agents experience outcomes as good (rewarding) or bad (punishing), then they may adjust their own behavioral propensities to seek rewards and avoid punishments. This agent-level learning ("reinforcement learning") in turn allows agents to sanction one another according to their own regulatory interests. More specifically, each agent is faced with one choice to *work* or *shirk* as well as choices to convey informal sanctions (e.g. *approval* or *disapproval*) to each agent in the network.

Like mimetic influence, this inductive influence is moderated by the social tie between sender and receiver. For example, a message received through a strong tie is subjectively more potent than an identical message received through a weak tie. Valence determines the qualitative force of both mimetic and inductive influence. If agent i has a positive tie toward agent j, then i will value j's social approval and use j as a positive referent or role model. Negative valence reverses the "meaning" of a message as it passes through the tie. That is, approval received from enemies will negatively reinforce (disaffirm) a behavior and disapproval received from enemies will positively reinforce a behavior. We thus do not assume that "approval" is an intrinsically rewarding outcome for behavior. The messages between connected agents are meaningful only in the context of their relationship.

In a network of N agents, each agent i has $N+1$ states: one compliance decision to "work" ($c_i = 1$) or "shirk" ($c_i = -1$) and N decisions to "approve" ($a_{ik} = 1$) or "disapprove" ($a_{ik} = -1$) of agent k, where k is an index over each of the N members. Note that agents must choose whether to approve of themselves (where $k = i$)—a choice we call "self-esteem"—but this choice does not have reinforcement consequences. Each agent has $N - 1$ ties toward other agents—varying in weight (w_{ij}) from -1.0 to $+1.0$. For exposition, we distinguish

[2] Our implementation of structural learning and mimetic influence differs from this simple illustration in two ways: First, for a network of N nodes, the single binary state in Fig. 1 is replaced with a vector of $N+1$ binary states. Second, the weights are not necessarily symmetric, so the influence network is a directed graph. These points will become clear in our explication of the model.

two qualities of this weight: *strength* (varying continuously from 0.0 to 1.0) and *valence* (taking a value of either -1 or $+1$). Each agent also has one reflexive tie that is fixed at $+1.0$. Of course an agent has always been 100% similar to itself, so a reflexive tie of unity seems to be a reasonable simplifying assumption. Reflexive ties play no role in the model (i.e. an agent does not imitate or reinforce itself), but they make clusters in the sociometric maps interpretable as cohesive cliques.

2.1. Mimetic influence

Mimetic influence allows that agents use their peers as positive or negative role models. Mimetic pressure to comply ($-1.0 \leqslant MC_i \leqslant 1.0$) on agent i is the average of the compliance choices ($c_j = \pm 1$) of all $N-1$ other agents j, weighted by the tie from i to j ($-1.0 \leqslant w_{ij} \leqslant 1.0$):

$$MC_i = \frac{1}{N-1}\sum_{j=1}^{N} w_{ij}c_j, \quad j \neq i \tag{1}$$

To the extent that members of i's positive reference group (other agents to which i is positively tied) agree on a given behavior, i receives mimetic pressure to engage in that behavior. The opposite is true for a negative reference group, or agents to whom i is negatively tied. With increasing agreement among i's friends to engage in a behavior and among i's enemies to engage in the opposite behavior, mimetic influence on i becomes stronger and we might say that a "group standard" has emerged in the reference group.

Mimetic influence works in the same way for the approval choices. Agent i's received mimetic pressure to approve of any agent k ($-1.0 \leqslant MA_{ik} \leqslant 1.0$) is an average of the choices to approve of k ($a_{jk} = \pm 1$) for all $N-1$ other agents j, weighted by the tie from i to j ($-1.0 \leqslant w_{ij} \leqslant 1.0$):

$$MA_{ik} = \frac{1}{N-1}\sum_{j=1}^{N} a_{jk}w_{ij}, \quad j \neq i \tag{2}$$

(Note that for self-esteem, $k = i$, while j is an index over each of $N-1$ other agents.)

2.2. Reinforcement learning and inductive influence

Each agent has an internal bias to adopt a -1 or $+1$ value for each of its $N+1$ binary states. Unlike most neural networks, nodes in this model may adjust their internal biases for adopting states based on subjective evaluations of previous outcomes. We call this process *reinforcement learning*. By this learning process, based on the "Law of Effect" [65], acts that were reinforced positively are repeated more often than acts that were reinforced negatively. At round t, any given agent independently examines all choices it made at the previous round, evaluates the events that followed those choices, and then adjusts its biases for these decisions to seek subjective rewards and avoid punishments. The evaluation is a weighted combination of the collective benefit received, the cost of working, and the set of social sanctions received (A) in the previous round.

The symbol C represents an individual share of the collective benefit, or $1/N$th of the sum of individual choices to work ($c_i = 1$) or shirk ($c_i = -1$):

$$C = \frac{1}{N}\sum_{i=1}^{N} c_i \tag{3}$$

The symbol A_i represents i's total approval by peers, specifically the average of approval choices ($a_{ji} = \pm 1$) of all $N-1$ other members j, weighted by i's attraction or repulsion to j (w_{ij}):

$$A_i = \frac{1}{N-1}\sum_{j=1}^{N} a_{ji}w_{ij}, \quad j \neq i \tag{4}$$

Both C and A_i thus vary in the range $[-1.0, 1.0]$. Actors must weigh these two sources of reinforcement—social approval and the collective good—to derive a total evaluation of the outcome. The parameter *affective dependence* represents the subjective importance assigned to received social feedback (relative to the collective good) in determining an agent's satisfaction with the outcome of a previous choice. Formally, agent i's satisfaction with the compliance decision (SC_i) takes into account the collective benefit (C), personal cost (e), and approval (A_i) outcomes associated with i's choice to work ($c_i = 1$) or shirk ($c_i = -1$) in the previous round:

$$SC_i = (1 - \beta_i)C + \beta_i A_i - ec_i \tag{5}$$

where β_i is i's affective dependence—ranging from 0.0 (agent i cares only about the collective good) to 1.0 (agent i cares only about social approval)—and e is a uniform cost of working that is imposed for any agent that chooses to work ($c_i = 1$), insensitive to the agent's affective dependence.[3]

The decision by i to approve of j is evaluated in the same way, but satisfaction with approval (SA_{ij}) is evaluated independently within each i–j dyad. That is, in evaluating the outcome of i's previous sanctioning behavior toward j, agent i considers only j's work effort and whether j approved of i (weighted by their tie):

$$SA_{ij} = (1 - \beta_i)c_j + \beta_i a_{ji} w_{ij} \tag{6}$$

In using satisfaction to adjust their biases to behave, agents see the origin as a reference point. They find levels of SC or SA above 0.0 to be satisfying and find levels below 0.0 to be dissatisfying. Agent i's internal bias toward compliance (BC_i) is reinforced accordingly by positive and negative satisfaction. When an agent chooses to work ($c_i = 1$) and is rewarded ($SC_i > 0$) or when an agent chooses to shirk ($c_i = -1$) and is punished ($SC_i < 0$), the agent's bias to work increases toward the maximum value of 1.0. However, when an agent chooses to shirk ($c_{it} = -1$) and is rewarded ($SC_i > 0$) or chooses to work ($c_{it} = 1$) and is punished ($SC_i < 0$), its bias to work decreases toward the minimum value of −1.0. We can express this dynamic relation as a difference equation for the change in i's bias to comply (ΔBC_i):

$$\Delta BC_i = SC_i c_i \theta(BC_i) \tag{7}$$

where the symbol Δ represents the change in a variable over the interval from time t to time $t + 1$. (We suppress time subscripts here for brevity.) A linear limiting function (θ) adjusts the increment of change, such that change slows as the variable (in this case BC_i) approaches its extreme values of −1.0 and 1.0.

$$\theta(x) = \begin{cases} \frac{1-x}{2} & \text{if } x \text{ is increasing} \\ \frac{1+x}{2} & \text{if } x \text{ is decreasing} \end{cases} \tag{8}$$

Incremental change in i's bias to approve of j (ΔBA_{ij}) is computed in the same way:

$$\Delta BA_{ij} = SA_{ij} a_{ij} \theta(BA_{ij}), \quad j \neq i \tag{9}$$

Agents' bias toward self-approval is fixed at unity ($BA_{ii} = 1.0$) by design, so negative self-esteem results only from external mimetic pressure (e.g. disapproval from friends or approval from enemies). Note that self-esteem does not directly influence an agent's behaviors, as an agent does not imitate itself or reinforce its own choices. Self-esteem does play a role in structural learning (described in Section 2.4). The possibility for agents to disapprove of themselves avoids wiring in the result where agents only make friends with those who approve of them.

2.3. Choice algorithm

In updating states ("choosing behaviors"), an agent must weigh the force of mimetic influence against its own internal bias, or the lessons of history. We represent agent i's relative weight of individual and social learning with parameter δ_i, *normative dependence*, ranging from 0.0 (agent i makes decisions based entirely on individual learning through reinforcement) to 1.0 (agent i adjusts behavior only through social learning).

[3] For convenience, we use positive and negative values to represent the binary states of incurring or escaping the cost of compliance. This parallels our use of positive and negative values to represent the binary compliance decision. An agent who complies contributes a positive amount ($c_i = +1$) to the public good but incurs a negative individual payoff ($-e$). Conversely, shirkers subtract from the public good ($c_i = -1$) but enjoy a positive individual payoff (e).

The behavioral propensity to comply (PC_i) is simply the sum of i's bias toward compliance (BC_i) and mimetic pressure on i to comply (MC_i), weighted by δ_i:

$$PC_i = (1 - \delta_i)BC_i + \delta_i MC_i \tag{10}$$

Normative dependence operates similarly for the approval choice, where agent i weighs bias (BA_{ij}) against mimetic influence (MA_{ij}) for the propensity to approve of agent j (PA_{ij}):

$$PA_{ij} = (1 - \delta_i)BA_{ij} + \delta_i MA_{ij} \tag{11}$$

For any decision, this propensity fits in the range $[-1.0, 1.0]$. An agent's probability to engage in a behavior (compliance or approval) corresponds to an approximate sigmoid function of this propensity. Specifically, this propensity in $[-1.0, 1.0]$ is compared to a uniform random number in the range $[-.5, .5]$; propensities toward either extreme thus yield a strictly determined outcome and responses become increasingly random as propensities approach the midpoint. The probability of choosing a $+1$ (or -1) state is 0.5 when the propensity is exactly at the midpoint of its range (0.0).

2.4. Structural learning: evolution of social networks

Structural learning is the tendency of the network to seek an overall configuration of weights that represents a local energy minimum (given a configuration of weights), due to the dissonance-reducing efforts of individual agents. Weights between agents start with uniform random values between -1.0 and 1.0, then adjust iteratively according to the similarity perceived by each partner. Note that the uniform random initialization means that the weight matrix is asymmetric at the start. Also, we see below that any two agents i and j may have different perceptions of their similarity and thus may not update their mutual weights (w_{ij} and w_{ji}) congruently. Although there will be a tendency toward symmetry in weights over time, this is not explicitly assumed.

In computing its perceived level of similarity with each peer j, agent i compares j's observed behavior —a binary vector of $\{-1, 1\}$ scores—with i's own internal propensities ("intentions" or "dispositions") to perform those same behaviors, a vector of real numbers in $[-1, 1]$. This implementation reflects an information asymmetry that is inherent in social interaction: We may know our own intentions and dispositions, even if our observed behavior is a "fluke" (an accident or experiment), but we cannot see peers' dispositions and must infer this information from their behavior.[4] As a result, agent i's strong dispositions (propensities near ± 1.0) are weighted more heavily in judging similarity than its ambivalent dispositions (propensities near 0.0). This assumption increases stability of the model by making it less sensitive to random fluctuations. The outcome is also highly plausible: In the social world, an anomalous act by one person is unlikely to cause her to reevaluate all of her friendships, though it may lead her friends to adjust their opinions of her.

Agent i uses three domains of decisions to judge its similarity with other agents j: its *compliance* (PC_i and c_j), its *self-approval* (PA_{ii} and a_{ji}), and *other-approval* (PA_{ik} and a_{jk} averaged over all $k \neq i$). Each of these domains yields a score for i's perceived similarity with j ranging from -1.0 (j's behavior is opposite to i's strong dispositions) to $+1.0$ (j's behavior corresponds perfectly to i's strong dispositions). If i has a middling or undecided disposition for a given behavior, this score will approach zero and thus the behavior will have little effect on the weight from i to j (w_{ij}). The change in weight from i to j (Δw_{ij}) corresponds to i's average level of similarity or dissimilarity with j across these three domains:

$$\Delta w_{ij} = \frac{1}{3}\left(PC_i c_j + PA_{ii} a_{ji} + \sum_{j=1}^{N} \frac{PA_{ik} a_{jk}}{(N-1)}\right)\theta(w_{ij}), \quad i \neq j, k \tag{12}$$

where PC_i is i's propensity for compliance and PA_{ik} is i's propensity to approve of k, c_j is j's decision to comply, a_{jk} is j's approval of k, and θ is the function defined in (8), which ensures that w_{ij} asymptotically approaches the extremes of its range $[-1.0, 1.0]$.

[4] This property of the model can be derived from the classic actor–observer effect [35], where observers overestimate the role of stable dispositions in others' behavior but regard their own behavior as relatively undiagnostic of their own dispositions. Judgments of similarity in the model are then driven by comparisons of one's own dispositions with dispositional inferences based on peers' behavior.

The limiting function θ is used for all continuous updates in the model—adjustment of bias toward compliance (BC) and approval (BA) as well as adjustment of weights (w) between agents—so propensities and weights will never reach the extreme values of their range under the dynamics proposed here. We nevertheless report their ranges as closed intervals because we regard those extreme values as within the defined state space of the model (i.e. permissible as initial conditions).

Persistent agreement or disagreement between i and j will build solid positive or negative relationships over time, but both agents are sensitive to substantial changes in the other's behavior, such as simultaneous change in several states. They are quick to "resent" friends or "forgive" enemies following a major change in agreement. If actors in the social world have a path-dependent tendency to forgive friends and distrust concession by enemies, then this model should conservatively underestimate their tendency to lock into such stable relationships.

3. Methods

Previous research, reported in Kitts et al. [37], has explored part of the parameter space of this model. Their research examined the level of group cooperation and cohesion across the full range of normative dependence—from pure reinforcement learning to pure imitation—over a range of group sizes, from 3 to 15 members. Their model allowed only positive ties ($w_{ij} > 0.0$), however, omitting xenophobia and disinfluence. A later paper [43] allowed for negative ties, but omitted inductive influence and reinforcement learning. These previous projects manipulated parameters systematically and mapped the response surface of the model at equilibrium. However, the crucial case where mimetic influence and reinforcement learning are both strong and where peers' approval and peers' work are both important sources of reinforcement does not typically converge to an equilibrium. This paper will focus on out-of-equilibrium dynamics within a single representative simulation run, where parameter values are held constant in this important region of the parameter space.

Given an initially random configuration of states and weights, agents will minimize dissonance in their relations (through dual processes of structural learning and mimetic influence) while also adjusting their behavioral propensities according to previous rewards and punishments (reinforcement learning). However, there is no guarantee that they will achieve a globally optimal state in either structural balance or reinforcement.

3.1. Parameter settings for the simulation

In order to investigate the important case where social learning and reinforcement learning operate with roughly equal strength to change behavior, we examine the midpoint of the range of normative dependence. The normative dependence parameter (δ) is heterogeneous over agents, with values drawn from a normal distribution (truncated at 0.0 and 1.0) with mean 0.5 and standard deviation 0.15.

In order to investigate the important case where social approval is important to agents' choices—but not so important that it overwhelms consideration of the collective good—we examine the midpoint of the range of affective dependence. This paper describes a set of individuals who, on average, assign approximately equal value to any peer's work and approval by a close friend (i.e. w_{ij} near 1.0). Values for the affective dependence parameter (β) are also heterogeneous over the agents, drawn from a truncated normal distribution with mean 0.5 and standard deviation 0.15.

In order to make the work choice into a theoretically interesting collective action problem, we assign a moderate cost of effort, $e = 0.2$. This particular value is arbitrary, but ensures that agents always have a regulatory interest in pressuring peers to work, while they always have personal inclinations to shirk. This conflict of interest thus makes inductive influence non-trivial.

Following Newcomb's [50,52] classic study of friendship network dynamics (and many subsequent formal analyses of his data), this paper will fix the group size (N) at 17 members. This group size allows for a reasonable distribution in parameter values among individual agents and is comparable to moderately large organizational work teams.

3.2. Method of analysis

We will view agents' positions and trajectories in sociometric space, examining the relational patterns that precede and follow changes in compliance behavior. In order to make this analysis tractable, we simplify the matrix of attraction scores (weights) to a two-dimensional map using a principal components analysis at each point in time.[5] The first two eigenvectors give agents' coordinates or relative location in this space; we present them graphically, weighting by their corresponding eigenvalues. Also, the largest eigenvalue (λ_1) indicates the variance explained by the first eigenvector, or the extent that a single latent ''issue'' appears to underlie variance in relations. If the first eigenvalue λ_1 approaches the total variance σ^2 and thus the quotient λ_1/σ^2 approaches 1.0, this indicates that the group has polarized into two internally cohesive and mutually antipathetic cliques.

The coordinates in this two-dimensional space are meaningful only relative to one another (subject to any arbitrary rotation, translation, or reflection). For visual clarity, we align the figures employing a Procrustes transformation of the principal components, then center the scatter of points on the origin and divide by the most extreme value—so the numerical coordinates vary in the range $[-1, 1]$. These transformations facilitate visually tracking the movement of agents through sociometric space without distorting their relative position.

4. Results

Although we use statistical analyses to distill patterns from an overwhelming volume of longitudinal data (308 compliance and approval choices and 272 social ties changing each round), our aim is to provide an analytic narrative of essential patterns of the evolving system of relations and behavior. These dynamic patterns are not observable from the broader computational experiments, which have focused on the mean response surface over the entire range of parameter values. The narrative includes a historical graph of compliance choices followed by a series of maps that summarize the network structure in two dimensions at distinct points in time (rounds 3, 5, 7, and 10). Fig. 2 presents the history of all 17 agents' compliance over the first 50 rounds of interaction.

Fig. 2 illustrates that the model does not tend toward homogenization or any other steady state. However, it does reveal dynamic concordance among clusters of agents, as if they deliberately act in concert; for example, over half of the agents $\{2, 5, 6, 9, 11, 12, 13, 15, 17\}$ switch to compliance en masse in round 20. This concordance is surprising, as agents receive no information about each other's behavior until after making their own choices for each round. Given the simple adaptive model of agents' decision process, it is not obvious that agents will be able to coordinate such a simultaneous jump to compliance. Even so, it is important that they act in concert, because a critical mass of simultaneous workers is required to make working sustainable under the reinforcement learning dynamic. How may these simple adaptive agents coordinate such concerted action? We may find some leverage on this question by examining the dynamic structure of social influence. It is convenient then that the computational model allows us unlimited insight into the relations between actors.

At initialization, the simulation is a cacophony of behavior and relations, with agents scattered arbitrarily in sociometric space, but the system rapidly self-organizes into a more coherent structure. That is, agents form sets of friends who tacitly agree in their evaluations of third parties. To give a qualitative view of this evolving structure, the following map shows a distribution of agents' position in sociometric space. Actors with similar sets of ''friends'' and ''enemies'' occupy proximate positions in this map. Polarization of the group into mutually antagonistic cliques would appear as two sets of structurally equivalent agents, stacked at opposite ends of the horizontal axis (Dimension 1). A selection of these maps describes the relationship between individuals' compliance choices and their positions in sociometric space. Fig. 3 shows this map at the end of the third round.[6]

[5] We first transform the sociomatrix into a (symmetric) correlation matrix R, where cell R_{ij} represents the correlation of the weight vectors for node i and node j. See Bonacich [6] for discussion of using eigenvectors of the sociomatrix for identifying structurally similar subgroups and see Nakao [49] for an empirical application to Newcomb's 17 fraternity members, over 15 temporal observations.

[6] An animation of 50 rounds of the simulation is available online at: http://faculty.washington.edu/kitts/SIMPAT.html.

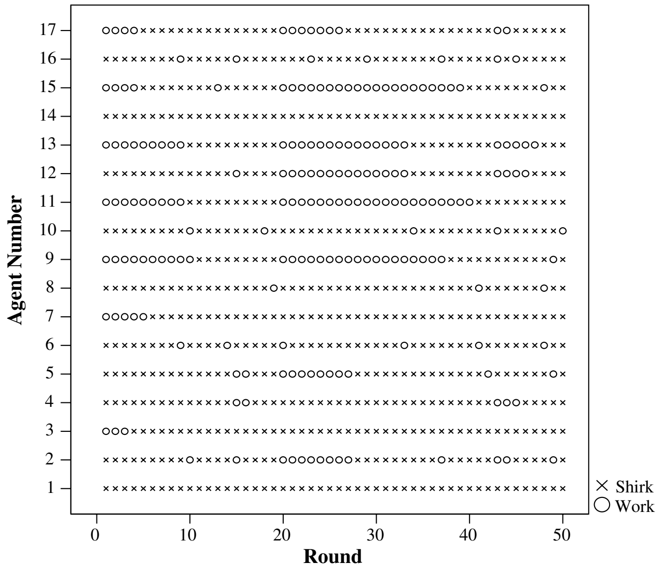

Fig. 2. Compliance choices of 17 agents over 50 rounds.

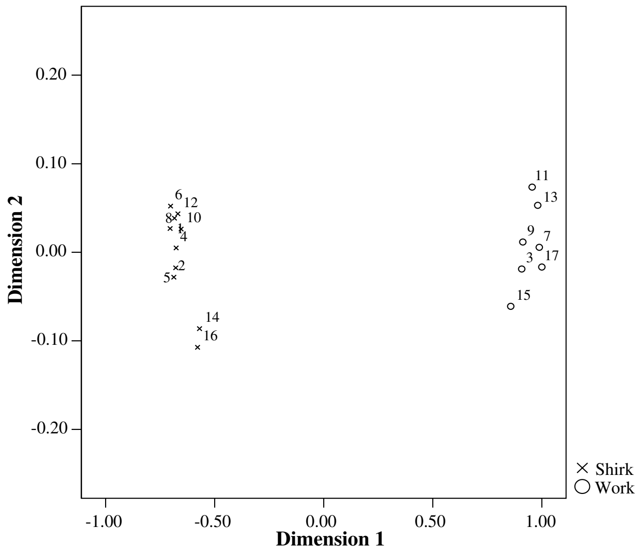

Fig. 3. Map of agent position in sociometric space at round 3.

See that the network segregates quickly. The first latent dimension explains 79.3% of the variance in transformed vectors of weights at this time. Further, this binary polarization maps onto compliance behavior exactly: All of the agents in the left wing are shirking and all of the agents in the right wing are working during this round.

Fig. 4 shows the fifth round, when the first dimension still explains three quarters of the variance in the matrix. This polarization largely represents factions of workers and shirkers, although Agent 3 switched to shirking during round 4 and is now beginning to drift from the worker clique toward the deviant shirkers in sociometric space. In round 5, tightly paired Agents 15 and 17 follow suit.

By round 7, Agent 3 is assimilated into an increasingly cohesive clique of shirkers, while Agents 15 and 17 continue to drift toward the shirkers in their social relations. Agent 7, who was on the sociometric margins of the worker clique, has defected and is now drifting in sociometric space. The three remaining cooperators {9, 11, 13} are a close triad, appearing as a tight cluster in Fig. 5.

The network is largely polarized again by round 9, but ambiguity arises at this point. In fact, Agent 7 seems to be challenging alliances and drawing members of both factions out through cross-cutting approval decisions. For example, Agent 7 may approve of some members of the working faction, maintaining partial agreement with both factions. The solidarity of the worker cluster is further disrupted by ephemeral trials in working (lasting only one round) by Agents {6, 16} in round 9 and by Agents {2, 10} in round 10.

Fig. 6 shows that these explorations in working fail to resurrect a cooperative clique, while Agents 11 and 13 fall into shirking, leaving Agent 9 as an isolated worker in the following round. Two rounds later, Agent 9 gives up and everyone shirks. Without the compliance choice as a salient distinction, the structural polarization breaks down and social relations fall into complete disarray. Agents are once again scattered across

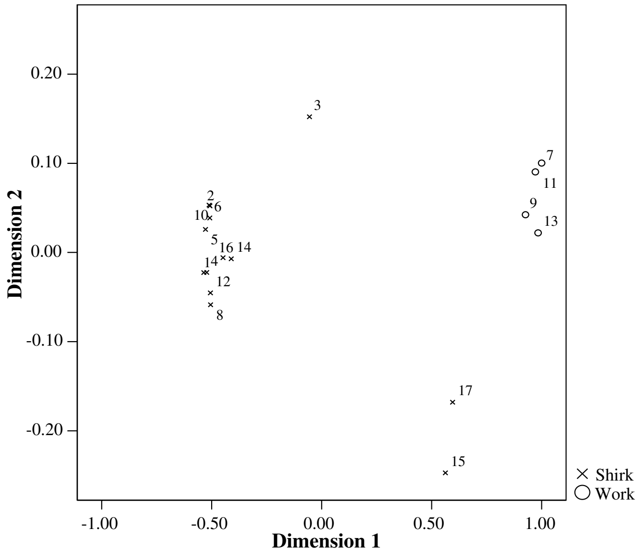

Fig. 4. Map of agent position in sociometric space at round 5.

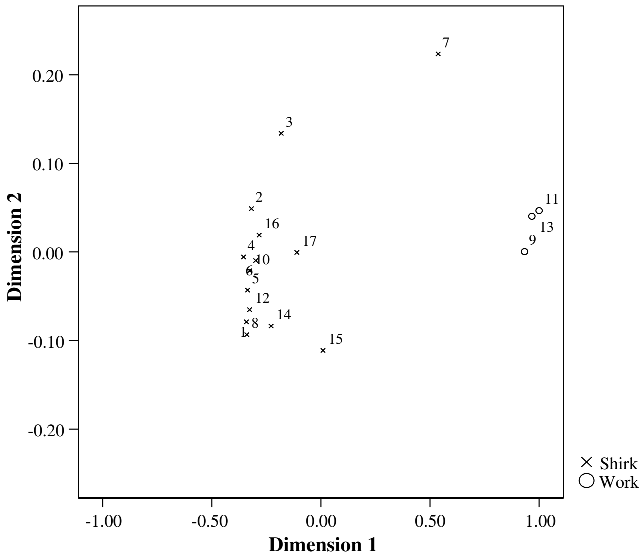

Fig. 5. Map of agent position in sociometric space at round 7.

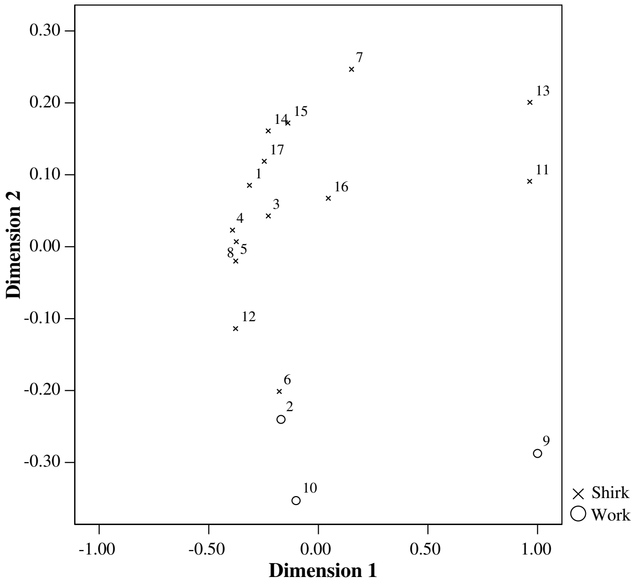

Fig. 6. Map of agent position in sociometric space at round 10.

sociometric space. The first dimension now explains less than a third of the variance in the matrix and scarcely more than the second.

In this condition of disorganized relations, social control is impossible to coordinate. Of course universal defection is not a globally stable equilibrium under the reinforcement learning dynamic. All agents remain sorely dissatisfied with the low level of the collective good, so negative reinforcement increases propensities to work. Isolated agents explore working over the next few rounds, until a faction $\{2, 4, 5, 12, 16\}$ simultaneously works at round 15. Even this is short-lived, however, as the clique is far too small to make the resulting production self-reinforcing. The agents continue to wander until round 20, when over half of the agents suddenly cooperate and create a new salient cleavage that stabilizes cooperation within a compliant clique.[7] The network maintains satisfying amity and enmity for several more rounds, until new subgroups break off and those cross-cutting cleavages undermine the primary polarization.

The qualitative pattern observed here recurs without end in this simulation and is representative of simulations in this region of the model's parameter space. Questions of whether such a pattern is inevitable or indefinite for this region cannot be strictly resolved using simulation. A simulation also cannot allow any inferences about the model's behavior in other regions of the parameter space. We know, for example, that a simulation would easily reach equilibrium in a very high range of normative dependence—where agents' personal bias and reinforcement learning are overwhelmed by mimetic influence. The model's response surface over these parameters has been investigated elsewhere, and we focus here on its dynamics in an interesting region where reinforcement learning and mimetic influence operate together. As our modest goal is to show a set of non-obvious macrolevel patterns that are *possible* consequences of a set of plausible microlevel dynamics, we need not prove that they will always obtain under the model.

5. Discussion and conclusion

This project has employed foundation assumptions from theories of social network dynamics and operant learning that have been extensively investigated in experimental research. Such elementary principles comprise

[7] Note that the production level is still not quite high enough here to exceed the cost of working and thus positively reinforce work efforts on its own. If a "worker clique" is able to maintain cooperation over any period of time, they must accomplish this through mimetic and inductive influence.

a rigorously specified agent-based model, which in turn generates group-level behavior dynamics that have been observed in naturalistic research. Our intriguing results warrant further investigation of the model, as we aim to account for emergent social phenomena using a set of basic but empirically grounded behavioral principles.

At the microlevel, we observe extensive variation in behavior among agents and over time that cannot be explained by agents' "attributes" (fixed parameter values). Mapping of sociometric structure presents an instructive portrait of the dynamics of group segmentation and social influence. Behavior and social relations are heavily intertwined in the model (as in the empirical world), but we do find some predictive leverage on how an agent will act next by observing its position in sociometric space (relative to other agents in the network) and also its trajectory in that social space.

At the macrolevel, these numerical investigations illustrate a cycle of schism and social breakdown. In this region of the parameter space, a simulated group will polarize into two cohesive factions, one that works and another that shirks. There is then attrition on the sociometric margins of the worker faction, among workers who differ from others in their patterns of friendship and social approval (e.g. by building ties to shirkers). Eventually, cooperation collapses altogether and the entire group shirks. As soon as everyone agrees on this important decision, their social structure paradoxically disintegrates. During these times of universal shirking, dissatisfaction with the low level of the collective good creates a maximum force toward changing compliance behavior in the reinforcement learning dynamic. However, individual agents' explorations in cooperation are not self-reinforcing, and coordination of concerted action toward a productive solution is impossible amid the structural ambiguity.

Gradually, the local dynamics of the network lead agents to coalesce into any number of cliques, evolving according to dyadic agreements in social approval choices. One of these cliques meanders in sociometric space, differentiating itself from the bulk of other members as the network self-organizes toward a simpler "us and them" structure. Once sufficiently alienated from the rest of the shirking masses and sufficiently cohesive to act as a single decision-making unit, a clique of agents spontaneously begins working as a team. The paradox here is that negativity and wayward cliques play an important role in the emergence and stability of cooperation.

The results also have intriguing implications for theories of structural balance [10]. Although the agents in this model seek to maintain balance in their behaviors with both positive and negative referents, this tendency is not "wired in" to the relations themselves. That is, two agents feel no direct need for consistency in their mutual evaluations of a third agent. The observed tendency toward structural balance is a byproduct of dyadic agreement on behavior (homophily). This reminds us that higher-order regularities such as triadic structural balance may emerge from more elementary dyadic processes.

Also, structural balance theory predicts that the system will reach a steady state only where a group is composed of no more than two factions. While our results are consistent with the notion that balance in dyadic relations is necessary for system stability, we observe a shifting field of cliques, not two global factions. Though either unanimity or binary polarization may indeed be necessary for system stability under the balancing dynamics, much of the mutual influence in this model appears to play out in smaller cliques of positive influence. Subgroups of two, three, or four appear to form more readily and even join or exit larger factions together. Such "friends" influence one another and may grow so similar in their relations to their environment that they change behavior and alliances in synchrony without direct coordination. Even while the binary nature of the behavioral choices implies a simple polarization into two factions, these smaller cliques of positive influence can destabilize global cleavages and prevent overall convergence. Such local dynamics would be missed in an analysis of global network stability.

The model does suggest that the group structure and level of cooperation are most stable when actors' relations converge to a single dimension of differentiation ("us" and "them") and a corresponding vector of "right" and "wrong" behaviors.[8] We thus provide formal insight to functionalist theories of social deviance [16,17], which observe that deviance reinforces social solidarity in real-world groups by catalyzing rituals of condemnation. In this case, the work choice provides a salient distinction that fosters balkanization of the

[8] Of course, convergence could also occur in the opposite fashion, as the total variance in the system shrinks to zero (homogeneity), but such an outcome is extremely unlikely under the model, because a widespread jump to compliance by anything less than the full group would quickly yield a strongly reinforced clique of shirkers who free ride on others' efforts.

group. As factions are galvanized in processes of mimetic and inductive influence, behavior within factions converges toward subgroup standards. As negative relations become more salient, the positive subgroups must "choose sides" and align with global factions. This emergent structure in turn allows agents to coordinate and reinforce costly collective action that would not be feasible in a mass of isolated individuals or small cliques.

Lastly, the model has interesting implications for social identity theory [64], which posits that members of even arbitrary groups will develop a bias to favor their in-group against an out-group. This model often produces a similar effect—where agents approve of others in their own clique and disapprove of outsiders—but this is an emergent property of principles of self-organization in the network, and occurs without a higher-order cognitive capacity. In fact, these agents have no concepts to allow them to process or store information about groups or categories of actors; they never "think" beyond their dyads. A few elementary assumptions at the individual and dyad level thus generate an effect analogous to group identity and in-group bias. This may then provide a parsimonious explanation for non-instrumental discrimination against out-group members, relying on only a few general principles of interaction in dyads. In effect, this shows a distributed representation of the formation of social categories, which are not reducible to local representations in the minds of individual agents.

In closing, let us consider what we can and cannot conclude as a result of this analysis of a single simulation run. A thorough understanding of the model's behavior requires systematic manipulation of the parameters and a robust representation of the model's response surface over that parameter space. Indeed, we have engaged this task in other papers cited, but here focus on the dynamics for a single simulation run at a theoretically important point in the parameter space where we had not found a stable equilibrium result. From this view, we cannot of course infer any relationships between the parameters and outcome variables, such as the level of cooperation or overall agreement. The view here has allowed us to see dynamic behavior that was not observable in the response surface view of the model. Notably, we show that a set of well-documented microlevel dynamics (homophilous sociometric choice, social influence, and reinforcement learning) can generate a set of well-documented macrolevel regularities (structural bifurcation, discrimination between subgroups, cycles of deviance and social stability). The parsimonious micromodel here can thus inform the more sophisticated or domain-specific theories that have previously been used to explain those macrolevel patterns. Mathematical models, computational experiments, and empirical research may work in tandem to elaborate and validate the suggestive points raised here.

Acknowledgments

I gratefully acknowledge the support of the National Science Foundation, SES-0433086 and IIS-0433637. This paper reports on a model developed in collaboration with Michael Macy and Andreas Flache and has benefited from their critical feedback. Also, I thank the editor and anonymous referee for helpful advice on the revision.

References

[1] O.T. Ahtola, Toward a vector model of intentions, in: B.B. Anderson (Ed.), Advances in Consumer Research, Association for Consumer Research, Ann Arbor, MI, 1976, pp. 477–480.
[2] R. Axelrod, The dissemination of culture: a model with local convergence and global polarization, in: R. Axelrod (Ed.), The Complexity of Cooperation, Princeton University Press, Princeton, NJ, 1997, pp. 148–177.
[3] W.S. Bainbridge, Neural network models of religious belief, Sociological Perspectives 38 (1995) 483–496.
[4] E. Berscheid, Opinion change and communicator-communicatee similarity and dissimilarity, Journal of Personality and Social Psychology 4 (1966) 670–680.
[5] S. Bochner, C.A. Insko, Communicator discrepancy, source credibility, and opinion change, Journal of Personality and Social Psychology 4 (1966) 614–621.
[6] P. Bonacich, Factoring and weighting approaches to status scores and clique identification, Journal of Mathematical Sociology 2 (1972) 113–120.
[7] D. Byrne, The Attraction Paradigm, Academic Press, New York, 1971.
[8] D. Byrne, W. Griffitt, A developmental investigation of the law of attraction, Journal of Personality and Social Psychology 4 (1966) 699–702.
[9] K. Carley, A theory of group stability, American Sociological Review 56 (1991) 331–354.

[10] D. Cartwright, F. Harary, Structural balance: a generalization of Heider's theory, Psychological Review 63 (1966) 277–293.
[11] D. Cartwright, A. Zander, Power and influence in groups: an introduction, in: D. Cartwright, A. Zander (Eds.), Group Dynamics: Research and Theory, Harper & Row, New York, 1968, pp. 215–235.
[12] D. Cartwright, A. Zander, Pressures to uniformity in groups: an introduction, in: D. Cartwright, A. Zander (Eds.), Group Dynamics: Research and Theory, Harper & Row, New York, 1968, pp. 139–151.
[13] P.S. Churchland, T.J. Sejnowski, The Computational Brain, MIT Press, Cambridge, MA, 1994.
[14] J.M. Cohen, Sources of peer group homogeneity, Sociology of Education 50 (1977) 227–241.
[15] J.A. Davis, Clustering and structural balance in graphs, Human Relations 20 (1967) 181–187.
[16] E. Durkheim, The Division of Labor in Society, Free Press, Glencoe, IL, 1960.
[17] K.T. Erikson, Wayward Puritans: A Study in the Sociology of Deviance, Wiley, New York, 1966.
[18] T.J. Fararo, J. Skvoretz, Unification research programs: integrating two structural theories, American Journal of Sociology 92 (1987) 1183–1209.
[19] L. Festinger, Informal social communication, Psychological Review 57 (1950) 271–282.
[20] L. Festinger, A Theory of Cognitive Dissonance, Stanford University Press, Stanford, CA, 1957.
[21] L. Festinger, A theory of social comparison processes, Human Relations 7 (1954) 117–140.
[22] L. Festinger, E. Aranson, Arousal and reduction of dissonance in social contexts, in: D. Cartwright, A. Zander (Eds.), Group Dynamics: Research and Theory, Harper & Row, New York, 1968, pp. 125–136.
[23] L. Festinger, S. Schachter, K. Back, Social Pressures in Informal Groups, Harper, New York, 1950.
[24] J.R.P. French, B. Raven, The bases of social power, in: D. Cartwright (Ed.), Studies in Social Power, Institute for Social Research, Ann Arbor, MI, 1959.
[25] N.E. Friedkin, Structural cohesion and equivalence explanations of social homogeneity, Sociological Methods and Research 12 (1984) 235–261.
[26] J.R. Harrison, G.R. Carroll, The dynamics of cultural influence networks, Computational and Mathematical Organization Theory 8 (2002) 5–30.
[27] D.O. Hebb, The Organization of Behavior: A Neuropsychological Approach, Wiley, New York, 1949.
[28] D.D. Heckathorn, Collective sanctions and the creation of prisoner's dilemma norms, American Journal of Sociology 94 (1988) 535–562.
[29] F. Heider, The Psychology of Interpersonal Relations, Wiley Publishers, New York, 1958.
[30] I.R. Hoffman, N.R.F. Maier, An experimental reexamination of the similarity-attraction hypothesis, Journal of Personality and Social Psychology 3 (1966) 145–152.
[31] G.C. Homans, The Human Group, Harcourt, Brace, & World, New York, 1950.
[32] J.J. Hopfield, Neural networks with graded response have collective computational properties like those of two-state neurons, Proceedings of the National Academy of Sciences 81 (1984) 3088–3092.
[33] J.J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, Proceedings of the National Academy of Sciences 79 (1982) 2554–2558.
[34] J.J. Hopfield, D.W. Tank, "Neural" computation of decisions in optimization problems, Biological Cybernetics 52 (1985) 141–152.
[35] E.E. Jones, R.E. Nisbett, The actor and the observer: divergent perceptions of the causes of behavior, in: E.E. Jones et al. (Eds.), Attribution: Perceiving the Causes of Behavior, General Learning Press, Morristown, NJ, 1972, pp. 37–52.
[36] D.B. Kandel, Homophily, selection and socialization in adolescent friendships, American Journal of Sociology 84 (1978) 427–436.
[37] J.A. Kitts, M.W. Macy, A. Flache, Structural learning: attraction and conformity in task-oriented groups, Computational and Mathematical Organization Theory 5 (1999) 129–145.
[38] B. Latane, Pressures to uniformity and the evolution of cultural norms, in: D.R. Ilgen, C.L. Hulin (Eds.), Computational Modeling of Behavior in Organizations: The Third Scientific Discipline, American Psychological Association, Washington, DC, 2000, pp. 189–215.
[39] P.F. Lazarsfeld, R.K. Merton, Friendship as a social process: a substantive and methodological analysis, in: M. Berger, T. Abel, C.H. Page (Eds.), Freedom and Control in Modern Society, Van Nostrand, New York, 1954, pp. 8–66.
[40] R. Lippitt, N. Polansky, F. Redl, S. Rosen, The dynamics of power, Human Relations 5 (1952) 37–64.
[41] A.J. Lott, B.E. Lott, Group cohesiveness as interpersonal attraction: a review of relationships with antecedent and consequent variables, Psychological Bulletin 64 (1965) 259–309.
[42] M.W. Macy, Natural selection and social learning in prisoner's dilemma: coadaptation with genetic algorithms and artificial neural networks, Sociological Methods and Research 25 (1996) 103–137.
[43] M.W. Macy, J.A. Kitts, A. Flache, S. Benard, Polarization in dynamic networks: a Hopfield model of emergent structure, Dynamic Social Network Modeling and Analysis, National Academies Press, Washington, DC, 2003, pp. 162–173.
[44] N.P. Mark, Beyond individual differences: social differentiation from first principles, American Sociological Review 63 (1998) 309–330.
[45] N.P. Mark, Culture and competition: homophily and distancing explanations for cultural niches, American Sociological Review 68 (2003) 319–345.
[46] P.V. Marsden, N.E. Friedkin, Network studies of social influence, Sociological Methods and Research 22 (1993) 127–151.
[47] R. Mazen, H. Leventhal, The influence of communicator-recipient similarity upon the beliefs and behavior of pregnant women, Journal of Experimental Social Psychology 8 (1972) 289–302.
[48] J.M. McPherson, L. Smith-Lovin, J.M. Cook, Birds of a feather: homophily in social networks, Annual Review of Sociology 27 (2001) 415–444.

[49] K. Nakao, Longitudinal approach to subgroup formation: reanalysis of Newcomb's fraternity data, Social Networks 15 (1993) 109–131.
[50] T.M. Newcomb, The Acquaintance Process, Holt, Rhinehart & Winston, New York, 1961.
[51] T.M. Newcomb, The prediction of interpersonal attraction, American Psychologist 11 (1956) 575–586.
[52] T.M. Newcomb, Stabilities underlying changes in interpersonal attraction, Journal of Abnormal and Social Psychology 66 (1963) 376–386.
[53] A. Nowak, Vallacher, Computational social psychology: a neural network approach to interpersonal dynamics, in: S. Read, L. Miller (Eds.), Connectionist Models of Social Reasoning and Social Behavior, Lawrence Erlbaum, Mahwah, NJ, 1997.
[54] R. Radloff, Opinion evaluation and affiliation, Journal of Abnormal and Social Psychology 62 (1961) 578–585.
[55] E.M. Rogers, D.K. Bhowmik, Homophily–heterophily: relational concepts for communication research, Public Opinion Quarterly 34 (1970–1971) 523–538.
[56] M.E. Rosenbaum, The repulsion hypothesis: on the nondevelopment of relationships, Journal of Personality and Social Psychology 51 (1986) 1156–1166.
[57] S.F. Sampson, A Novitiate in a Period of Change: An Experimental and Case Study of Social Relationships, PhD thesis, Department of Sociology, Cornell University, 1968.
[58] E.E. Sampson, C.A. Insko, Cognitive consistency and performance in the autokinetic situation, Journal of Abnormal and Social Psychology 68 (1964) 184–192.
[59] S. Schachter, Deviation, rejection, and communication, Journal of Abnormal and Social Psychology 46 (1951) 190–207.
[60] S.H. Schwartz, R.E. Ames, Positive and negative referent others as sources of influence: a case of helping, Sociometry 40 (1977) 12–21.
[61] R. Secord, C. Backman, Social Psychology, McGraw-Hill, New York, 1964.
[62] G. Smeaton, D. Byrne, S.K. Murnen, The repulsion hypothesis revisited: similarity irrelevance or dissimilarity bias? Journal of Personality and Social Psychology 56 (1989) 54–59
[63] E. Stotland, A. Zander, T. Natsoulas, Generalization of interpersonal similarity, Journal of Abnormal and Social Psychology 62 (1961) 250–256.
[64] H. Tajfel, J.C. Turner, The social identity theory of intergroup relations, in: S. Worchel, W.G. Austin (Eds.), Psychology of Intergroup Behavior, Nelson-Hill, Chicago, IL, 1986.
[65] R. Thorndike, The Fundamentals of Learning, Teachers College, Columbia University, New York, 1932.
[66] L.M. Verbrugge, The structure of adult friendship choices, Social Forces 56 (1977) 576–597.