Comput Math Organ Theory (2007) 13: 333–353
DOI 10.1007/s10588-007-9015-x

Shape up or ship out: social networks, turnover, and organizational culture

James A. Kitts · Paul T. Trowbridge

Published online: 1 May 2007

Abstract This paper considers a formal model of cultural transmission in organizations, examining the interplay of structured social influence and organizational demography. A set of focused and fine-grained computational experiments elucidates this model's assumptions, facilitates deeper explanations for some of its behavior, and explores the robustness and scope conditions of previously published conclusions. In doing so, this investigation highlights several important issues in the design and evaluation of computational experiments.

Keywords Social networks · Turnover · Organizational culture · Organizational demography · Simulation · Social influence

Extensive research has documented the importance of organizational culture, but has made little progress in specifying the processes by which manifestations of organizational culture emerge and are maintained over time. A distinct research tradition in organizational demography has investigated recruitment and turnover, developing general models of the population dynamics of employees, managers, or other members of organizations. These ongoing dynamics of turnover make the maintenance of organizational culture problematic: When and how do cultures persist in organizations, even as the people who enact these cultures exit and are replaced by others? Addressing this question can provide some leverage for our limited understanding of the dynamics of organizational culture.

J.A. Kitts (✉)
Graduate School of Business, Columbia University, 3022 Broadway, New York, NY 10027-6902, USA
e-mail: James.Kitts@mail.gsb.columbia.edu

P.T. Trowbridge
Department of Sociology, University of Washington, 202 Savery Hall, Box 353340, Seattle, WA 98195-3340, USA
e-mail: trow@u.washington.edu

In an ambitious research program, Richard Harrison and Glenn Carroll (1991, 1998, 2001a, 2001b, 2002, 2006) have investigated the stability and convergence of organizational culture using an integrated model of organizational demography and socialization. Of course, effective socialization has implications for recruitment and turnover, while rates of recruitment and turnover surely affect the dynamics of socialization. Both processes depend in important ways on the structure of social contacts among organization members, which also change over time. Because our intuitions are ill-equipped to comprehend interacting dynamic systems, these authors use formal theory to grapple with this vexing interplay of socialization, turnover, and social networks.

In this paper we elucidate the assumptions of the Harrison-Carroll model using mathematical and graphical lenses and then systematically investigate a range of its behavior using a series of computational experiments. We are particularly interested in the counterintuitive conclusion from Harrison and Carroll (2002) that turnover increases the average strength of social ties between organization members, an outcome that we will call *cohesiveness*. This result is intriguing in light of our empirical intuition that individual social ties tend to grow stronger over time, and thus high levels of turnover that continually restart this process should diminish group cohesiveness, ceteris paribus. In considering the dynamics of cohesiveness in this model, we aim to improve our understanding of this important research program more generally, and also to demonstrate the effectiveness of using narrowly-focused and fine-grained computational experiments to identify underlying causes of model behavior.

The authors have extensively elaborated on the implications of the model not only for basic research in organizational behavior but also for management practice. Because previous work (Harrison and Carroll 2006) has provided a rich variety of empirical applications and recommendations—from negotiating corporate mergers and hiring university faculty to disrupting terrorist cells—we do not discuss external validation in this paper. Instead, we use a rigorous formal analysis of the model to offer a deeper understanding of its behavior. Our contribution thus does not speak to or depend on the adequacy of the Harrison-Carroll model as an account for real-world phenomena, which is a question to be addressed by empirical research.

1 Model

In order to examine the dynamic interplay of organizational demography and socialization, Harrison and Carroll consider a single stylized dimension of culture, or "enculturation," described as the fit of each organization member to management's ideal cultural template. They assume that individuals' enculturation scores change through in-group socialization, which molds members to fit their local organizational culture. In the original (1991) model, this socialization reflected three primary forces—an upward pull from management toward the ideal score, a downward pull due to natural decay, and a lateral pull by surrounding peers. This peer influence operated as a uniform pull toward the population mean level of enculturation. The distribution of enculturation scores also depends on turnover among members, as the scores of departing members are replaced by the scores of new recruits.

Recent work (Harrison and Carroll 2001a, 2002, 2006) has extended the model to allow that influence is not equal across members, but may vary in force according to the strength of interpersonal ties. Focusing on the dynamics of enculturation through dyadic influence and through turnover, the extended model disregards the pulls by management and decay. In this project, we follow the recent work on the extended structural model, and so leave the dynamics of management socialization and decay for later work.

Formally, a set of N organization members can be described as an $N \times 1$ vector C, with a distribution of culture over members, where C_i denotes member i's enculturation score. Dyadic influence is weighted by social ties, which vary in strength between 0 and 1. The model represents the set of all such ties as an $N \times N$ matrix S, where S_{ji} denotes the strength of member j's tie from member i (determining j's power to influence i), and $i \neq j$. An organization begins with a random initialization of both C and S across all members and these scores are reinitialized for each member who exits (representing replacement by a new recruit). For clarity, we will first describe the dynamic model, then separately detail Harrison and Carroll's assumptions about initial distributions of enculturation and influence among founding members and new recruits.

1.1 Dynamics of interpersonal influence

With the distribution of organizational culture given at the start of each simulation round (interpreted substantively as one month in calendar time), each member's score changes according to the difference equation (1). Influence on member i is conditioned by the strength of the dyadic tie from each other member $j\,(S_{ji})$:

$$\Delta C_i = a \sum_{j=1}^{N} \frac{S_{ji}(C_j - C_i)}{N-1}, \quad i \neq j. \tag{1}$$

The incremental change in C_i (or ΔC_i) in an iteration of the model is proportional to the present discrepancy between enculturation scores for i and j, so it diminishes as $(C_j - C_i)$ approaches zero. This dyadic influence is aggregated over all other members j and divided by the total number of other members $(N-1)$, so that each member is influenced by a weighted mean of all other members' enculturation scores. The parameter a determines the rate of change in enculturation, such that $a = 0$ implies perfect inertia, and thus no cultural change over time for any members. Harrison and Carroll hold a fixed at 0.05.

1.2 Dynamics of social structure

The matrix of weighted social relations, S, changes over time according to a simple rule of homophily (Lazarsfeld and Merton 1954): Members grow "closer" to (that is, come to be more strongly influenced by) others who are relatively similar to them in culture. Here, the tie strength S_{ji} is a decreasing function of the absolute cultural distance between i and j and an increasing function of i's mean distance to all other members. First, Harrison and Carroll find this mean distance (absolute deviation) between member i and all other members, as shown in (2):

$$D_i = \sum_{j=1}^{N} \frac{|C_j - C_i|}{N-1}, \quad j \neq i. \tag{2}$$

The transformation $1 - D_i$ can then be interpreted as i's mean similarity to all other members. For each dyad, a change factor (W_{ji}) is computed as the following function of the similarity between member j and member i (or $1 - |C_j - C_i|$) minus i's mean similarity to all other members ($1 - D_i$):

$$W_{ji} = (1 - |C_j - C_i|)^4 - (1 - D_i)^4, \quad j \neq i. \tag{3}$$

Both terms are raised to the fourth power, so the incremental change in the strength of influence is steeper if i and j are more similar and if i is less similar to all others. The change in S, or ΔS_{ji}, is determined as follows:

$$\Delta S_{ji} = \begin{cases} bW_{ji}(1 - S_{ji}) & \text{if } W_{ji} > 0, \\ bW_{ji}S_{ji} & \text{if } W_{ji} < 0. \end{cases} \tag{4}$$

When j is more similar to i than j is to other members on average ($W_{ji} > 0$), then the strength of the tie (S_{ji}) increases towards 1. When j is less similar to i than j is to other members on average ($W_{ji} < 0$), then S_{ji} decreases toward zero. Including the present value of S_{ji} in the equation limits the range of S_{ji} by making ΔS_{ji} approach zero as S_{ji} approaches the limits (0, 1). The parameter b determines the rate of change in S, such that $b = 0$ implies perfect inertia and thus no change in relations among members over time. Harrison and Carroll hold b fixed at 0.075.

1.3 Dynamics of membership turnover

Harrison and Carroll allow that organizational culture changes through member turnover as well as social influence. Although other papers have modeled organizational growth and decay (Harrison and Carroll 1991, 2001b), this version of the model assumes a fixed size (N). The values for C and S are reinitialized for member i (and i's ties to all others, j) when the member exits, representing immediate replacement by a new recruit.

They model an individual-level hazard rate, RD_i, as a stochastic function of alienation from the organization's culture (i's squared deviation from the group mean, $\overline{C}$) with two parameters, ER and AR:

$$RD_i = ER + AR(C_i - \overline{C})^2. \tag{5}$$

The parameter AR represents selective attrition due to individuals' alienation from organizational culture; higher values of AR make the exit process more sensitive to alienation. Members with poorer fit to the local mean—that is, members in the tails of the distribution of enculturation scores—will have a greater propensity to exit the organization when $AR > 0$. The parameter ER reflects all other factors that may influence the baseline likelihood of turnover at any given time. Just as with a and b, the parameters *ER* and *AR* may be set to zero to deactivate turnover dynamics in the model.

A member exits the organization when $-\log(U) < RD_i$, where U is a random number from a uniform distribution in (0, 1). Harrison and Carroll interpret RD_i as a hazard (i.e. the expected number of exits per unit time) instead of using a simple probability because they envision turnover as an event process occurring in continuous time. Thinking of RD_i as a hazard rate allows that RD_i may theoretically exceed 1.0 (unlike a probability) so the expected duration could then be shorter than a single round. Members should then be able to exit and be replaced in the middle of a round, interrupting influence processes at that time. The model treats time as intrinsically discrete in several other ways, however. Members' states are updated synchronously; all update their enculturation C and social relations S scores simultaneously, with respect to last round's values for all peers. Further, exiting members effectively stay until the end of the round regardless of their hazard rate of exit. This can be seen in that members who are exiting in round t exert or receive as much influence as do members who are not exiting in that round, while incoming recruits do not exert or receive any influence until the round after the member they are replacing exits. Given the low values of *RD* observed in simulations under the model, however, we do not expect that this specification substantively affected any of the results that we are interested in here.[1]

2 Mathematical analysis

Harrison and Carroll are interested in the model's behavior where all three processes are operating in tandem, so they run simulations under the integrated model. Here, we briefly consider a formal analysis of the extended portions of the model in isolation, as this will give us a deeper insight into why the integrated model behaves as it does.

Harrison and Carroll express the model as a set of recursion relations, where each dynamic variable (e.g. C, S) at time $t + 1$ is given as a function of its value at the previous time period (t) plus a set of other factors that may increase or decrease its value. First, we express these dynamics as difference equations, where the left side of the equation gives the change in the variable (e.g. ΔC, ΔS) for a given unit of time. This algebraically equivalent form makes the conditions for equilibrium more transparent. We may simply set this change to zero and solve the resulting equation to determine the conditions for a steady state in discrete time—considering this part of the model in isolation.

Examining the influence part of the model (i.e. change in culture), we solve (1) as described above for the conditions that yield a steady state in all members' enculturation scores. First, we see an obvious and trivial condition: No members' enculturation scores will change if the rate parameter $a = 0$. This feature of the model is important because it allows us to deactivate all influence in the model, forcing members to retain their initial enculturation scores, but it does not reveal an interesting steady

[1] Running simulations up to 3,000 iterations with baseline turnover and selective attrition set higher than the values that Harrison and Carroll used ($ER = 0.05$; $AR = 1.0$), the highest RD value we ever observed for any member was $\sim$0.38.

state of the model. The conditions for a steady state with $a > 0$ require some discussion. First, see that within each dyad i-j there will be zero influence from j to i if either $S_{ji} = 0$ (their social tie is null) or $C_j - C_i = 0$ (their enculturation levels are identical).

There will be a steady state for all enculturation scores in the organization if—for all i–j dyads—either $C_j - C_i = 0$ or both $S_{ji} = 0$ and $S_{ij} = 0$. We may define this case more succinctly using some terms from graph theory. Let us define a *component* of a graph as a maximal connected subgraph of nodes. A connected subgraph is a set of members who are connected to one another directly or indirectly through relations between other members. A maximal connected subgraph is one that is not subsumed by any larger connected subgraph. Notice that any non-null graph contains one or more components, including the case where the entire graph is connected (i.e. there is only one component) or the graph contains no relations at all (i.e. each node constitutes its own component).

With these terms, we may say that the influence process will reach a nontrivial equilibrium under one condition: when there is cultural homogeneity among members within every component of the S graph. If the entire social network is connected this means that equilibrium will obtain only when the entire organization is culturally homogeneous. If S is disconnected, such that two or more subgroups cannot influence one another directly or through any indirect path, then equilibrium will obtain when there is homogeneity within each distinct component of the graph. This result obtains immediately for the trivial case where no member is connected to any other. Under Harrison and Carroll's assumptions—that S begins as a uniform random matrix, the influence component of the model in (1) will typically create a force toward convergence of the group culture toward uniformity.

Now we may consider the dynamic of relational change (ΔS). Inspecting (4) reveals that any network may be a trivial equilibrium where rate parameter $b = 0$. Again, this property allows us to deactivate structural dynamics in the model, but does not reveal an interesting steady state of the model. Assuming $b > 0$, then, we see that a steady state will obtain for S within each i–j dyad if any of the following three conditions is true: (1) $W_{ji} > 0$ and $S_{ji} = 1$; (2) $W_{ji} < 0$ and $S_{ji} = 0$; (3) $W_{ji} = 0$. The first two conditions indicate that the relation will not change if forces push it farther when it is already at its limits. The case $W_{ji} = 0$ is more complicated. This condition will be true for all i–j dyads if and only if the discrepancy in enculturation scores between any two members equals the mean discrepancy for each of those members with all other members, i.e. $|C_j - C_i| = D_i$ and $|C_i - C_j| = D_j$. In other words, for the process of relational change to cease, all members must be equidistant from all other members in cultural space. Under Harrison and Carroll's representation of culture as a single dimension and group size as larger than two members, we may further say that this condition will obtain only if all members' enculturation scores are identical.

If universal cultural homogeneity is an equilibrium of the dynamic for cultural change in (1) and also is an equilibrium of the dynamic for relational change in (4) then it is a steady state of the integrated model in the absence of turnover. However, it is not unique in this respect. The socialization dynamic (ΔC) and structural dynamic (ΔS) will both exhibit fixed points if the network is bifurcated into culturally homogeneous components, each of which exerts zero influence on the others. Further, the conditions for a steady state in both cultural and structural dynamics will be approached over time for any case where S is partitioned into components that are *relatively* homogeneous: All members within each component must be closer to one another in culture than they are to the broader set of organization members, on average, and they must be more distant in culture to others outside the component than they are to the broader set of organization members, on average. Because change in social relations is deterministic, this latter condition will guarantee that ‘weak ties’ (Granovetter 1973) to members outside their clique will not emerge, and thus they will be influenced only by fellow clique members. As a result, each clique will converge toward its own subculture.

We have described Harrison and Carroll’s integrated model of socialization and demography, detailing the joint dynamics by which members’ enculturation levels change due to influence, the influence relations between members change due to cultural proximity, and the membership composition of the group changes due to turnover. Focusing on the first two (deterministic) dynamics here has allowed a simple mathematical analysis to highlight forces in the model leading toward either a unified homogeneous group culture or bifurcation into homogeneous subcultures that fail to influence one another. In the Harrison-Carroll model, these outcomes are prevented by membership turnover. Stochastic turnover continually reintroduces cultural heterogeneity into the population and also joins any disconnected subgroups (preventing segmentation into cliques). Including turnover also makes the model mathematically intractable, particularly because of selective attrition, which creates a dependence of turnover on the relationship between individual members’ enculturation score and the group mean culture. Following Harrison and Carroll, we now employ computational experiments using simulations under the model.

3 Computational experiments

3.1 Initial values: enculturation scores (C) and social relations (S)

In order to conduct computer simulations by recursively evaluating the dynamic relations (1), (4), and (5), we must specify initial conditions for all variables. We begin by describing the assumed distributions of scores over individual enculturation (C) and dyadic influence (S) in the Harrison-Carroll model. Members enter the organization with a specified mean level of culture, CM, plus a standard normal disturbance, ε, with a dispersion (standard deviation) adjustable by the parameter CS:

$$C_i = CM + CS(\varepsilon). \tag{6}$$

Harrison and Carroll use (6) to compute the initial enculturation scores for founding members of new organizations as well as the scores for new recruits following turnover events.[2] They interpret higher values of CM as indicating higher selectivity on the part of management. Values for C are capped at [0, 1], truncating this distrib-

[2]Harrison and Carroll describe a single parameter CM, which they manipulate at levels 0.3, 0.5, and 0.7 to represent different levels of recruitment selectivity, but they use two distinct parameters in the

ution at the limits of C, such that any values outside this range are stacked at 0 and 1. We investigated this tendency and did not find that stacking of values on 0 and 1 is a pervasive problem for the value of CS used by Harrison and Carroll ($CS = 0.15$), as long as $CM = 0.5$. However, greater dispersion would yield W-shaped ($CS = 0.25$) or U-shaped ($CS = 0.5$) distributions. The stacking also becomes severe when CM is not 0.5 (Harrison and Carroll 2001a, 2002, 2006). We note that this bimodal or trimodal distribution of C seems more likely to generate an equilibrium of multiple cliques, if b is high enough that relations adapt quickly to the initial clustering of enculturation scores.

To determine the structure of influence among the members, Harrison and Carroll initialize the matrix S of dyadic influence with random numbers from a specified distribution in the interval (0, 1), with zeros on the main diagonal. This implementation assumes that a member does not influence itself ($S_{jj} = 0$) and that initial influence is always positive ($0 < S_{ji} < 1$). Harrison and Carroll do not assume that influence is symmetric ($S_{ji} = S_{ij}$), though homophily will tend to lead to symmetry over time.

A crucial goal of Harrison and Carroll's extended model is to investigate the implications of a relationship between timing of entry and the structure of social influence for the dissemination of culture. They consider two scenarios, one in which social relations are distributed randomly across pairs of members and another in which the timing of organizational entries into cohorts plays an important role in determining the strength of interpersonal contacts. In the "random influence" model, Harrison and Carroll initialize S as uniform in (0, 1) and later replace dyadic ties for any exiting member with random values from the same distribution. Thus, relations connecting all types of new members (both founders and recruits) are generated by the same uniform random process. This random influence model serves as a baseline for the investigation of cohort-dependent influence.

We focus on the "cohort influence" model, which assigns lasting importance of the timing of recruitment for the calculation of initial influence levels. Equation (7) determines the initialized strength of influence between members j and i, or S_{ji}, as a function of j's tenure (t_j) and i's tenure (t_i). This function takes different forms depending on whether one or both members are *incumbents* ($t > 0$), *recruits* ($t = 0$; initialized at a turnover event), or *founders* ($t = 0$; initialized at the start of the simulation):

$$S_{ji} = \begin{cases} \max[U^2, 1 - \alpha(t_j + \varepsilon)] & \text{incumbent } j \text{ influence on recruit } i, \\ \max[U^2, 1 - \beta(t_i + \varepsilon)] & \text{recruit } j \text{ influence on incumbent } i, \\ \max[U^2, 1 - .05|\varepsilon|] & \text{recruit } j \text{ influence on recruit } i, \\ U & \text{founder } j \text{ influence on founder } i \end{cases} \tag{7}$$

for all $j \neq i$, where $\alpha = 0.01$, $\beta = 0.02$, U is a random value from a uniform distribution in (0, 1), ε is a random value from a standard normal distribution, and S is confined to the range (0, 1] by capping high values at 1.0. There is no capping of initialized S scores at zero because a uniform random variable in (0, 1) will not yield a value below zero.

experiment—one (CM) for the mean C value of the founding members (0.5) and one (SEL) for the mean C values of the replacement recruits in the different conditions (0.3, 0.5, and 0.7). We investigate only the case where $CM = SEL = 0.5$, so the distinction is unimportant here.

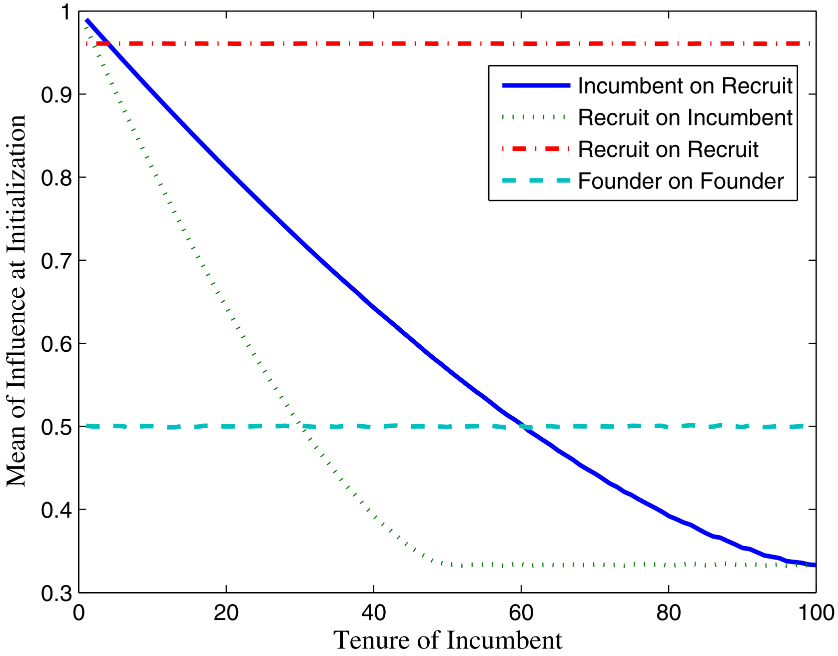

Fig. 1 Expected values for initialized social influence, S with $\alpha = 0.01$, $\beta = 0.02$

First note that incumbent-to-recruit, recruit-to-incumbent, and recruit-to-recruit influence ties are initialized as the larger of either a squared uniform random number (with an expected value of 1/3)[3] or the value of a second function of t and ε. The first term allows that there is a stochastic bottom cap with an expected value of ∼.33 for the initialization of influence between incumbents and entrants or between co-recruits, but this is not true for ties between co-founders. We will explore the implications of this assumption shortly, as we systematically compare the four initialization functions in (7).

Because Harrison and Carroll are interested in the mean and variance of S at the end of simulations, it will be instructive to consider the mean and variance of S at initialization for these various conditions. Given these numerical values, this model assumes that incumbents influence recruits more strongly than recruits influence incumbents, but that both of these influences are decreasing with the incumbent's tenure at the newcomer's arrival. Harrison and Carroll note that this "cohort influence" assumes that most of the important peer influence occurs within cohorts and that influence diminishes monotonically with increasing discrepancy in tenure. However, we will see that there is no such bonding effect for the founding cohort. Founding members typically have low influence on one another and will (on average) share greater influence with members who are recruited in the early rounds of the simulation.

Figure 1 demonstrates the assumptions in (7) by plotting the expected value of the initialized strength of influence for incumbent-to-recruit, recruit-to-incumbent,

[3]Note that a uniform distribution in (0, 1) raised to the power γ has an expected value of $\frac{1}{1+\gamma}$.

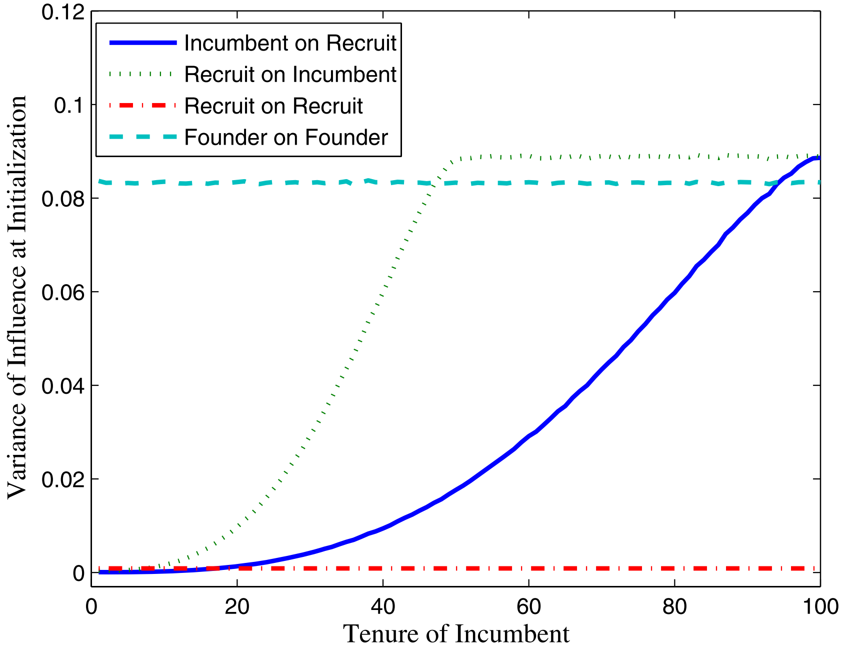

Fig. 2 Expected variances of initialized social influence, S with $\alpha = 0.01$, $\beta = 0.02$

recruit-to-recruit, and founder-to-founder. We include recruit-to-recruit and founder-to-founder strength of influence as reference lines.

Examining Fig. 1, we can summarize the following observations about the expected values of the strength of social influence ties (S) at initialization:

1. On average, incumbents have greater influence on recruits than recruits have on incumbents.
2. After the incumbent's tenure exceeds ∼50 rounds (just over 4 years), an incoming recruit's expected influence on the incumbent reaches 1/3 (the expected value of the squared uniform random variable) and remains at this level indefinitely as the incumbent's tenure increases further.
3. After the incumbent's tenure exceeds ∼100 (just over 8 years), the incumbent's expected influence on incoming recruits also reaches the same bottom limit and remains constant after that point. Thus, observation 1 above is no longer strictly true when there is a great discrepancy in tenure.
4. Recruits have a high initial mean influence on other recruits that enter at the same time (i.e. in the same "cohort").
5. However, the expected influence of recent entrants (e.g. "sophomores") on new recruits ("freshmen") is stronger than the expected influence between new recruits. Most surprisingly, freshmen influence sophomores more strongly than freshmen influence each other, on average.[4]

[4]This inversion of the cohort effect in the first few rounds of tenure does not seem to fit Harrison and Carroll's description of the model, but it is very modest and we do not suppose that it has any substantive impact on results.

6. Founding members are initialized with a relatively low (0.5) mean influence on one another.

Figure 2 shows the variance of dyadic S scores for recruits, incumbents, and founders, as initialized by (7). The figure reveals the following patterns:

7. Co-recruits have practically zero variance in their influence on one another. Combined with observation 4 above, this implies that individuals within a cohort are all close "friends" and have no significant substructures ("cliques").
8. Similarly, there is near-zero variance in influence between new recruits and recent entrants. Variance in influence does not grow steeply until the discrepancy in tenure is quite high (about 20 rounds for recruit-to-incumbent influence and much longer for incumbent-to-recruit influence).
9. Founding members have relatively high variance in their influence on each other, as they are initialized with a uniform random structure.

Harrison and Carroll choose parameter values that qualitatively fit intuitions or previous research findings in organizational demography. However, the contribution of the model does not depend so much on the realism of the particular assumptions, but on the theoretical insights derived from the model. The importance and external validity of these theoretical insights is a question for future empirical research, so we aim instead to investigate the internal logic of the model. We will use narrowly focused computational experiments to provide deeper explanations of the model's behavior, investigate the robustness of conclusions, and identify scope conditions for predictions.

3.2 Experimental design

Harrison and Carroll conduct "virtual experiments" to illuminate the potentially complex dependence of model predictions on input assumptions. The virtue of applying an experimental method to simulation is the ability to focus on a particular process, drawing inferences on the model's behavior with maximal internal validity. For example, an investigator may map a model's behavior while manipulating one or two parameters systematically and holding all others constant across conditions. Because this method removes the risk of 'spurious' inference due to confounding variables, it renders moot the use of multivariate statistical models to partial out such effects.

Harrison and Carroll hold most parameters constant in their experiments but manipulate several of them as independent variables—examining two or more levels of each—creating a set of qualitatively distinct experimental "conditions." More specifically, they assume fixed values for $a = 0.05$, $b = 0.075$, $CS = 0.15, \alpha = 0.01$, and $\beta = 0.02$, while varying parameters ER, AR, CM, N, and also manipulating the structure of influence according to the *random influence* and *cohort influence* protocol. These manipulations comprise a $3 \times 3 \times 4 \times 2$ factorial design:

- Turnover
  - None—$AR = 0$; $ER = 0$
  - Low—$AR = 0.4$; $ER = 0.02$
  - High—$AR = 0.4$; $ER = 0.04$
- Selectivity (CM) $= \{0.3, 0.5, 0.7\}$

- Size $(N) = \{5, 15, 25, 50\}$
- Influence = {*Random*, *Cohort-Based*}

Further, note that the turnover manipulation is actually a paired manipulation of two parameters of the turnover function—*ER* (base turnover) and *AR* (sensitivity to alienation in the exit process)—two forces that may have different implications for cultural convergence. Thus, Harrison and Carroll manipulate *five* independent variables in the experiment that we investigate here. See that this defines a five-dimensional space across which the model's behavior (its 'response surface') may vary. By implementing such coarse manipulations (only 2–4 levels of each independent variable), Harrison and Carroll aim to map a five-dimensional space with a set of only 72 steps. Such a broad scope may give a valuable first look at a model's range of behavior, but a narrower and deeper focus may be required to observe and account for interactions and other nonlinearities.

In making inferences about the model's response surface, we must be able to distinguish variability in model outcomes that is due to manipulations of parameter values or initial conditions from variability in outcomes that is due to stochasticity in the model's behavior. Harrison and Carroll guard against sampling error on the response surface by performing 100 simulations in each experimental condition. Performing such replicate runs allows them to derive more accurate estimates of the mean outcomes in the 72 experimental conditions, but the resulting analyses convey no information about the model's behavior in the regions of the space *not* sampled. Given constraints on computation time, researchers must strike a balance that provides sufficient observations at each sampled point to derive accurate estimates at each point along with a fine enough sampling resolution to describe the overall shape of the response surface.

Harrison and Carroll (2002) pool across the various experimental conditions and use inferential statistics (regression analysis) to investigate changes in mean outcomes as parameters are manipulated. They examine four dependent variables—*mean influence*, *variance of influence*, *influence centrality*, and *net influence range*. In pooling across so many parameter manipulations (also pooling the 'cohort influence' model with the 'random influence' model) and using OLS regression, Harrison and Carroll depart from the advantages of experimental control and aim instead for statistical control.

We focus on the findings on their first dependent variable, the "mean influence" among members, which is simply the sum of all dyadic influence strengths S (ignoring the diagonal) divided by the total number of directed relations, or $N(N-1)$. In order to distinguish this average strength of social bonds in the group (S) from the actual observed influence among members (ΔC), we refer to average tie strength as *cohesiveness*. Harrison and Carroll (2002) examine cohesiveness across twelve experimental conditions, including three levels of turnover (*no-turnover*, *low-turnover*, and *high-turnover*) and four group sizes (5, 15, 25, and 50). In the cohort influence model, they demonstrate a roughly linear increase of cohesiveness across these three levels of turnover and no effect of group size.

3.3 New experiments

We will focus on a much narrower range of conditions, providing a complementary view with new insights into the model's behavior. First, we limit our study to a single dependent variable, group cohesiveness. Second, we focus on the cohort influence model and limit experimental manipulations to only two independent variables at a time (rather than five) and we hold other parameters constant by design (rather than by estimating partial regression coefficients). This experimental control allows us to make stronger inferences about the model's dependence on our manipulations. Third, rather than comparing only 2–4 extreme values on each parameter, we perform very fine-grained manipulations and display estimated response surfaces graphically. If manipulating a parameter has no effect on the model behavior, the response surface will be flat on that dimension. Conversely, if the observed response surface demonstrates a marked trend as a parameter is manipulated, we conclude that the manipulation has affected the outcome. Note that we will not provide tests of statistical significance on those inferences. In this context, such tests can at best show that enough runs have been performed that trends on the response surface are probably not due to sampling variability on the model's stochastic response. To guard against sampling error without introducing inferential error due to misspecification of a parametric model, we simply perform a very large number of independent replications. As the number of replications of a true computational experiment becomes large, the estimated mean response surface becomes indistinguishable from the true mean response surface.

Our first experiment examines cohesiveness as we manipulate the two parameters of the turnover process, *ER* and *AR*. In order to independently investigate the roles of selective attrition and random turnover, we vary *AR* from 0 to 1.0 by fine increments and orthogonally manipulate *ER* from 0 to 0.05 by fine increments, including all three of Harrison and Carroll's turnover conditions as special cases. Specifically, we set $ER = \{0.0, 0.001, 0.002, \ldots, 0.05\}$ and set $AR = \{0.0, 0.02, 0.04, \ldots, 1.0\}$, yielding a 51×51 factorial design and an array of 2,601 unique experimental conditions. We replicate the full set of conditions 500 times, yielding 500 observations at each point in parameter space. The resulting 1,300,500 independent observations of the model's behavior allow us to map the response surface with both fine resolution and high confidence. For the first experiment, we hold group size constant in the middle of the range that Harrison and Carroll examined ($N = 25$).

Following Harrison and Carroll, we allow the simulation to iterate for 300 rounds in time (interpreted as 25 years) and compute outcome variables as the average over the final 24 iterations. Figure 3 maps cohesiveness over the ranges of parameter manipulations, where the height of the surface at each vertex indicates the average cohesiveness over the 500 replications at that point in parameter space.[5]

The levels of cohesiveness reported by Harrison and Carroll (2002) at the three points that they investigated—($AR = 0$, $ER = 0$), ($AR = 0.4$, $ER = 0.02$), ($AR = 0.4$, $ER = 0.04$), labeled in Fig. 3 as No, Lo, and Hi, respectively—are very similar to the values reported here at those same points. However, sampling only these three points had made the relationship of turnover to cohesiveness appear roughly linear in their

[5]The steps of the simulation program are fully specified as pseudocode in the online appendix http://www.soc.washington.edu/users/kitts, so they may be replicated on any platform. The original Harrison-Carroll simulation program was developed in BASIC and executed on a personal computer. Because our research design included over 10,000,000 simulations (requiring approximately 1,000 hours on a single 3 GHz CPU), we executed the simulations in MATLAB 7.0 on a Mosix cluster.

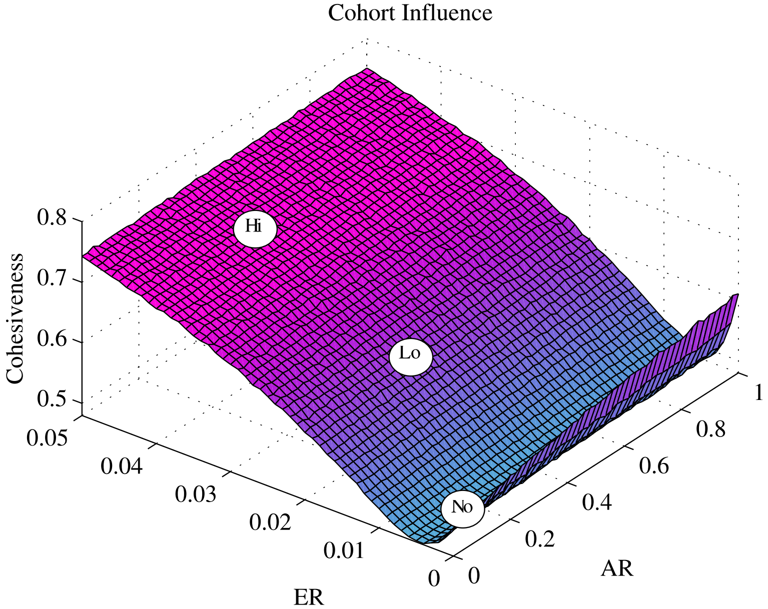

Fig. 3 Cohesiveness by base turnover (*ER*) and selective attrition (*AR*), with $N = 25$, $a = 0.5$, $b = 0.075$, $CM = 0.5$, $CS = 0.15$, $\alpha = 0.01$, $\beta = 0.02$

investigation. Our experiment affirms that these three points do indeed capture an overall positive relationship between turnover and cohesiveness in this region of the parameter space. Surprisingly, Fig. 3 based on finer manipulations suggests that the true relationship (under the model) of *ER* to cohesiveness is not only nonlinear but nonmonotonic. Increasing baseline turnover in the low-turnover region diminishes cohesiveness but increasing turnover farther increases cohesiveness.

Selective attrition (*AR*) has little effect on cohesiveness over the range examined by Harrison and Carroll. When $ER = 0$, the height of the cohesiveness surface (average tie strength, *S*) climbs modestly from ∼0.53 to ∼0.61 as *AR* increases from 0.0 to 1.0. Selective attrition has a monotonically positive effect on cohesiveness over the whole surface, but random turnover swamps this effect and makes it negligible whenever $ER > 0$. Given that *AR* has no important role either as a main effect or as an interaction with *ER* in this region, we now follow Harrison and Carroll in holding *AR* at 0.4.

Next we manipulate group size, *N*. Harrison and Carroll implemented four group sizes (5, 15, 25, 50), then pooled all cases and included (log) size in a regression analysis. We instead repeat the above *ER* manipulation over the range of *N* from 3 members to 50 members. We set $ER = \{0.0, 0.001, 0.002, \ldots, 0.05\}$ and set $N = \{3, 4, 5, \ldots, 50\}$, yielding a 51×48 factorial design and an array of 2,448 experimental conditions. We replicate the full set of conditions 500 times, giving 1,224,000 independent observations of the model's behavior. This experiment shows that group size has no main effect on cohesiveness and no interaction with *ER*, lacking even the modest trend that was apparent for *AR*. We do not show a plot of this response surface because projecting the relationship of *ER* to cohesiveness over the range of *N* does

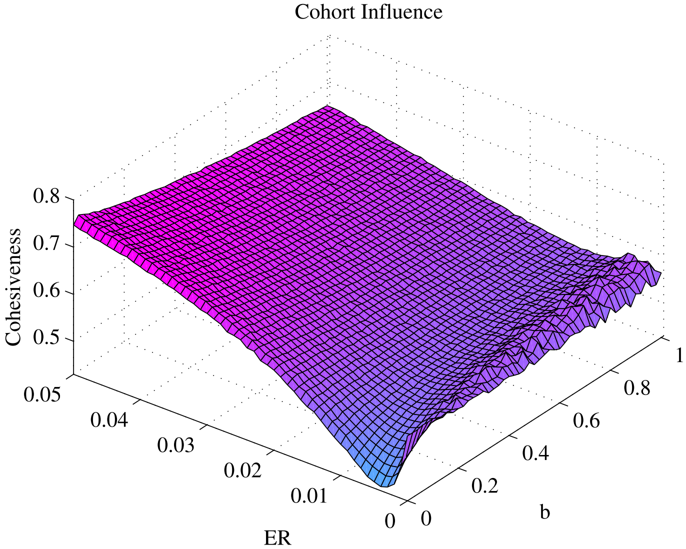

Fig. 4 Cohesiveness by base turnover (*ER*) and rate of structural change (b), with $N = 25$, $a = 0.5$, $AR = 0.4$, $CM = 0.5$, $CS = 0.15$, $\alpha = 0.01$, $\beta = 0.02$

not show us anything new. This plot is included with other supplemental figures in an online appendix http://www.soc.washington.edu/users/kitts.

Next we manipulate the parameter a, which sets the rate that culture adjusts in response to peer influence in (1). Harrison and Carroll fix this parameter at 0.05, leading culture to adjust in fairly small increments over successive iterations. We will manipulate this parameter over a much larger range and observe the model's behavior. We set $ER = \{0.0, 0.001, 0.002, \ldots, 0.05\}$ and set $a = \{0.0, 0.02, 0.04, \ldots, 1.0\}$, yielding a 51×51 factorial design and an array of 2,601 experimental conditions. We replicate the full set of conditions 500 times, giving 1,300,500 independent observations of the model's behavior.

Just as for N, varying the dynamic of social influence from very slow (low-a) to very rapid (high-a) does not alter the overall shape of the relationship between *ER* and cohesiveness. Most telling, the nonmonotonic effect of *ER* on cohesiveness obtains even for the extreme case of $a = 0$, where enculturation scores are fixed for all members over their entire membership histories. Again, we omit the $ER \times a$ response surface here, but include it in the online appendix.

Lastly, we manipulate the parameter b, which sets the rate that social relations S adjust in response to cultural similarities in (4). Harrison and Carroll set this parameter to 0.075, leading culture to adjust in small increments over successive iterations. We will manipulate this parameter using the same protocol as the $ER \times a$ experiment. Just as for the previous experiments, the same nonmonotonic relationship of turnover to cohesiveness obtains even at $b = 0$, where social relations scores are fixed at initialization for each dyad and are not allowed to change over time. However, Fig. 4 shows that increasing b to a high level does alter the effect of turnover on observed cohesiveness.

The overall positive effect of turnover remains over the entire range of b examined, but the initial dip at low turnover disappears when structural change is rapid (high b). We will revisit this finding after we have given a richer explanation of the model's behavior.

We can decompose the relationship of turnover to cohesiveness by setting parameters (ER, AR, a, and b) to zero—selectively eliminating sub-processes in the model—and observing which model components are necessary and sufficient to generate the observed pattern. The fact that we see the nonmonotonic effect of turnover on cohesiveness even when we deactivate influence dynamics by setting $a = 0$ in (1), deactivate structural dynamics by setting $b = 0$ in (4), and eliminate selective attrition by setting $AR = 0$ in (5) helps us understanding this robust nonlinear relationship. If turnover's observed effect on cohesiveness cannot be attributed to social influence, dyadic structural change, or selective attrition, the only remaining explanation is the initialization protocol of cohort influence, specified in (7).

To see how this occurs, consider the case with no change in culture ($a = 0$) or social ties ($b = 0$). With zero turnover, the average weight of social ties is then the expected value of a uniform distribution in $(0, 1)$, or 0.5, and this is constant over time because the founding members stay in the organization and ties do not change. Given a sufficiently long time horizon, even an infinitesimally low turnover rate will eliminate the founding members; given that each new recruit will arrive after a lengthy delay, the recruit's social influence ties in both directions (from the recruit to incumbents and vice versa) will be initialized at very low values in (7) and thus cohesiveness will be low overall. Increasing the turnover rate will then shorten the wait between recruits and thus increase the initialized influence between recruits and incumbents in both directions. A higher turnover rate also increases the probability of multiple turnover events within a single round (creating a strong cohort bonding effect). Both outcomes increase cohesiveness according to (7), as suggested by Fig. 1. It is thus apparent that, in the absence of structural change or cultural influence, the protocol for initializing influence ties (S) results in a drop in cohesiveness as turnover increases from zero to a negligible positive value, followed by an increase in cohesiveness as turnover rises farther.

Now we may consider the conditions in which this pattern disappears in Fig. 4, specifically where parameter b is much higher than it was in Harrison and Carroll's (2002) cohort influence experiments. Note that increasing the b parameter allows homophily to serve as a competing determinant of network structure (S). When structural change is slow (low b), the initialization protocol drives the relationship between turnover and cohesiveness. Allowing the strength of ties to change more quickly over time attenuates this pattern, as the initialized values for the strength of social ties have less impact on the observed structure at the end of the observation period.

When $ER = 0$, increasing b into a much higher range typically leads networks to bifurcate into cultural factions that have negligible on influence on one another. This partitioning of S into disconnected subgroups occurs because social ties fade to zero faster than conformity leads individuals' culture to converge toward homogeneity. This network segmentation in turn diminishes overall cohesiveness. Turnover ($ER > 0$) spans those factions with cross-cutting ties and leads the network to coalesce. Even a tiny level of turnover is sufficient to eventually yield this result when structural change is rapid (high-b). The result of this rapid adjustment in social relations is a marked increase in observed cohesiveness when ER exceeds zero.

4 Discussion

Our replication of Harrison and Carroll's model has focused on their observed relationship of turnover and cohesiveness, or the average strength of social influence ties (S). We focus on this part of their investigation not only because it is an important and interesting theoretical relationship, but because it serves as a helpful showcase for a number of points about the methodology of computational experiments.

Our highly-controlled replication with fine-grained manipulations has supported their conclusion that baseline turnover (ER) increases cohesiveness under the model. We have further demonstrated that this relationship between ER and in-group influence is more interesting than previously reported, as it includes robust nonlinearities—a nonmonotonic shape and also an interaction with the rate of structural change (b).

Investigations of the cohort influence model so far have revealed that the observed nonmonotonic relationship of ER to cohesiveness is robust as we vary the AR, N, and a, parameters. Notably, we have shown that it obtains even when we eliminate selective attrition by setting $AR = 0$, eliminate dynamic social influence by setting $a = 0$, and eliminate structural dynamics by setting $b = 0$ in the model. In doing so, we have discovered that the relationship of ER to cohesiveness in the cohort-influence condition reflects the distribution of new members and the initialization protocols for cohort influence given in (7). This does not mean their result is an "artifact" of the initialization, only that it does not result from the dynamics of selective attrition, socialization, or structural change (adjustment of ties) in the model. If this complex relationship of turnover to cohesiveness is supported empirically (an issue that we will not address), then our experiments have shown that the random turnover component of Harrison and Carroll's cohort influence model is sufficient to account for the pattern.

We have also shown that this relationship of turnover to cohesiveness is sensitive to the relative rate of structural change (b): High rates of structural change increase observed cohesiveness markedly when turnover is a very low positive value. Dynamic adjustment of ties (increasing cohesiveness as organizational culture converges) solves the problem of weak networks due to sparse entry timing, while the turnover itself solves the problem of network bifurcation.

4.1 Directions for future research

Any fruitful modeling project will yield a host of unanswered questions that guide an ongoing research agenda. Our deep and narrow focus on a particular finding in Harrison and Carroll (2002) has illuminated many features of the overall model and has given us an opportunity to discuss important issues of design for computational experiments. We uncovered more intriguing relationships in this set of fine-grained experiments than we were able to discuss here. The online appendix http://www.soc.washington.edu/users/kitts includes response surfaces for all four of Harrison and Carroll's dependent variables (*mean influence*, *influence variance*, *influence centrality*, and *net influence range*) from these same experiments. It also includes companion figures generated under the *random influence* protocol, for comparison to the cohort-influence experiment.

An important methodological question that we have not addressed yet is dynamic stability. Although earlier papers in Harrison and Carroll's research program focused on a simpler model that reportedly reached equilibrium within a practical period of time, their extended model does not stabilize within the observation period over most of the parameter space we investigated. Would results differ substantively if we let the simulations iterate longer? This is a particular concern because we are manipulating rate parameters: If there is an observed correlation (across experimental conditions) between a manipulated rate parameter and observed cohesiveness at the end of the simulation, this may indeed show that increasing the parameter yields greater (or lesser) cohesiveness in the long run. But if the response surface is changing over time, then an early snapshot of the response surface may not represent the model's later behavior. It may be that the parameter has no effect on the stable outcome (i.e. the stable response surface is flat), but that the parameter affects the waiting time until the outcome is realized. In this case, apparent patterns in the early response surface may be spurious.

For example, results in the cohort influence model that are due to the relative preponderance of founding members versus later recruits should weaken as time passes, because founding members will eventually be replaced under any nonzero turnover. In general, any observed patterns in the very low turnover region may be sensitive to the width of the observation window. In order to assess the robustness of our reported patterns to questions of dynamic stability, we replicated the experiments here over 3,000 iterations instead of 300 iterations.[6] The results we have reported in the $ER \times AR$, $ER \times N$, and $ER \times a$ experiments are similar over this longer window, although the initial drop in the cohesiveness surface as ER rises from zero becomes very steep because any positive level of turnover typically leads to loss of all founding members during the longer observation period.

Over the 3,000 round observation window, the results reported in the $ER \times b$ experiment were similar for $ER = 0$ and for almost the entire positive range of turnover, but the very-low-turnover region of Fig. 4 exhibits different patterns. Cohesiveness in the region where $b = 0$ and ER is very low falls even lower over a long period as more founding members are replaced by sporadic recruits who have weak ties to incumbents and have no way to build more positive ties over time. Near the value of b used in the Harrison-Carroll experiments ($b = 0.075$), results for 3,000 iterations are similar to the results with 300 iterations. Where b is much higher, however, the boost in cohesiveness at very low levels of turnover becomes much steeper over a longer time horizon, approaching a stepfunction as ER exceeds zero. Above this jump in cohesiveness, the relationship of ER to cohesiveness becomes similar to the one that we have observed in all of the other experiments. That is, the typical nonmonotonic effect of ER on cohesiveness actually appears in the high range of b; it just needs a longer time to be observable.

[6]Researchers may be interested in a model's behavior over a particular time horizon rather than its long-term behavior. Given Harrison and Carroll's interpretation of iterations as calendar months, the time span of 3,000 iterations (250 years) is of little substantive interest. However, running the model longer has helped us understand the dynamic processes underlying the response surfaces observed on a shorter time scale.

A reviewer asks about the possibility of periodic orbits, where the model revisits a closed cycle of states with a fixed period. Indeed, under extremely restrictive conditions ($a = 1$, $b = 0$, $ER = 0$, $AR = 0$) this outcome is theoretically possible; for example, we may observe a period-two orbit where two nodes parrot each other's enculturation states or a period-m orbit where m nodes pass enculturation states around a ring of influence.[7] In addition to the above parameter restrictions, such model behavior can only occur when S contains particular structures. Specifically, a graph containing at least one strongly connected component of two nodes (having mutual ties of 1.0) or a strictly directed cycle of more than two nodes (having a chain of ties of strength 1.0) will support an indefinite chain of influence with periodic orbits.

Notably, the model does not allow for heterogeneity among members in their baseline propensity to exit (ER) or their sensitivity to alienation (AR), as these are fixed system-level parameters. If the model had allowed for heterogeneity among members in these parameters, we might expect robust temporal patterns as a result (Kitts 2007): For example, heterogeneity in ER across members would lead to an asymptotic decrease in turnover over time, due to differential attrition among high-ER members. Where there is heterogeneity in AR, we should expect selection against high-AR members on the tails of the C distribution. Members who fit the group poorly (are distant from the mean culture) and are highly sensitive to alienation should exit, whereas members who fit well or who have low-AR should be more likely to stay. As a result, the variance in C will diminish over time especially among members with high values of AR, while the mean AR value will fall over time especially among members with values of C far from the mean.

Our use of three-dimensional surface plots to present results provides a much finer picture of the relationship of input parameters to resulting model behavior than we could obtain from a traditional coarse-grained experiment. With a sufficient number of replications, such plots effectively describe the mean response surface where we are interested in the main effects and an interaction of two independent variables.[8] We have reported such an analysis for four pairs of independent variables ($ER \times AR$, $ER \times N$, $ER \times a$, $ER \times b$) but this method does not consider interactions that are not explicitly included in the experimental design. For example, we may wonder if a faster rate for structural dynamics (b) would alter the effect of one of the other parameters (such as AR or a) on cohesiveness; we cannot address the question of $AR \times b$ or $a \times b$ interactions without running more simulations and manipulating those parameters orthogonally. This repeated observation of two-dimensional slices of the model's behavior is a narrow but rigorous way of exploring the range of a model's behavior. We may investigate three-way interactions (e.g. $ER \times a \times b$) by running a full $ER \times a$ experiment at several levels of b, and may visualize the results by animating the surface plot as b is manipulated.[9] Such an experiment would be a promising extension to the present study, but higher order interactions become very

[7]Here 'node' may refer to individual members or sets of members that are culturally and structurally equivalent.

[8]This particular visualization is obviously inappropriate if we are interested in something other than the mean model response, such as the variability in results, the distribution of outcomes over multiple equilibria, or the dynamic stability of observations.

[9]See Kitts (2006) for examples of visualizations of three-way interactions in computational experiments.

computationally demanding with fine-grained manipulations such as those we use here.

Lastly, future work should bring this research program into closer dialog with alternative formal models of the spatial and relational diffusion of culture (Axelrod 1997; Carley 1991; Epstein and Axtell 1996; Fararo and Butts 1999; Krackhardt 2001; Latane 2000; Macy et al. 2003; Mark 1998). Although the goals of these various modeling programs may differ, such exercises in model alignment (Axtell et al. 1996) may augment each research program independently as well as contribute to integration of scholarship. It would be particularly fruitful to investigate parallels between these abstract enterprises and basic models of network influence theory (Friedkin and Johnson 1990) and dynamic social impact theory (Latané 1996), which have been most grounded in empirical research.

Acknowledgements The first author gratefully acknowledges the support of the National Science Foundation, SES-0433086. Richard Callahan, Richard Harrison, Glenn Carroll, Carter Butts, and anonymous referees provided helpful input for this project.

References

Axelrod R (1997) The dissemination of culture: a model with local convergence and global polarization. In: Axelrod R (ed) The complexity of cooperation. Princeton University Press, Princeton, pp 148–177
Axtell R, Axelrod R, Epstein JM, Cohen MD (1996) Aligning simulation models: a case study and results. Comput Math Organ Theory 1:123–141
Carley K (1991) A theory of group stability. Am Sociol Rev 56:331–354
Carroll GR, Harrison JR (1998) Organizational demography and culture: insights from a formal model and simulation. Adm Sci Q 43:637–667
Epstein JM, Axtell R (1996) Growing artificial societies: social science from the bottom up. MIT Press, Cambridge
Fararo TJ, Butts CT (1999) Advances in generative structuralism: structured agency and multilevel dynamics. J Math Sociol 24:1–65
Friedkin NE, Johnsen EC (1990) Social influence and opinions. J Math Sociol 15:193–206
Granovetter MS (1973) The strength of weak ties. Am J Sociol 78:1360–1380
Harrison JR, Carroll GR (1991) Keeping the faith: a model of cultural transmission in formal organizations. Adm Sci Q 36:552–582
Harrison JR, Carroll GR (2001a) Modeling organizational culture: demography and influence networks. In: Cooper CL, Cartwright S, Earley PC (eds) The international handbook of organizational culture and climate. Wiley, New York, pp 185–216
Harrison JR, Carroll GR (2001b) Modeling culture in organizations: formulation and extension to ecological issues. In Lomi A, Larsen ER (eds) Dynamics of organizations: computational modeling and organization theories. AAAI Press, Menlo Park, pp 37–62
Harrison JR, Carroll GR (2002) The dynamics of cultural influence networks. Comput Math Organ Theory 8:5–30
Harrison JR, Carroll GR (2006) Culture and demography in organizations. Princeton University Press, Princeton
Kitts JA (2006) Collective action, rival incentives, and the emergence of antisocial norms. Am Sociol Rev 71:235–259
Kitts JA (2007) Bailing out: an elementary model of disaffiliation. Unpublished
Krackhardt D (2001) Viscosity models and the diffusion of controversial innovations. In: Lomi A, Larsen R (eds) Dynamics of organizations. MIT Press, pp 243–268
Latané B (1996) Dynamic social impact: the creation of culture by communication. J Commun 46:13–25
Latane B (2000) Pressures to uniformity and the evolution of cultural norms. In: Ilgen DR, Hulin CL (eds) Computational modeling of behavior in organizations: the third scientific discipline, pp 189–215
Lazarsfeld PF, Merton RK (1954) Friendship as a social process: a substantive and methodological analysis. In: Berger M, Abel T, Page CH (eds) Freedom and control in modern society. Van Nostrand, New York, pp 8–66

Macy MW, Kitts JA, Flache A (2003) Polarization in dynamic networks: a Hopfield model of emergent structure. In: Dynamic social network modeling and analysis. National Academies, Washington
Mark N (1998) Beyond individual differences: social differentiation from first principles. Am Sociol Rev 63:309–330

James A. Kitts is an Assistant Professor in the Graduate School of Business at Columbia University. Having earned his Ph.D. in Sociology from Cornell University, he is broadly interested in the dynamics of cooperation and competition among organizations and among their members. He has recently studied the emergence of norms in work teams, polarization and factionalism in social networks, and the demography and ecology of social movement organizations.

Paul T. Trowbridge is a graduate student in Sociology at the University of Washington. He presently studies social networks and coalitions among social movement organizations.