Title
Theory development with agent-based models

Journal
Organizational Psychology Review, 5(4)

ISSN
2041-3866

Authors
Smaldino, PE
Calanchini, J
Pickett, CL

Publication Date
2015

DOI
10.1177/2041386614546944

Peer reviewed

Theory Development with Agent-Based Models

Paul E. Smaldino
Johns Hopkins University

Jimmy Calanchini
University of California, Davis

Cynthia L. Pickett
University of California, Davis

In press at *Organizational Psychology Review*

Abstract

Many social phenomena do not result solely from intentional actions by isolated individuals, but rather emerge as the result of repeated interactions among multiple individuals over time. However, such phenomena are often poorly captured by traditional empirical techniques. Moreover, complex adaptive systems are insufficiently described by verbal models. In this paper, we discuss how organizational psychologists and group dynamics researchers may benefit from the adoption of formal modeling, particularly *agent-based modeling*, for developing and testing richer theories. Agent-based modeling is well-suited to capture multi-level dynamic processes and offers superior precision to verbal models. As an example, we present a model on social identity dynamics used to test the predictions of Brewer's (1991) optimal distinctiveness theory, and discuss how the model extends the theory and produces novel research questions. We close with a general discussion on theory development using agent-based models.

Keywords: group dynamics, agent-based modeling, formal models, optimal distinctiveness, social identity

Introduction

"A mathematical theory may be regarded as a kind of scaffolding within which a reasonably secure theory expressible in words may be built up. ... Without such a scaffolding verbal arguments are insecure." – J.B.S. Haldane (1964)

There are many challenges that face group dynamics researchers. Group dynamics are inherently complex phenomena that emerge through repeated interactions, and may include recursive relationships among numerous variables at multiple levels of analysis (Smaldino, 2014a). The human mind, as powerful as it is, faces limitations in conceptualizing complex systems (Resnick, 1994), and the methodologies commonly applied to the study of groups are constrained in their ability to fully capture the dynamic elements of group behaviors (Cronin, Weingart, & Todrova, 2011). Cross-sectional studies, for example, present static snapshots of phenomena with coarse-grained temporal resolution. Longitudinal studies may incorporate an explicit temporal dimension but often fail to account for other dynamic elements of groups, such as feedback loops and more fine-grained temporal interactions. It has been suggested, not entirely in jest, that group dynamics research is often in reality the study of group statics (McGrath, 1986).

This paper argues that many of these challenges can be addressed through the use of *agent-based models* (ABMs). ABMs constitute a subclass of *formal models*, all of which use mathematical equations or computational algorithms to precisely specify the elements of a system, the relationships between those elements, and the subsequent dynamics of the states of and relationships among those elements. Other formal modeling approaches include coupled differential equations, game theoretic analyses, cellular automata, and multinomial processing trees. Each of these has unique strengths and limitations, but a thorough review of the universe of modeling approaches is beyond the scope of this paper. Instead, we focus specifically on agent-based models and their utility for theory development within the domain of group dynamics. Broadly, an ABM is a computational representation of a social system in which the individual actors are explicitly represented, and their operations and interactions are simulated to generate and (potentially) predict complex individual-level and group-level phenomena. ABMs operate on the principles that simple behavioral rules can generate complex behavior and that the whole may be greater than the sum of its parts.

We are not the first to advocate for formal modeling in the context of group dynamics. Indeed, many researchers have called upon the field to adopt formal modeling (e.g., Cronin et al., 2011; Harrison, Lin, Carroll, & Carley, 2007; Hughes, Clegg, Robinson, & Crowder, 2012; McGrath, Arrow, & Berdahl, 2000; Weinhardt & Vancouver, 2012). However, this paper differs from other discussions of formal modeling in organizational psychology in several important ways. Previous papers have generally discussed formal modeling as a way to more concretely specify theory (e.g., Weinhardt & Vancouver, 2012). Implicit in such discussions is the assumption that theory *development* should transpire as it has traditionally done so, through verbal reasoning and empirical analysis. In contrast, we argue here that agent-based models' utility for theory development merits their addition to the organizational psychologist's toolkit as a complement to (but not replacement for) other analytic techniques. Moreover, previous work has only discussed agent-based models as part of a broad overview of many formal modeling techniques available to group dynamics researchers (e.g., Cronin & Weingart,

2011; McGrath et al., 2000; Weinhardt & Vancouver, 2012). Because we focus specifically on agent-based modeling, we present a more detailed discussion of their ability to represent key elements of group dynamics, such as organizational structure and individual heterogeneity. Additionally, we extend previous work by describing a concrete example of an agent-based model resulting from a formalization of optimal distinctiveness theory (Brewer, 1991). Using this example as a framework, we discuss implications of agent-based modeling for theory development in group dynamics. Finally, we discuss how agent-based modeling may be paired with empirical analysis in order to lend increased precision in both theory development and hypothesis testing.
Understanding the processes of group dynamics and social organization involves a several challenges. Our aim in this paper is to show how agent-based models are a valuable tool in surmounting them.

Dynamic Challenges of Group Dynamics Research

There are at least three major challenges faced by group dynamics researchers, each of which can be at least partly addressed through the adoption of ABMs. The first of these is that group dynamics involve multiple levels of organization: individual, group, and structural/situational. Individual-level processes – cognition, perception, and behavior – affect the way individuals interact with one another in group contexts, and these interactions between individuals can create emergent group-level behaviors, social norms, and institutions (Smaldino, 2014a). Meanwhile, structural and situational factors characterize the contexts in which individuals and groups operate, as well as the behaviors and social norms that constrain individual- and group-level processes. Many research approaches focus on only one level of organization, and therefore fail to capture complex interactions between these levels that may give rise to phenomena which are intractable from a single-level perspective (Wimsatt, 1974). In contrast, agent-based models can easily capture multiple levels of organization.

Another challenge to studying dynamic processes is the presence of feedback between and within the levels of organization at which groups operate (Ilgen, Hollenbeck, Johnson, & Jundt, 2005; Marks, Zaccaro, & Mathieu, 2000). One individual's output becomes another individual's input, whose output in turn becomes input for the first individual and/or others. Between levels of organization, an individual's behavior can influence group-level processes, which in turn feed back as stimuli for the individual and her social partners. Conversely, group properties can serve as stimuli for individual behaviors, a process often called 'downward causation' (Campbell, 1974; Smaldino, 2014b). There are myriad permutations of such multi-directional causal pathways in a dynamic system, which can give rise to hard-to-predict behavioral cascades. Such recursivity is difficult to capture with many commonly used empirical methods. For example, researchers who are interested in group processes often take the individual out of the group context entirely by studying participants in isolation (e.g., in cubicles) and then try to recreate a group context using implied or imagined others, thereby eliminating feedback effects. A major benefit of agent-based models is their ability to account for feedback, and even to specify conditions for the presence or absence of cascades (Schelling 1978; Bak 1996).

A third challenge group dynamics researchers face is in modeling a system with memory. Individuals in groups do not exist in timeless, isolated psychological and environmental states. Rather, current states are functions of previous states plus any other input received from the environment in the interim. Memory is the dependence of a current state on previous states. In this context, memory does not necessarily refer to individual brain states, but rather to the general ability of a system to retain information about its previous states. Groups can also have memory: collective knowledge can be stored at the group level in forms like norms, traditions, or manuals (Wilson, Goodman, & Cronin, 2007). Indeed, the ability of human cultures to build upon their collective knowledge is the foundation of complex culture, and a marvel of sociobiological computation (Smaldino & Richerson, 2013). Memory complicates the study of group dynamics because many group processes are path dependent (Cronin et al., 2011). As such, a given phenomenon can have different downstream effects depending on when it happens and what has happened before it. Agent-based models can be used to study the sensitivity of a system to path dependence, and to explore perturbations to social organization at various stages of group processes.

Limitations of Verbal Models

In the usual course of affairs, social and behavioral scientists understand a system through the use of verbal (or mental) models (Johnson-Laird, 1983). In other words, the scientist has an idea, described in words, of the relevant parts of the system and how they fit together to translate inputs and starting states into outputs and end states. Sometimes someone will propose a new verbal model, which will become the basis for new empirical tests. This approach has served the psychological sciences well for decades. However, as we become interested in increasingly complex phenomena, our ability to verbally conceptualize a working model for the behavior of the system may reach a limit. The conceptualization and communication of a complex system using words alone is often insufficient because humans are often unable to see how complex group-level patterns can emerge from individual behaviors (Resnick, 1994). Moreover, verbal models are often difficult to falsify, because the flexibility of language allows for many interpretations of a given phrase (Popper, 1963). Faced with such limitations, researchers studying the psychology and behavior of groups are increasingly turning to tools and ideas from the fields of dynamical systems (Vallacher & Nowak, 1997) and complex adaptive systems (Miller & Page, 2007), in which the use formal models is a standard technique.

Formal Models

Although the ultimate focus of this paper is on agent-based models, it is important to recognize that ABMs are a type of formal model. Because organizational psychology has been slow to adopt many modeling techniques, we will first attempt to convey more generally the broad advantages for formal modeling. Formal models work by specifying idealized and simplified pictures of the part of the world under consideration. Mathematical analysis or computational simulations can then be used to examine the consequences of those specifications. In reference to their ability to help us understand complex systems, the Nobel prize-winning physicist and complexity science pioneer Murray Gell-Mann has suggested that formal models serve as “prostheses for the imagination” (Brockman, 1995). Formal models can help yield insight into complex processes by distilling down the essence of hypotheses, illuminating boundary conditions for theories, and eschewing the vagueness of language in favor of a precise, mathematical specification of theoretical details (Bedau, 1999; Epstein, 2008; Schank, 2001; Wimsatt, 1987). A frequent and fortuitous consequence is that the explicit articulation of previously implicit or unexamined assumptions may reveal processes or outcomes unrecognized under less specific formulations.

As many models of psychological phenomena are quite complex, we will illustrate with a classic example from ecology: the Lotka-Volterra model of predator-prey dynamics (Brauer & Castillo-Chávez, 2001). This model specifies two animal species: a prey species with a positive rate of growth in the absence of predators, and a predator species with a negative growth rate in the absence of prey. The number of predators negatively influences the number of prey, and the number of prey animals positively influences the number of predators. These relationships can be specified as coupled differential equations, whose solutions represent the resulting dynamics of the two species populations. Depending on the parameters denoting the strength of the various relationships, the two populations may crash, oscillate continuously, or settle into a stable equilibrium. Although the basic model ignores complexities like spatial structure, climate, third-party species, and life history, it can still yield insight into the processes of inter-species competition and provide a basis for more predictive models.

Although formal models are not a substitute for careful empirical research, they provide a crucial complement to more commonly used methodologies. This relationship is well recognized in many fields, and mathematical and computational models have become *de rigueur* in disciplines that study social behavior, including economics, political science, ecology, sociology, and evolutionary anthropology. Yet, despite the ability of formal models to provide a more rigorous theoretical foundation for dynamic group processes, the field of organizational behavior has been slow to embrace their potential (for notable exceptions, see Berdahl, 1998; Crowder, Robinson, Hughes, & Sim, 2012; Fang, Lee, & Schilling, 2010; MacCoun, 2012; Anderson & Lewis, 2013). Although other researchers have also called upon the field to adopt formal modeling (e.g., Cronin et al., 2011; Harrison, Lin, Carroll, & Carley, 2007; Hughes, Clegg, Robinson, & Crowder, 2012; McGrath, Arrow, & Berdahl, 2000; Weinhardt & Vancouver, 2012), formal models in general remain relatively scarce in the group dynamics literature. This scarcity may in part reflect a hesitance to simplify the details of dynamic processes into a system of equations. Many models, such as those found in game theory, assume fixed or simple contingent behavior patterns (e.g., tit-for-tat) to the exclusion of a rich spectrum of other possible responses and processes, and analytic models based on coupled differential equations are necessarily limited in their ability to capture environmental complexity and individual heterogeneity. However, advancements in modeling techniques, software development, and computing power over the most recent decades now allow complex behaviors, perception, and cognition to be built into formal models.

Agent-Based Models

In the psychological sciences, relationships of interest are often too complex for straightforward mathematical analysis, and may require a computational specification and analysis through simulation. This is an advantage of ABMs, as their analysis concerns the data generated from iterated simulations of agent interactions, which can encompass quite complex relationships. Some well-known examples include artificial neural network models of cognition (in which the neurons are the agents; e.g., Hopfield, 1982; McClelland et al. 1986), Kalick and Hamilton’s (1986) model of mate preferences, and Nowak, Szamrej, and Latané’s (1990) model of dynamic social impact.

In contrast to most traditional research methods, ABMs are expressly designed to accommodate the three challenges to studying group dynamics outlined above. ABMs are explicitly multi-dimensional and can efficiently and simultaneously model dynamics at multiple levels of organization while incorporating feedback processes and system memory. ABMs are therefore especially well-suited to study complex systems whose dynamics are inherently dependent on heterogeneous actors and organizational structure, which are difficult to model using other formal techniques such as differential equations (Smith & Conrey, 2007). Users can specify rules that describe individual agent behavior, agent-agent interaction, and agent-environment interaction. Various programs and toolkits have been developed to facilitate constructing ABMs. Prominent examples include NetLogo (Wilensky, 1999) and the MASON simulation library (Luke, Cioffi-Revilla, Sullivan, & Balan, 2005).

An ABM is typically comprised of a population of agents in an environment. Each agent is discrete and autonomous, pursuing goals based on local, limited information. Agents can learn from, have memory of, or otherwise be affected by interactions with other agents and the environment. Behavior rules are postulated theoretically and implemented symbolically as a series of computational algorithms. Once initial parameters are specified, the model can run for a given amount of time or until some desired state (such as an equilibrium) has been reached. The end result is the pattern of outcomes that emerge in the population over time. Of primary interest to the study of group dynamics, outcomes can arise from interdependent processes, such that each agent is affected not only by its own decisions, but also by environmental conditions, group-level processes, and the behaviors and responses of other agents in the environment. This is a true strength of ABMs: few other methodologies are as well-suited to represent and capture multi-level, interdependent, recursive processes. Additionally, ABMs can generate numerical data like other quantitative research methods and, if suitable, they can allow users to visualize the movement of individual agents and the patterns that these movements produce. In such cases, agents literally move and change onscreen in real time. Put another way, if other research methods take static snapshots of group dynamics, ABMs create motion pictures. (For more detailed discussions on the nature of ABMs in other disciplines, see Epstein, 1999; Gilbert & Terna, 2000; Schank, 2001; Macy & Willer, 2002; Smith & Conrey, 2007).

Some people may be concerned that explicitly articulating every detail of a model amounts to “setting yourself up for success.” That is, if everything is set up according to the specifications of the theory, the only possible outcome is the one predicted by the theory. However, it is important to remember that individual-level behavioral rules do not necessarily translate directly into simple predictions of outcomes at the group level, because structural organization, feedback, and system memory all affect group-level processes. One of the strengths of ABMs is the ability to directly incorporate those properties of complex social systems, creating the potential for unexpected outcomes to arise despite complete knowledge of initial conditions (Epstein, 1999).

ABMs are especially well-suited to the study of group dynamics for a variety of reasons. They can simultaneously represent multiple levels of organization, and can easily accommodate recursion and system memory. We will now introduce an ABM from our own research on social identity dynamics to demonstrate how ABMs can be used to develop theory. We will give careful attention to the processes of developing and analyzing the model, in order to provide a potential guide for future model development.

Modeling Social Identity Dynamics

Social identity is an important component in the psychological study of groups. Specifically, individuals' social identities help to balance opposing needs to belong and to be distinct (Brewer, 1991; Brewer & Pickett, 2002; Vignoles, 2011). Individuals' feeling of comfort in a social context – that is, achieving the right balance of belongingness and distinctiveness – depends not only on their own social identities, but also on the identities and behaviors of those in their environment. The dynamics of social identity are directly relevant to organizational psychology. For example, when social identity is salient and high performance is perceived to benefit the organization, social identification is positively related to task performance and work motivation (van Knippenberg, 2000). Additionally, Shepherd and Haynie (2009) demonstrated that one facet of social identity, the need for distinctiveness, was sensitive to economically-imposed pressures and led entrepreneurs to experience diminished psychological well-being as a result of insufficient feelings of belonging. (For a review of social identity processes in organizational contexts, see Hogg & Terry, 2001.)

Optimal Distinctiveness Theory

Within psychology, there has been a long-standing interest in understanding processes of group formation and change. Group living has been a hallmark feature of human societies throughout our evolutionary history, and the formation of groups is essential to human survival in most ecological contexts. In addition to providing "safety in numbers," groups also provide a powerful organizational system that can confer great survival advantages. Social categorization and group boundaries help define the limits of cooperative interdependence. Group members can provide aid to mutually acknowledged ingroup members and trust that these ingroup members will reciprocate. In other words, groups can be thought of as bounded communities of mutual trust and obligation (Brewer, 2007; Brewer & Caporael, 2006).

Because group membership involves reciprocal obligation, it has been argued that individuals have developed a preference for groups that are moderately-sized, such that the number of individuals to which one is obligated is manageable (Brewer & Caporael, 2006): a group that is too small is limited in its ability to generate synergistic benefits, but a group that is too big is vulnerable to freeloading members. As such, there are both lower and upper limits on ideal group size. At the psychological level, *optimal distinctiveness theory* (ODT; Brewer, 1991) proposes that these ecological pressures – to belong to groups and at the same time limit their size – have resulted in two oppositional psychological motives: the need for inclusion and the need for differentiation. According to ODT, when these two motives operate together, they are hypothesized to give rise to a preference for "optimally distinct" identities, i.e., membership in groups that are moderately-sized in relation to the population in which the group resides.

Support for the existence of these motives and a preference for optimal distinctiveness has been found across numerous empirical studies (see Leonardelli, Pickett, & Brewer, 2010 for a review). These studies have documented a general preference for groups that are moderately sized. All things being equal, individuals (at least in Western, industrialized settings) identify more strongly with groups that represent a numerical minority of the total observable population than with groups that represent a majority of the population (Abrams, 1994; Blanz, Mummendey, & Otten, 1995; Brewer & Weber, 1994; Ellemers & van Rijswijk, 1997; Simon & Brown, 1987; Simon & Hamilton, 1994).

When individuals update their social identities in response to changes in their social environment (e.g., switch political parties; join the Shriners; go vegetarian), the relative size of every group in the environment also changes, which in turn alters the landscape for every other individual who shares that environment. As individuals leave and join groups, the level of the distinctiveness of each group changes relative to all the other groups in the social environment, which consequently changes the distinctiveness of every individual in every group. Thus, individual-level changes in social identity create opportunities for dynamic cascades and feedback throughout the environment. Though the ODT literature acknowledges such a fluid nature of the social environment (e.g., Turner, Oakes, Haslam, & McGarty, 1994), most research has focused on one individual responding to a static environment. Because there are potentially limitless causal pathways and recursive loops in such a dynamic social landscape, it can be challenging to form – let alone test – anything more than the most simple hypotheses about such a complex social system. Thus, the impact of individual-level preferences for optimal distinctiveness has been studied almost exclusively at the individual level, and group-level outcomes had remained unexplored until recently. This lack of exploration of group-level optimal distinctiveness was due, in large part, to the tools that are traditionally available to researchers interested in the psychological study of groups. As noted previously, lab-based experimentation is an excellent method for testing causal relationships at the level of the individual. However, the lab is not amenable to examining the interactions among multiple individuals over extended periods of time.

When conceptualized from a group dynamics perspective, it is easy to see why laboratory studies on isolated individuals are insufficient for fully capturing the dynamic processes that are logical outcomes of the predictions of ODT. Moreover, because ODT research has thus far only examined individuals, it is largely unknown how these preferences shape group formation and change. An explicit assumption of the theory has been that individuals generally maintain an equilibrium between their opposing needs for differentiation and for assimilation by switching to more “optimal” social identities when the need arises (Leonardelli et al., 2010). However, the group-level consequences of such an assumption had not been tested. As such, it was not clear whether individuals following this type of heuristic for social identity choice would actually sustain either a social equilibrium, in which individuals are relatively stable in their social identities, or even an individual equilibrium in which individuals are satisfied in an “optimally distinct” social identity.

Building an ABM of Social Identity: The SID Model

To address this gap, Smaldino and colleagues (2012) developed an agent-based model in which agents make social identity decisions (i.e., whether to select or abandon a particular group membership) based on the social identities of their neighbors (hereafter referred to as the SID model). The model was built upon the assumption that agents make these decisions in an attempt to satisfy their preference for groups that are optimally distinct. The overarching goal was to observe the resultant dynamics of group formation as a function of variations in group size preferences and the size of social neighborhoods.

Designing an ABM of a complex social system requires the researcher to articulate the specific components of that system and to specify the properties of those components and how they relate to one another. There are myriad ways in which any particular system might be decomposed into parts, and there is not necessarily a uniquely best decomposition (Kauffman, 1971). The questions being asked about the system will often drive how the system is modeled. The model designer should therefore strive to create a model system that captures the aspects of the real-world system in which she is interested, with enough complexity to create a credible analogy between the model and real-world systems, but otherwise as simple as possible so as to maximize generalizability and minimize obscuring artifacts. The complexity of a model may depend on how much is already known about the system in question, as well as the complexity of the theoretical mechanisms under investigation. A complicated theory may demand a complicated model. Yet simple models can often be extremely revealing. The overarching lesson in the history of complexity science is that complex or unexpected behavior often emerges from an instantiation of relatively simple relations and behaviors.

Because ODT had not been previously studied using an ABM, the SID model is by necessity as simple as possible in order to test the basic assumptions of the theory. Each agent has only two attributes: membership in a particular group (i.e., its social identity) and a preference for that identity to have a particular degree of distinctiveness (i.e., its optimal distinctiveness). Later, we will also discuss the introduction of a third attribute: spatial location. Each agent follows two simple behavior rules. First, the agent observes the social identities of all of the other agents in its neighborhood, noting how many individuals (including itself) possess each social identity. Second, if the agent observes a social identity that is closer to its optimal degree of distinctiveness than its current social identity, it switches to that social identity. Thus, an agent will stay with its current social identity only when it observes no other social identities that are closer to its optimal degree of distinctiveness.

The formal description is necessarily more precise, if perhaps less rhetorically pleasing. A population of N agents is initialized so that each agent is randomly assigned to one of k social identities. Each agent has a preference for optimal distinctiveness, d^*, which is the desired proportion of the agents in its neighborhood with whom the agent shares a social identity (e.g., if $d^* = 0.2$, then each agent would be maximally satisfied if exactly 20% of the agents in its neighborhood shared its social identity). For simplicity, we specified that d^* is the same for all agents and does not change over time. The model dynamics proceed in discrete time steps. At each time step, each agent, in random order, assesses the distinctiveness, d_g, of each social identity g in its neighborhood, which is equal to the proportion of agents who have that social identity (i.e., $d_g = n_g/N$, where n_g is the number of agents with social identity g). The agent then switches to the social identity with a distinctiveness closest to d^*, so that it minimizes $|d_g—d^*|$, the difference between its optimal and actual degrees of distinctiveness. If two or more social identities are equally close to optimal, one is selected at random unless one of these is the agent's current social identity, in which case it retains its current social identity. Each agent acts in turn, so that one agent's choice may directly influence the choices of agents that "move" after it. This process continued either until a stable equilibrium was reached, in which no agent had a better option than its current social identity, or for 5,000 time steps, whichever came first. For the initial runs, all agents in the population were defined to share the same neighborhood, and a stable equilibrium was always reached. Next, we explored the role of the structure of social networks. Agents were situated in on a lattice and only assessed the social identities of those agents within a limited radius (their local neighborhood; see below). In these cases, an equilibrium was not always reached, as agent choices sometimes led to perpetual feedback and identity switching, and the results reflect the average situation after 5,000 time steps.

Model results. Simulations demonstrated that, in a well-mixed population in which each agent knew and responded to the social identities of every other agent (i.e., all agents shared the same big neighborhood), all agents ended up in groups that were overly large and therefore sub-optimally distinct (Figure 1A). This outcome happened at all levels of desired optimal distinctiveness, d^* (see Smaldino et al. (2012) for a detailed explanation of the causal dynamics). At first glance, it therefore might appear that the assumptions of ODT do not lead to individuals who are satisfied with their degree of distinctiveness. However, in considering this result, it became apparent that individuals in real populations rarely have full knowledge of everyone in their environment, but rather are organized in social networks with limited knowledge beyond their local connections.

Though ODT does not explicitly articulate network constraints in shaping distinctiveness, we nevertheless elaborated upon the well-mixed initial model by introducing a simple structure, such that agents responded only to the social identities of their neighbors. Agents were placed on a square lattice with periodic boundaries (so as to eliminate edge effects, the agents on the far right edge are neighbors with the agents on the far left, and similar for the top/bottom edges). The lattice was a 50×50 grid containing 2,500 agents, one to a cell. An agent's *local neighborhood* was defined as all the agents within a square centered on the agent and extending *r* cells in each of the four cardinal directions (up, down, left, right), such that a neighborhood consisted of $(2r + 1)^2$ agents. When this change was implemented, stable, spatially contiguous groups emerged in which all agents achieved high levels of satisfaction, that is, their groups were close to optimally distinct (Figure 1). As such, the SID model identified one boundary condition to optimal distinctiveness previously unspecified by ODT: network structure, which is analogous to spatial, social, or organizational distance in the real world.

The inclusion of a network structure proved to be critical to our model of social identity dynamics. We followed up on this finding by probing what the model could tell us about the effects of different types of network structures on agent satisfaction. We systematically varied the size of each agent's local neighborhood, *r*. Figure 2 shows agents' average distinctiveness at equilibrium as a function of local neighborhood size for three different values of d^*. These results indicate that agents' satisfaction with their distinctiveness is best achieved through social neighborhoods of intermediate size, which complements and extends ODT's predictions regarding preferences for moderately-sized groups. When a neighborhood is too small, agents may become trapped either in homogeneous groups or in groups in which they are unique. However, once neighborhoods exceed a certain size, agents again end up in groups that are too large, leaving them with an unsatisfied need for increased distinctiveness.

Implications for Theory Development

The SID model is a step toward a more precise body of theory linking individual psychology with group behavior. At a basic level, the model provides a clear demonstration of how individual decisions can feed back into stimuli for others – a point that previously existed only theoretically in the optimal distinctiveness literature. The SID model also highlights the previously unarticulated role that network structure and size plays in how individuals assess and modify their distinctiveness. One novel prediction generated by the model is that a moderate neighborhood size is ideal for generating optimally distinct groups. The question of how such moderately sized local neighborhoods are formed – perhaps through a combination of physical space, social structure, and cognitive constraints – also represents a new and potentially illuminating direction for both empirical and theoretical research.

The findings of the SID model represent a theoretical advance in the study of optimal distinctiveness. Empirical work remains to be done to investigate how individuals perceive their relevant networks, and how neighborhood size and population density influence perceptions and decisions related to distinctiveness. As we have stated before, modeling does not and cannot replace carefully-designed empirical studies. Ideally, lines of communications between modelers and empiricists should remain open and active in both directions. Models can demonstrate possibility and suggest new directions for research, but a single model cannot solve the problems of an entire theory any more than can a single experiment. There are assumptions inherent in any model – just as there are for any other empirical method – and future work can test the robustness or boundaries of these assumptions. Nevertheless, ABMs are showing themselves to be an essential tool for theory development in the study of group dynamics and related fields. In the next section, we discuss general strategies for researchers hoping to develop theory with ABMs.

Theory Development with Agent-Based Models

A scientific theory provides a scaffold for making sense of data, and facilitates the generation of pointed questions to enable researchers to make even *better* sense of the world (Schank, May, & Joshi, 2014). We focus here on developing theory in the science of group dynamics and organizational psychology.

Just as with a laboratory experiment, a case study, or most other research paradigms, the primary value of a formal model lies in its ability to function as an analogy. A mouse immune system is analogous to the human response to disease. A laboratory experiment is analogous to a real-world social setting. As an analogy, a model is intended to account for key processes, components, and variables as specified by a hypothesis. It is not supposed to be – nor can it be – an exhaustive or literal representation of reality. No methodology can be that. However, it can often be useful to compare a model to outcome information (data) from another source, which can include laboratory experiments, survey results, etc. – which of course are also snapshots and analogies. Models can clarify the terms of a hypothesis, provide alternatives to null hypothesis testing (Schank, 2001), and illuminate boundary conditions and future research questions by their failure to adequate reproduce empirical results (Wimsatt, 1987). Consider also the similarity between an ABM and a thought experiment. Thought experiments are often the seed of scientific hypotheses and theories. The mind envisions a scenario and its consequences, and new questions and possibilities arise. However, the human mind, as powerful as it is, faces limitations in conceptualizing complex systems. ABMs can function as prostheses for the imagination, increasing our power to develop theory.

Given that model building is (or should be) guided by theory, comparing a model to reality can be conceived of as a test of the theory at its base. In making such a comparison, the question becomes, "Does the model, as an instantiation of theory, generate data that are congruent with data from another relevant source?" Both "yes" and "no" answers to this question can be instructive. A "yes" is essentially a multi-modal convergent validation of theory in that the theory instantiated as a formal model can produce outcomes congruent with data from other sources. A "no" may indicate that the theory is wrong or that the model is mis-specified and does not accurately reflect theory. However, a "no" does not necessarily invalidate theory, but rather can help to expand or build theory. Articulating a theory as a formal model often imposes an exactness which verbal formulations lack. Put another way, formal modeling requires the researcher to fill in theoretical gaps to a greater degree than many other research designs. Previously unspecified details may emerge as important moderators or boundary conditions unconsidered by verbal instantiations of a theory, and as such represent touchstones for theory development.

Data generated by a model may not match data from another source for a variety of reasons. For example, model data may reflect emergent properties of dynamic processes that other research modalities may not capture. Similarly, some group dynamics may depend on causal loops that emerge only as these processes feed back upon themselves until they reach a tipping point. Longitudinal studies and other modalities that incorporate a temporal dimension are typically limited by practical constraints such as funding and attrition and therefore may not adequately capture group dynamics that unfold over long periods of time. Models, however, can represent timescales from the very short to the very long. As such, model data should not necessarily be expected to be congruent with data collected in a less comprehensive manner, but instead may represent outcomes beyond the scope of data collected or generated with other research methods.

The link between individual and collective behavior is not always intuitive. ABMs have shown us that individual-level behavioral rules often do not translate directly into simple predictions of large-scale outcomes, because structural organization, feedback, and system memory all affect group-level processes. In other words, group dynamics are more than the sum of their constituent parts. One of the strengths of ABMs is the ability to directly incorporate those properties of complex social systems, creating the potential for unexpected outcomes to arise despite complete knowledge of initial conditions (Epstein, 1999).

Resources for Theory Development with ABMs

There are three types of resources, all important, that we recommend for researchers who wish to add ABMs to their methodological toolkit. First, seek out published models, which illustrate how other researchers have applied ABMs to tackle problems in theory development. We have presented one example of an ABM applied to a question from the behavioral sciences, but there are many more. Some modeling topics of interest to group dynamics researchers include opinion dynamics (Nowak, Szamrej, & Latané, 1990; Axelrod, 1997; Deffuant, Neau, Amblard, & Weisbuch, 2000; Centola & Macy, 2007), mate choice (Kalick & Hamilton, 1986; Simão & Todd, 2002; Smaldino & Schank, 2012), cooperative group formation (Gray et al. 2014; Smaldino & Lubell 2014), and distributed social cognition (Stasser, 1988; Smith & Collins, 2009; Smith, 2014).

Second, read philosophical literature on how to think about models. There are many discussions that go far beyond the present discourse on that matter (e.g., Bedau, 1999; Beisbart, 2012; Epstein, 1999; 2008; Kauffman, 1971; Parke, in press; Schank, 2001; Weisberg, 2007; Wimsatt, 1987). Such readings can help the aspiring modeler in both model design and interpretation, as well as providing broad support for the epistemic value of computational models. Finally, take time to learn to model. Of course, becoming an expert at model design and programming is not a requirement for every researcher using modeling. Just as experimental inquiry utilizes experts in statistical analysis, exploiting ABMs may similarly involve specialists. Nevertheless, it will benefit anyone employing ABM to have some experience with building and testing models. Thankfully, it has never been easier to learn to learn the requisite programming to construct an ABM. For example, the NetLogo modeling environment (Wilensky, 1999) was specifically designed to be accessible to non-programmers, and has a shallow learning curve. NetLogo is free, open-source, and comes bundled with a large library of demo models and a detailed tutorial. For those requiring more formal instruction, a recent textbook on agent-based modeling (Railsback & Grimm, 2012) teaches modeling techniques along with a practical guide to programming models in NetLogo. Additionally, an introduction to ABM using NetLogo (Janssen, 2010) has recently been published free online (www.openabm.org/book/1928/games-gossip). For more advanced programmers who desire more precise control over the details of their models, MASON (Luke et al., 2005) is an open source simulation library written in Java, is extensively documented with numerous demo models and tutorials, and is maintained by an active community of users and developers. Several other software tools and tutorials are also available depending on a user's needs and expertise.

Combining Agent-Based Modeling with Other Methods

In addition to adding rigor to thought experiments, ABMs may also be combined with data as a more systematic alternative to null hypothesis testing common in behavioral research. This allows for fine-grained model comparison in which the models being tested propose explicit mechanisms rather than mere statistical relationships. Here will discuss two examples of this application of ABM. Because the ABM literature on organizational psychology is still sparse, the examples will be drawn from elsewhere – the first on human crowd dynamics, the second from animal behavior. Though neither study comes directly from the field of organizational psychology, both involve group dynamics and social organization.

Moussaïd, Helbing, and Theraulaz (2011) investigated the dynamics of pedestrians maneuvering toward an exit through a crowd. Though other models of pedestrian dynamics had been previously proposed, their model was based on psychological theories of decision making and social norms rather than on physics-based approaches that lack psychological underpinnings, and involved a formal instantiation of two simple decision heuristics. Their model was tested against empirical data derived from video analyses of pedestrians moving through a corridor in which a confederate was either standing still or moving in the opposite direction, as well as data from crowds of different densities moving through a bottleneck. They were able to fit their model to match all these cases. The wide range of conditions against which the model was tested provides strong support for their theories of perception and decision making among pedestrians. Importantly, validating their model with real-world data allowed them to generate new questions concerning more complex scenarios, and provides predictive power for future industrial design and emergency preparation.

Moussaïd and colleagues (2011) used empirical data to validate their model, but it is also possible to generate a range of models and then use data-fitting to select the best among models. In this way, the data can be used directly to select which theory, formalized as an instantiation of an ABM, is best supported empirically. Schank (2008) investigated the development of social behavior in rat pups using a combination of experiments and modeling. He filmed the pups moving in an open arena alone and in groups of eight, at both seven and ten days of age. The pups' positions, directional headings, and surroundings were catalogued at regular time intervals. At the same time, he developed an ABM of probabilistic movement based on the individual's surroundings (e.g., relative location of walls, other pups), and a genetic algorithm (Mitchell, 1997) was implemented to optimize model parameters to fit the data. The best-fit model parameters showed that the group behavior of seven-day-old pups could be explained as the aggregation of many pups, each behaving as if it was alone. In other words, the pups reacted to one another no differently than they reacted to a wall, suggesting that their huddling behavior involved no social cognition. In contrast, the group behavior of 10-day-old pups was *not* explained by the parameters that best fit their individual behavior. 10-day-old pups behaved differently alone and in groups, supporting the hypothesis that social cognition comes online in rats around 10 days of age. In this instance, the agent-based model functioned as a tool for data analysis, providing more precise hypotheses than would be possible by verbal theories alone.

Conclusion

Many important group-level phenomena result from repeated interactions among multiple individuals that are influenced by feedback across several levels of organization. Insight into these processes can therefore be difficult to achieve solely through the use of traditional research methods that focus on intentional actions by isolated individuals. Researchers interested in group dynamics may therefore benefit from concepts and techniques used to understand complex dynamical systems.

The purpose of this paper was to elaborate to researchers studying group dynamics and organizational psychology on the benefits of adopting agent-based modeling for developing theory in the study of groups. In doing so, we have largely described the process of modeling as straightforward and concrete. Like most things in science, however, the true path to success is rarely linear, and the process of investigation via modeling may be quite recursive. At each stage of specification, new questions, ideas, and implementations may present themselves. In contrast to an experiment that includes human subjects, and in addition to the other strengths of modeling discussed in this paper, it is generally quite easy to perform additional tests or to add new parameters to computational models without disrupting the flow of data collection in the lab or involving institutional review boards. We hope that the practical, conceptual, and theoretical benefits of modeling discussed in this paper will inspire researchers to add formal models to their methodological toolkit.

References

Abrams, D. (1994). Social self-regulation. *Personality and Social Psychology Bulletin, 20*(5), 473-483.

Anderson, E. G., & Lewis, K. (2013). A dynamic model of individual and collective learning amid disruption. *Organization Science, 25,* 356-376.

Axelrod, R. (1997). The dissemination of culture. *Journal of Conflict Resolution, 41*, 203–226.

Bak, P. (1996). *How nature works: The science of self-organized criticality.* New York: Copernicus.

Bedau, M. A. (1999). Can unrealistic computer models illuminate theoretical biology? In *Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop*, (pp. 20–23). Orlando, FL: GECC.

Beisbart, C. (2012). How can computer simulations produce new knowledge? *European Journal for Philosophy of Science, 2,* 395–434.

Berdahl, J. L. (1998). The dynamics of composition and socialization in small groups: Insights gained from developing a computational model. In M. A. Neale, E. A. Mannix, & D. H. Gruenfeld (Eds.), *Research on managing groups and teams* (Vol. 1, pp. 209-227). Greenwich, CT: JAI.

Blanz, M., Mummendey, A., & Otten, S. (1995). Positive-negative asymmetry in social discrimination: The impact of stimulus valence and size and status differentials on intergroup evaluations. *British Journal of Social Psychology, 34(4),* 409-419.

Brauer, F., & Castillo-Chávez, C. (2001). *Mathematical models in population biology and epidemiology.* New York: Springer.

Brewer, M. B. (1991). The social self: On being the same and different at the same time. *Personality and Social Psychology Bulletin, 17*(5), 475-482.

Brewer, M. B. (2007). The importance of being *we*: Human nature and intergroup relations. *American Psychologist, 62*(8), 728-738.

Brewer, M. B., & Caporael, L. (2006). An evolutionary perspective on social identity: Revisiting groups. In M. Schaller, J. Simpson, & D. Kenrick (Eds.), *Evolution and social psychology* (pp. 143-161). New York: Psychology Press.

Brewer, M. B., & Pickett, C. L. (2002). The social self and group identification: Inclusion and distinctiveness motives in interpersonal and collective identities. In J. Forgas & K. Williams (Eds.), *The social self: Cognitive, interpersonal, and intergroup perspectives* (pp. 255–271). Philadelphia: Psychology Press.

Brewer, M. B., & Weber, J. G. (1994). Self-evaluation effects of interpersonal versus intergroup social comparison. *Journal of Personality and Social Psychology, 66*(2), 268-275.

Brockman, J. (1995). *Third culture: Beyond the scientific revolution.* New York: Touchstone.

Campbell, D. T. (1974). 'Downward causation' in hierarchically organised biological systems. In F. Ayala & T. Dobzhansky (Eds.), *Studies in the philosophy of biology* (pp. 179–186). University of California Press.

Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology, 113,* 702-734.

Cronin, M. A., Weingart, L. R., & Todorova, G. (2011). Dynamics in groups: Are we there yet? *The Academy of Management Annals*, *5*(1), 571-612.

Crowder, R. M., Robinson, M. A., Hughes, H. P. N., & Sim, Y. W. (2012). The development of an agent-based modeling framework for simulating engineering team work. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, 42(6), 1425–1439.

Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems, 3*, 87–98.

Ellemers, N., & Van Rijswijk, W. (1997). Identity needs versus social opportunities: The use of group-level and individual-level identity management strategies. *Social Psychology Quarterly, 60*(1), 52-65.

Epstein, J. M. (1999). Agent-based computational models and generative social science. *Complexity, 4(5),* 41–60.

Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation, 11(4),* 12.

Fang, C., Lee, J., & Schilling, M. A. (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science, 21,* 625–642.

Gilbert, N., & Terna, P. (2000). How to build and use agent-based models in social science. *Mind & Society, 1,* 57–72.

Gray, K., Rand, D. G., Ert, E., Lewis, K., Hershman, S., & Norton, M. I. (2014). The emergence of "us and them" in 80 lines of code: Modeling group genesis in homogeneous populations. *Psychological Science, 25,* 982-990.

Haldane, J. B. S. (1964). A defense of beanbag genetics. *Perspectives in Biology and Medicine*, *7(3),* 343–359.

Harrison, J.R., Lin, Z., Carroll, G.R., & Carley, K.M. (2007). Simulation modeling in organizational and management research. *Academy of Management Review*, 32, 1229-1245.

Hogg, M. A., & Terry, D. J. (Eds.). (2001). *Social identity processes in organizational contexts*. Psychology Press.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, 79,* 2554–2558.

Hughes, H. P. N., Clegg, C. W., Robinson, M. A., & Crowder, R. M. (2012). Agent-based modelling and simulation: The potential contribution to organizational psychology. *Journal of Occupational and Organizational Psychology*, *85(3),* 487–502.

Ilgen, D.R., Hollenbeck, J.R., Johnson, M., & Jundt, D. (2005). Teams in organizations: From input-process output models to IMOI models. *Annual Reviews of Psychology, 56*, 517–543.

Johnson-Laird, P. N. (1983). *Mental models: Toward a cognitive science of language, inference, and consciousness.* Cambridge, MA: Harvard University Press.

Kalick, S. M., & Hamilton, T. E. (1986). The matching hypothesis reexamined. *Journal of Personality and Social Psychology, 51,* 673–682.

Kauffman, S. A. (1971) Articulation of parts explanation in biology and the rational search for them. In R. C. Buck & R. S. Cohen (Eds.), *PSA 1970* (pp. 257–272). Philosophy of Science Association.

Leonardelli, G. L., Pickett, C. L., & Brewer, M. B. (2010). Optimal distinctiveness theory: A framework for social identity, social cognition, and intergroup relations. In M. P. Zanna & J. M. Olson (Eds.), *Advances in experimental social psychology* (Vol. 43, pp. 66–115). New York: Academic Press.

Luke, S., Cioffi-Revilla, C., Sullivan, K., & Balan, G. C. (2005). MASON: A multi-agent simulation environment. *Simulation, 81*, 517–527 .

MacCoun, R. J. (2012). The burden of social proof : Shared thresholds and social influence. *Psychological Review, 119,* 345-372.

Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology, 28,* 143-166.

Marks, M. A., Zaccaro, S. J., & Mathieu, J. E. (2000). Performance implications of leader briefings and team-interaction training for team adaptation to novel environments. *Journal of Applied Psychology, 85*, 971–986.

McClelland, J. L., Rumelhart, D. E., & the PDP Research Group. (1986). *Parallel distributed processing, Vol. 2: Psychological and biological models*. Cambridge, MA: MIT Press.

McGrath, J.E. (1986). Studying groups at work: Ten critical needs for theory and practice. In P. Goodman and Associates (Eds.), *Designing effective work groups* (pp. 362–392). San Francisco: Jossey-Bass.

McGrath, J. E., Arrow, H., & Berdahl, J. L. (2000). The study of groups: Past, present, and future. *Personality and Social Psychology Review, 4,* 95–105.

Miller, J. H., & Page, S. E. (2007). *Complex adaptive systems: An introduction to computational models of social life*. Princeton, NJ: Princeton University Press.

Mitchell, M. (1997). *An introduction to genetic algorithms.* Cambridge, MA: MIT Press.

Moussaïd, M., Helbing, D., & Theraulaz, G. (2011). How simple rules determine pedestrian behavior and crowd disasters. *Proceedings of the National Academy of Sciences, 108,* 6884–6888.

North, D.C. (1990). *Institutions, institutional change and economic performance*. Cambridge, UK: Cambridge University Press.

Nowak, A., Szamrej, J., Latané, B. (1990). From private attitude to public opinion: A dynamic theory of social impact. *Psychological Review, 97*, 362–376.

Parke, E. (in press). Experiments, simulations, and epistemic privilege. *Philosophy of Science.*

Popper, K. (1963). *Conjectures and refutations.* London: Routledge.

Railsback, S. F., & Grimm, V. (2012). *Agent-based and individual-based modeling: A practical introduction*. Princeton, NJ: Princeton University Press.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Resnick, M. (1994). *Turtles, termites, and traffic jams: Explorations in massively parallel microworlds*. MIT Press.

Schank, J. C. (2001). Beyond reductionism: Refocusing on the individual with individual-based modeling. *Complexity, 6(3)*, 33–40.

Schank, J. C. (2008). The development of locomotor kinematics in neonatal rats: An agent-based modeling analysis in group and individual contexts. *Journal of Theoretical Biology, 254,* 826–842.

Schank, J. C., May, C. J., & Joshi, S. S. (2014) Models as scaffold for understanding. In L. R. Caporael, J. R. Griesemer, & W. C. Wimsatt (Eds.), *Developing scaffolds in evolution, culture, and cognition* (pp. 147–167). Cambridge, MA: MIT Press.

Schelling, T. C. (1978). *Micromotives and macrobehavior.* New York: W. W. Norton.

Shepherd, D., & Haynie, J. M. (2009). Birds of a feather don't always flock together: Identity management in entrepreneurship. *Journal of Business Venturing, 24,* 316–337.

Simão, J., & Todd, P. M. (2002). Modeling mate choice in monogamous mating systems with courtship. *Adaptive Behavior, 10,* 113–136.

Simon, B., & Brown, R. (1987). Perceived intragroup homogeneity in minority-majority contexts. *Journal of Personality and Social Psychology, 53*(4), 703-711.

Simon, B., & Hamilton, D. L. (1994). Self-stereotyping and social context: The effects of relative in-group size and in-group status. *Journal of Personality and Social Psychology, 66*(4), 699-711.

Smaldino, P. E. (2014a). The cultural evolution of emergent group-level traits. *Behavioral and Brain Sciences,* in press.

Smaldino, P. E. (2014b). Group-level traits emerge. *Behavioral and Brain Sciences,* in press.

Smaldino, P. E., & Lubell, M. (2014). Institutions and cooperation in an ecology of games. *Artificial Life, 20,* 207-221.

Smaldino, P. E., Pickett, C., Sherman, J., & Schank, J. (2012). An agent-based model of social identity dynamics. *Journal of Artificial Societies and Social Simulation, 15(4),* 7.

Smaldino, P. E., & Richerson, P. J. (2013). Human cumulative cultural evolution as a form of distributed computation. In P. Michelucci (Ed.), *Handbook of human computation* (pp. 979–992). New York: Springer.

Smaldino, P. E., & Schank, J. C. (2012). Human mate choice is a complex system. *Complexity, 17*(5), 11–22.

Smith, E. R. (2014). Evil acts and malicious gossip: A multiagent model of the effects of gossip in socially distributed person perception. *Personality and Social Psychology Review*.

Smith, E. R., & Collins, E. C. (2009). Contextualizing person perception: Distributed social cognition. *Psychological Review, 116,* 343–364.

Smith, E., R. & Conrey, F. R. (2007). Agent-based modeling: A new approach for theory building in social psychology. *Personality and Social Psychology Review, 11*, 87–104.

Stasser, G. (1988). Computer simulation as a research tool: The DISCUSS model of group decision making *Journal of Experimental Social Psychology, 24,* 393-422.

Turner, J. C., Oakes, P. J., Haslam, S. A., & McGarty, C. (1994). Self and collective: Cognition and social context. *Personality and Social Psychology Bulletin, 20*(5), 454–463.

Vallacher, R. R., & Nowak, A. (1997). The emergence of dynamical social psychology. *Psychological Inquiry, 8,* 73–99.

Van Knippenberg, D. (2000). Work motivation and performance: A social identity perspective. *Applied Psychology*, *49*(3), 357-371.

Vignoles, V. L. (2011). Identity motives. In S. J. Schwartz, K. Luyckx, & V. L. Vignoles (Eds.), *Handbook of identity theory and research* (pp. 403–432). New York: Springer.

Weinhardt, J. M., & Vancouver, J. B. (2012). Computational models and organizational psychology: Opportunities abound. *Organizational Psychology Review, 2(4)*, 267–292.

Weisberg, M. (2007). Who is a modeler? *British Journal for the Philosophy of Science, 58,* 201–233.

Wilensky, U. (1999). *NetLogo*. Evanston, IL: Northwestern University, Center for Connected Learning and Computer-Based Modeling. Retrieved from http://ccl.northwestern.edu/netlogo/

Wilson, J.M., Goodman, P., & Cronin, M.A. (2007). Group learning. *Academy of Management Review, 32*, 1041–1059.

Wimsatt, W. C. (1974) Complexity and organization. In K. Schaffner & R. S. Cohen (Eds.), *PSA 1972* (pp. 67–86). Philosophy of Science Association.

Wimsatt, W. C. (1987). False models as means to truer theories. In M. H. Nitecki & A. Hoffman (Eds.), *Neutral models in biology* (pp. 23–55). New York: Oxford University Press.

Figure Captions

Figure 1.

Results of the SID model from the well-mixed and spatial versions. (A) Agents' average distinctiveness at the end of a simulation run, as a function of their optimal distinctiveness preferences (*d**) and neighborhood size (*r*). Data is averaged over 30 runs for each parameter condition. The well-mixed model represents the case in which each agent's neighborhood is the entire population. The dotted grey line indicates an ideal case in which the average distinctiveness is equal to agents' optimal distinctiveness. (B – E) Spatial organization of agents at equilibrium for example runs with different optimal distinctiveness preferences (*d**) and local neighborhood sizes (*r*). Each agent's color represents a particular social identity. $\langle d \rangle$ refers to the average distinctiveness across the population. All of these runs settled to highly stable spatial patterns except (C), in which around 25% of agents switched identities each time step.

Figure 2.

Average distinctiveness in the population as a function of neighborhood size for three values of *d** (indicated by the dotted lines). Data are averaged over 30 simulation runs for each parameter condition.

Figure 1.

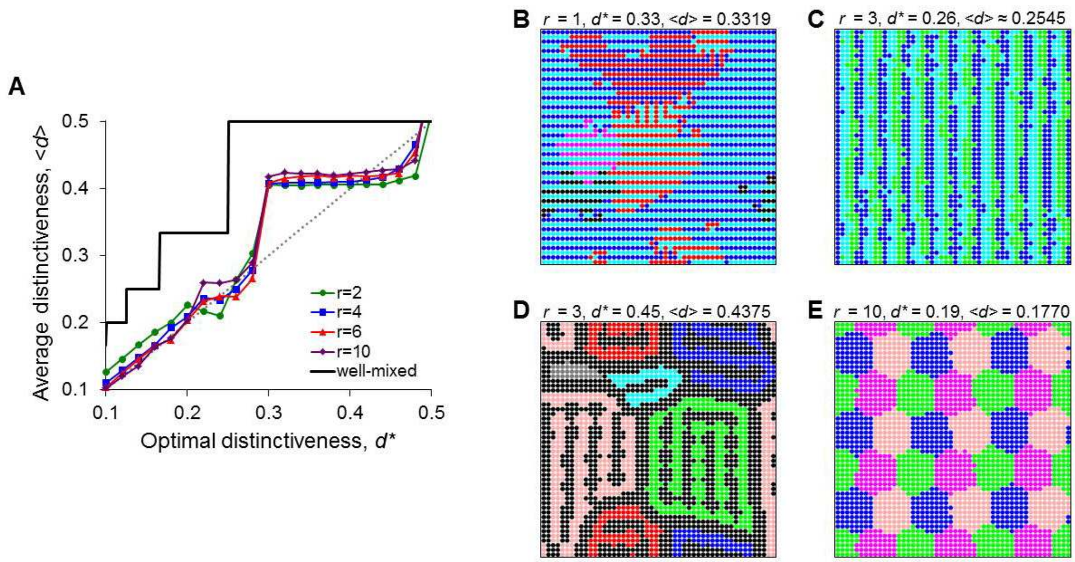
A
Average distinctiveness, <d>
Optimal distinctiveness, d*
0.1
0.2
0.3
0.4
0.5
r=2
r=4
r=6
r=10
well-mixed
B r = 1, d* = 0.33, <d> = 0.3319
C r = 3, d* = 0.26, <d> ≈ 0.2545
D r = 3, d* = 0.45, <d> = 0.4375
E r = 10, d* = 0.19, <d> = 0.1770

Figure 2.

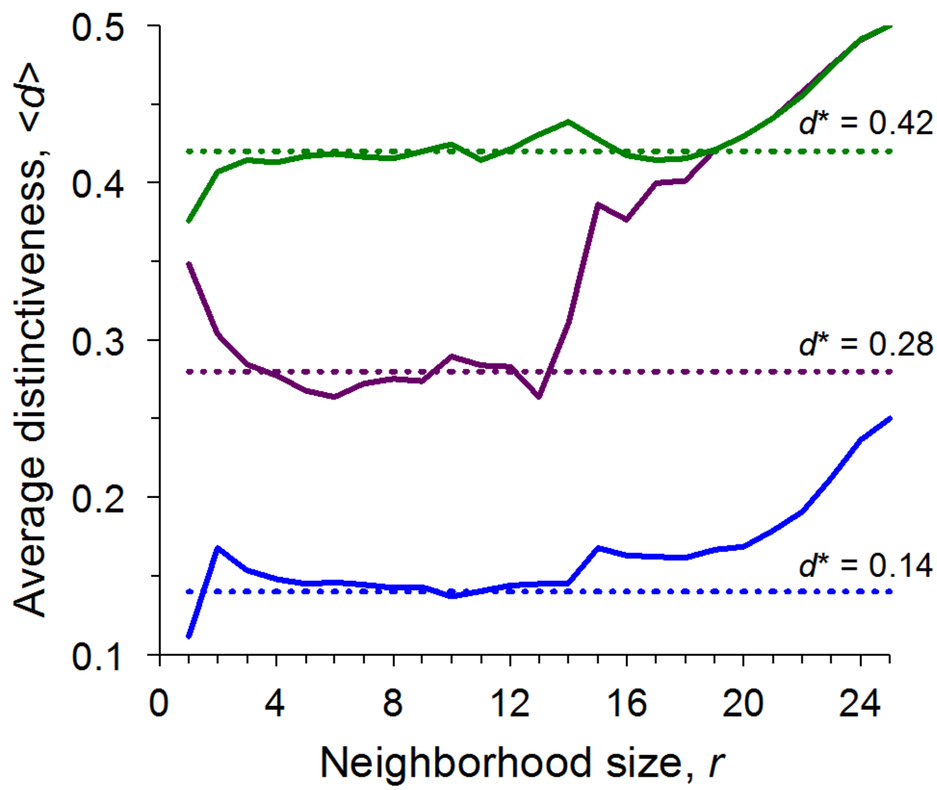
0.5
0.4
0.3
0.2
0.1
Average distinctiveness, <d>
d* = 0.42
d* = 0.28
d* = 0.14
0
4
8
12
16
20
24
Neighborhood size, r