

# On the Coevolution of Stereotype, Culture, and Social Relationships: An Agent-Based Model

Social Science Computer Review  
201X, Vol XX(X), 1–17  
© The Author(s) 2013  
Reprints and permission:  
sagepub.com/journalsPermissions.nav  
DOI: 10.1177/0894439313511388  
ssc.sagepub.com



Kenneth Joseph<sup>1</sup>, Geoffrey Morgan<sup>1</sup>,  
Michael K. Martin<sup>1</sup>, and Kathleen M. Carley<sup>1</sup>

AQ1

## Abstract

The theory of constructivism describes how shared knowledge, representative of cultural forms, develops between individuals through social interaction. Constructivism argues that through interaction and individual learning, the social network (who interacts with whom) and the knowledge network (who knows what) coevolve. In the present work, we extend the theory of constructivism and implement this extension in an agent-based model (ABM). Our work focuses on the theory's inability to describe how people form and utilize stereotypes of higher order social structures, in particular observable social groups and society as a whole. In our ABM, we formalize this theoretical extension by creating agents that construct, adapt, and utilize social stereotypes of individuals, social groups, and society. We then use this model to carry out a virtual experiment that explores how ethnocentric stereotypes and the underlying distribution of culture in an artificial society interact to produce varying levels of social relationships across social groups. In general, we find that neither stereotypes nor the form of underlying cultural structures alone are sufficient to explain the extent of social relationships across social groups. Rather, we provide evidence that shared culture, social relations, and group stereotypes all intermingle to produce macrosocial structure.

## Keywords

agent-based model, agent simulation, constructivism, stereotypes, social schemas

## Introduction

The process by which people interact, exchange information, and consequently learn is the central component of Carley's (1990, 1991) theory of constructivism. Constructivism argues that individual learning from interactions takes place on two levels. First, social interactions bring us new knowledge, knowledge that represents bits of larger cultural forms we collect over time. Second,

AQ2

---

<sup>1</sup> Computation, Organization and Society, Institute for Software Research, Carnegie Mellon University, Pittsburgh, PA, USA

### Corresponding Author:

Kenneth Joseph, Computation, Organization and Society, Institute for Software Research, Carnegie Mellon University, Wean Hall, Room 4127, Pittsburgh, PA 15213, USA.  
Email: kjoseph@cs.cmu.edu

as humans receive and share knowledge with interaction partners, we “learn” a perception of what we expect them to know. Paired with the assumption of homophily (Lazarsfeld & Merton, 1954; McPherson, Lovin, & Cook, 2001), that people tend to interact with others similar to them, constructivism explains how social relationships evolve as, via interaction, the knowledge two actors believe themselves to share increases.

This approach to the coevolution of knowledge and social relationships has considerable explanatory power over the dynamics of social networks (Lizardo, 2006; Pachucki & Breiger, 2010) and has proved to be an effective tool for social simulation (Carley & Hill, 2001; Hirshman, Charles, & Carley, 2011). However, the theory of constructivism relies on two assumptions that make modeling large-scale social systems difficult. First, constructivism assumes that humans are able to retain a perception of each individual they know to exist. Second, the theory assumes that our perceptions of individuals are compartmentalized—what we know and learn about a specific individual does not inform our perception of anyone else.

In reality, while humans may hold a persistent perception of a small set of individuals, cognitive limitations prevent us from retaining specific views of each person we know to exist. Rather, humans make the most of our cognitive abilities by incorporating what we learn of individuals into generalized beliefs of what groups of people are likely to know. Our perception of a specific individual is, consequently, a mishmash of what we have gleaned from prior experiences with him or her and others we, wittingly or unwittingly, determined were similar (Hilton & von Hippel, 1996; Schaller & Lataner 1996). Obvious examples of “similar” include race or gender, but we will define similarity in the present work as two actors that both belong to one of the two higher order social structures—a *social group* (e.g., Tajfel & Turner, 1979) or the *generalized other*, defined by Mead (1925) as the cognitive depiction of society as a whole.

These perceptions of higher order social structures are the basis for social stereotypes. Stereotypes are a useful cognitive tool in that they allow us to navigate an impossibly large social world to find others we believe we can benefit from interacting with. However, while useful, these stereotypes are error-prone. In particular, our stereotypes of others are frequently biased by a belief that those in our own groups are much better candidates for beneficial interaction than those outside of them. This bias is often referred to as ethnocentrism, favoritism of one’s own group at the expense of others (e.g., Hartshorn, Kaznatcheev, & Shultz, 2012).

Constructivism, in its original form, cannot directly address the issue of ethnocentrism because it does not explain how we form perceptions of social groups or the generalized other. In the present work, we extend constructivism to model these perceptions. We describe how homophily can be based on perceived knowledge similarity while still accounting for the fact that humans constantly make inferences using social stereotypes. The crux of our extension is a new model of implicit social cognition (Greenwald & Banaji, 1995) that utilizes the ideas of cognitive schemas, a general cognitive model that includes mechanisms for stereotyping (e.g., Rumelhart, 1978), and the instantiation of schemas due to the activation of concepts (Collins & Loftus, 1975).

We develop our extensions to constructivism in a simple yet elegant agent-based model (ABM) that connects three levels of the social world. At a sociostructural level, agents are placed into social groups. At the individual level, agents use the group affiliation of their interaction partners to create and refine schemas of other agents, social groups, and the generalized other. Finally, at the dyadic level, agents use schemas to determine the knowledge of others. They then use homophily to determine with whom they will interact. Importantly, all of this occurs within an ABM that is both more computationally efficient and more cognitively plausible than those implementing previous instantiations of constructivism, a fact we detail further in other work (Morgan, Joseph, & Carley, n.d.).

Using this model, we can begin to explore how ethnocentric stereotypes affect intergroup relationships in a society. To this end, we carry out a virtual experiment to understand how varying the degree of ethnocentrism in an artificial society affects the formation of social relationships across

social groups under three different models of the underlying cultural structure. As culture can be thought of as a collection of commonly shared facts (Axelrod, 1997; Carley, 1991; Lizardo, 2006), we represent cultural forms as bits of knowledge that spread throughout the society. Results show that the true distribution of underlying knowledge in a society can serve to combat ethnocentrism as agents learn, interact, and update their social stereotypes. However, it can also serve to exacerbate even small amounts of ethnocentrism, suggesting the dynamic interplay between cultural forms and stereotypes in large social systems.

The rest of this article is structured as follows. In the second section, we detail related work across a variety of scholarly domains. In the third section, we detail our model, and in the fourth section we provide details of our virtual experiment. The fifth section gives results of the virtual experiment, and the sixth section concludes with a discussion of the broader implications of our study, its limitations, and observations of avenues for future work.

## Related Work

We rely on work covering how cognitive processes within the individual, interaction patterns along the dyad, and societal-level structure interact to produce cultural forms and social networks. We now situate our work vis-related theory at each of these levels of sociality.

### *The Individual*

A plethora of work exists on how social stereotypes form and develop—we refer the reader to the reviews of Greenwald and Banaji (1995) and Hilton and von Hippel (1996) for recent introductions (and note that we cover only the iceberg’s tip here). In the present work, we blend the *prototype* and *exemplar* models (Hilton & von Hippel, 1996) of stereotyping behavior, an approach in line with calls in the social psychology literature (Hamilton & Mackie, 1990). Prototype theory suggests that humans have representations of social groups in their cognition and use these representations to make inferences about individuals within these groups (Hilton & von Hippel, 1996). Exemplar theory suggests that instead of cognitive representations of social groups, people hold perceptions of idealized individual group members, which in turn serve as *exemplars* of social groups (Garcia-Marques & Mackie, 1999; Smith & Zárate, 1992). In the model presented here, agents hold perceptions of the knowledge of individuals, social groups, and the generalized other. As dictated by prototype theory, agents can update their perceptions of higher order social structures through interaction. As dictated by exemplar theory, agents can use their knowledge of individuals to construct perceptions of social groups.

Despite their differences, both exemplar and prototype theory rely on the idea that cognitive representations exist as *schemas*. Rumelhart (1978) defines schemas as “data structures” comprised of variables that represent what our minds expect of a given situation. Schema are *instantiated* when they fit to a given environment—that is, our mind uses the set of schemas whose variables best match those presented by our current situation. Once instantiated, a schema can “fill in” information about what we should expect variables we cannot observe to be like. Instantiated schemas are then updated to reflect what we have learned from the present situation, information that will be incorporated into our perception the next time the schema is used. Our approach to combining exemplar and prototype models of stereotyping is based on this underlying concept of a schema. Agents hold schemas of individuals, social groups, and the generalized other and instantiate them based on their “fit” to the situation at hand. Agents use schema to “fill in” their perception of individual’s underlying knowledge and update schemas instantiated via interaction with the new knowledge that they learn.

Unfortunately, the idea that schema “fit” a specific environment is difficult to model via schema theory alone. Activation theory (Collins & Loftus, 1975) provides a useful mechanism for understanding how this fit is determined. It suggests that humans hold a set of *concepts* in our cognition that have a specific level of activation. The activation level of a concept is increased when we think about the concept and decreases when we do not. A schema, it is argued, is instantiated not when its distribution of characteristic variables matches the environment, but when the concepts it is associated with reach a certain threshold of combined activation. Anderson (e.g., Anderson, 2007) and his colleagues have formalized activation theory in adaptive control of thought–rational (ACT-R; Anderson, Matessa, & Lebiere, 1997), a computational model of the mind. We use an approximation of ACT-R’s activation equations to model schema activation levels, described further in Morgan et al. (n.d.).<sup>1</sup>

### The Dyad

In the present work, we assume that activation is driven solely through interaction. Anderson and Schooler (1991) have explored this viewpoint, where they show that people are more likely to contact those they have recently interacted with in a way consistent with the activation equations in ACT-R. More specifically, we assume interaction activates a single concept, the one an agent holds of a specific individual. While only a single concept is triggered, activation theory states that activation of a particular concept *spreads* to a host of related concepts. In our model, when an agent interacts with an individual in a given social group, activation of the individual may spread to activate this group as well. Thus, schemas for both the individual and of the higher order social structures may be instantiated, and consequently may be updated, upon interaction. However, as two concepts differentiate over time, the level of activation that spreads between them dissipates. The spread of activation from a concept of an individual to the social groups she is in therefore decreases over time in our model. This process is related to the social psychological concept of *decategorization* (Wilder, 1986, as cited by Dovidio & Gaertner, 2010), whereby humans can come to regard members of social groups solely as individuals, rather than representatives of the social groups they are in.

While activation is an important consequence of a social interaction, it is not the only thing that results from one. Constructuralism argues that interactions also cause agents to exchange knowledge. This exchange allows agents to learn what others know, allowing them a lossy perception of the knowledge of those they have interacted with. Our extension to constructuralism defines how agents generalize what they learn during interaction to higher order social structures. Specifically, via our blend of exemplar theory and prototype theory, an agent may update his schematic representation of social groups and the generalized other as he learns new information about individuals he perceives to be in one and/or the other.

This learning process provides the link between dyadic interaction and the development of an agent’s stereotypes of higher order social structures. However, to allow agents to constantly learn new things about social groups would contradict recent research, suggesting that stereotypes of groups tend to harden and become fixed over time (Gregg, Seibt, & Banaji, 2006). In our model, we assume that as a schema of a higher order social structure persists, it gradually becomes more rigid, until eventually it does not bend even in the face of direct contrary evidence. As ongoing research exists on the malleability of group stereotypes over time (Dovidio & Gaertner, 2010), this presents an important variable to modify in our exploration of the parameter space.

### The Society

Several ABMs have focused on the formation of stereotypes of higher order social structures (Hales, 1998; Hartshorn et al., 2012). Such works seem, however, to be focused on game theoretic models of

social behavior, a focus different from the information diffusion–based interests of the present work. To this end, several models have also been built to describe how dyadic interaction patterns change over time with the diffusion of knowledge throughout an artificial society. This dynamism between cultural flow and dyadic interaction has led to understandings of how cultures can both merge and separate (e.g., Axelrod, 1997; Flanche & Macy, 2011; Watts & Strogatz, 1998).

AQ3

However, models such as these often assume that while dyadic interaction patterns can change, they can only do so within the bounds of some rigid, underlying social network structure. For example, Centola and Macy (2007) assume both that there exists some static social network agents may not stray outside of (the interaction, or “access” network; Flanche & Macy, 2011) and that within this network actors’ preferences for interaction change as cultural forms flow through the network. The authors find that modifications to this static network structure result in unique flow patterns for complex contagions in small world networks.

While such work is enlightening, Pachucki and Breiger (2010), in their recent review of work at the intersection of culture and social networks, rightfully argue that culture and network structure are cyclically and dynamically intertwined. They suggest that a causal, static network structure prevents a true understanding of how cultural forms develop. In contrast to the works mentioned above, we thus assume that the social network is dynamic and need not be explicated to study how social structure affects cultural tendencies. Instead, we rely on the principles of homophily and cognition to espouse or prevent new relations at the dyad level while maintaining a static *social group* structure. This assumption falls in line with macroscopic structural ideas portrayed in much of social psychology (a prominent example being Tajfel & Turner, 1979) and in sociology, perhaps most notably by Blau (1977).

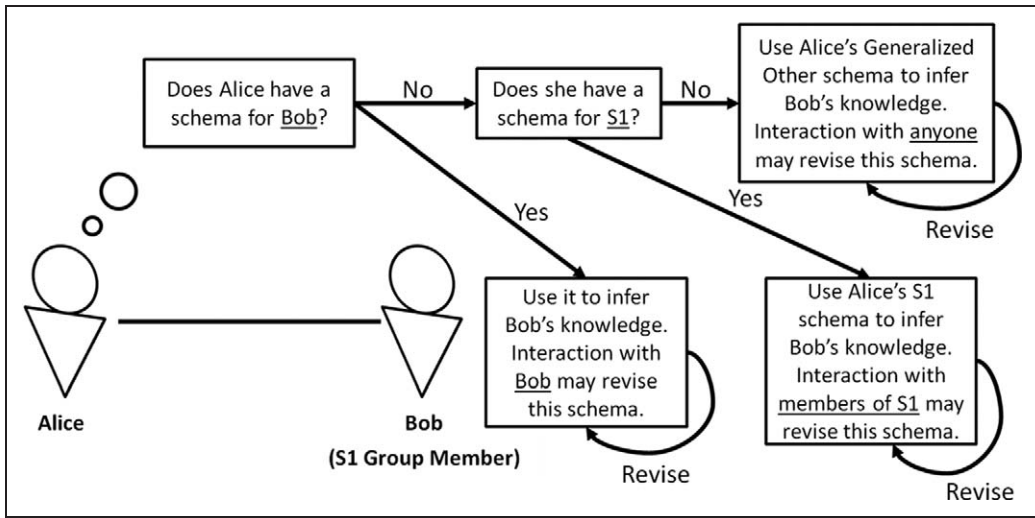
## Model

Like Soar (Laird, Newell, & Rosenbloom, 1987), our model is a knowledge-level model. It assumes all actors are privy to a set of knowledge, represented as a collection of bits (0s or 1s). Agents are initialized with a specific collection of bits by the modeler. In the present work, agents are also initialized to be in one of the four equally sized social groups.

After these initialization steps, the simulation proceeds in a turn-based fashion. At the beginning of each turn, each agent determines his probability of interacting with all others based on how similar he perceives their knowledge to be to his. Note that agents only consider bits set to 1 (knowledge that they know) when determining similarity. Therefore, an agent with a knowledge set of 1100 will believe they have a similarity of zero with an alter perceived to have a knowledge set of 0000. After determining a probability of interaction with all other agents, each agent then selects interaction partners and, via interaction, passes some knowledge he knows (bits that are set to 1) and receives knowledge his partners know. How agents determine the knowledge of others and what occurs in agent cognition as a result of an interaction are the core processes of the model. For full technical details, including all equations and further information on the algorithms involved, we refer the reader to Morgan et al. (n.d.).<sup>2</sup>

### *Model of Agent Cognition*

Agents are able to develop schemas at three different social “tiers”—the “specific other” (individual agents), the social group, and the generalized other. An agent determines what each other agent knows by using the most specific tier of his schematic representation of that agent. Thus, if Alice holds a schema of Bob as a specific other, Alice uses this schema to perceive what Bob’s knowledge is. If Alice has no specific other schema of Bob but both knows he is in social group S1 and has a schema for that group, she uses it to determine Bob’s knowledge. Finally, if Bob is not a specific



**Figure 1.** Alice needs to infer Bob's knowledge to determine her likelihood of interacting with him.

other and Alice does not have a schema for S1, Alice will construct her expectation of Bob's knowledge from her perception of the generalized other. This process is modeled in Figure 1, where Alice is attempting to form a perception of Bob's knowledge.

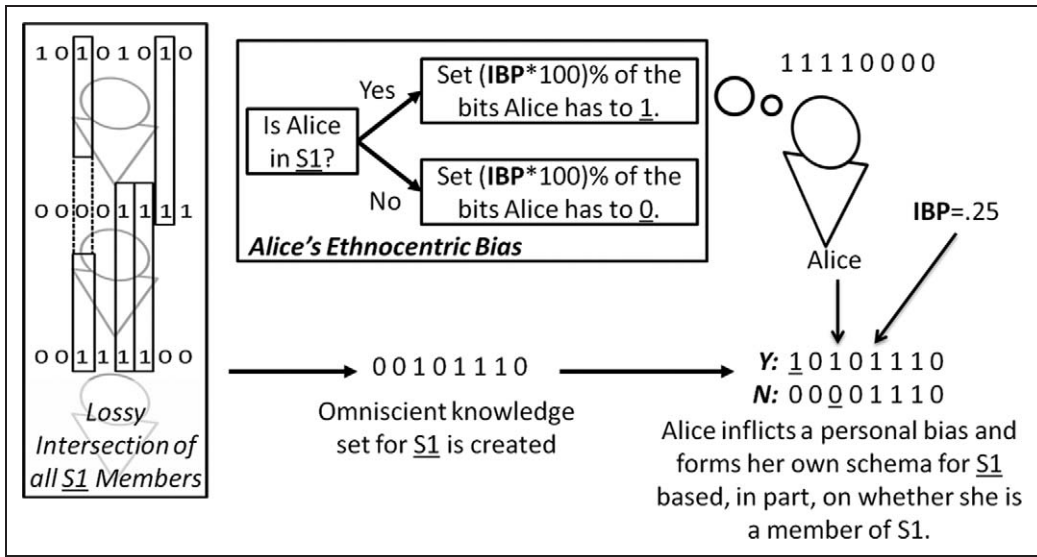
Agents begin the simulation with no specific other schemas, but do begin the simulation with two pieces of information about social groups. First, each agent is aware of all members of their own social group. Consequently, if both Bob and Alice were to be in the same social group, say S1, Alice would know Bob was in S1. However, agents can only learn the social groups of those outside their own via interaction. This model assumption assures that agents are not omniscient of macrosocial structure at the beginning of the simulation. Agents also begin with a schema of all social groups and a schema of the generalized other. These schemas are initialized through a two-step process, depicted in Figure 2.

First, an "omniscient" schema for each social group and the generalized other is constructed by considering the knowledge of all agents in the simulation. This is done using the principle of *lossy intersection*, by which a knowledge bit is set to 1 for a schema if a simple majority ( $50\% + 1$ ) of the group has that bit. This process is related to schema instantiation as described by Rumelhart (1978), where humans tend to perceive their environment as being representative of the "average" of each variable within a schema. From this omniscient schema, agents then each construct their own more or less biased schema of each social group.

The agent's schemas for social groups may be biased by ethnocentrism. The level at which an agent inflicts this ethnocentric bias is determined by the *initial bias parameter* (IBP). The IBP gives the probability that an agent will believe at the start of the simulation that any bit he knows is also known by everyone in his group, and similarly as the probability of his belief that any bit he knows is *not* known by members of *other* social groups. At a value of zero, agents therefore have no bias in their schemas for social groups—their initial schema is true to the omniscient schematic representation constructed via the initial lossy intersection procedure. At a value of 1, an agent's social group schemas depict his perception that all members of the social groups he is in know everything he knows, and all members of groups he is not in know nothing he knows.

It is important to note, however, that bias does not affect the generalized other, and thus agent perception of the generalized other fits any actual consensus that exists in the knowledge structure of the society. This decision was made because human perception of society has been discussed in





**Figure 2.** Alice constructs a schema for S1.

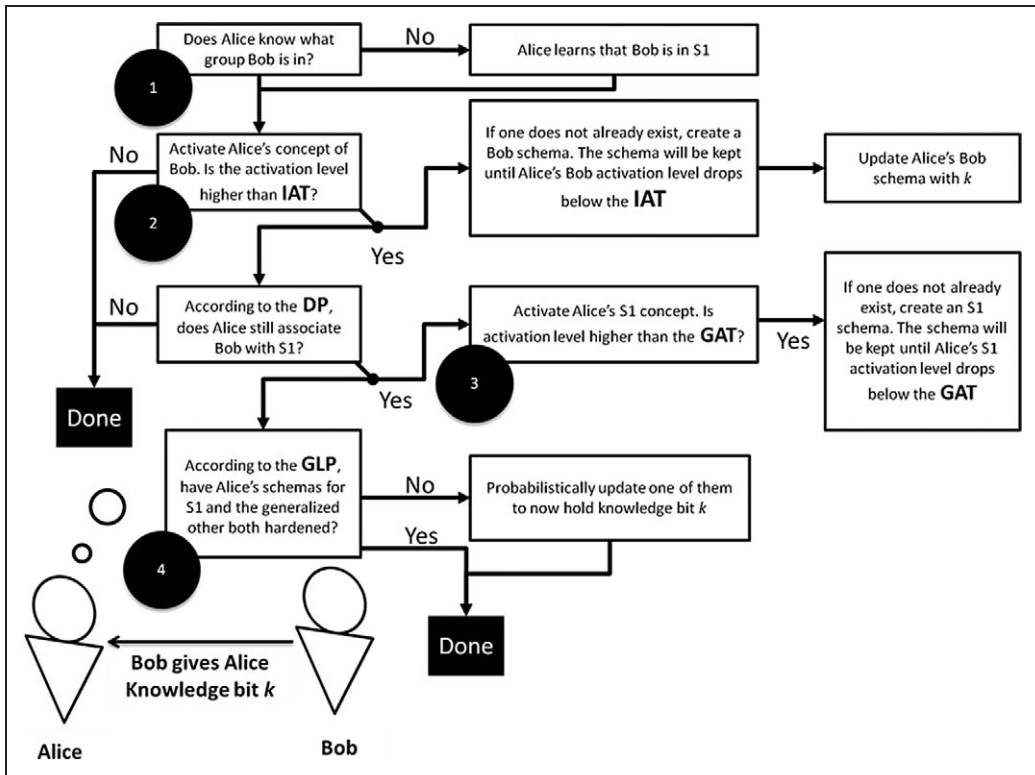
the related literature with respect to its nature as both an in-group (Turner, Hogg, Oakes, Reicher, & Wetherell, 1987) and an out-group (Mead, 1925). As will be seen, this modeling strategy had consequences on our results, and future work may examine alternative methods for initializing agent's generalized other schema.

### The Interaction Process

Interactions result in a series of four steps in our model of agent cognition. To explicate these steps, we will assume that an interaction has occurred between Alice and Bob, that Bob has passed Alice knowledge bit  $k$ , and that Bob is in social group S1. This interaction is depicted in the process model in Figure 3. In the figure, black circles represent the beginning of the four steps we will discuss, and model parameters are in large, boldface text. We will focus here on how Alice's cognition is affected by this interaction, but note that the same will be true for Bob as well.

The first step once an interaction has occurred is for Alice to learn the social group Bob is in (if she does not know it already). Via this information, Alice can use her schematic representation of S1, if she has one, when determining what she expects Bob to know later in the simulation. Note that this first step assumes an agent is able to perceive the group affiliation of all others in the society immediately and without error. After learning the social group Bob is in, the second step is for the level of activation for Alice's concept of Bob to increase. If the activation level of her concept of Bob reaches the *individual activation threshold* (IAT), her schema of him is instantiated. If Alice does not already have a schema of Bob to instantiate, she will construct one based on what she knows of S1 and the generalized other. If she does not have a schema for S1, then Alice relies entirely on her perception of the generalized other to initialize a schema for Bob. If Alice does have a schema for S1, she will construct her schema for Bob by taking bits from her schema of S1 and bits from her generalized other schema probabilistically based on the activation levels of these two schemas. Once instantiated, Alice's schema for Bob is then updated with the fact that Bob knows  $k$ .

It is important to recall that activation levels for concepts decay each turn—exponentially, as dictated by ACT-R's activation equations (Anderson, 2007). Therefore, if Alice does not interact with



**Figure 3.** The four-step process that occurs in Alice's mind when Bob gives Alice a piece of knowledge.

Bob for a long period of time, the activation level of her concept for him may eventually drop below the IAT. If this occurs, we assume the concept of Bob is too weak to cause an instantiation of Alice's schema for him. Though the concept retains its activation level, we assume that her schema for Bob is essentially "forgotten." Alice must therefore reconstruct a schema for Bob the next time her concept of him is activated above the IAT. This mechanism is done in part for computational efficiency, but also is grounded in empirical evidence (Morgan et al., n.d.).

The final two steps to occur upon interaction are relevant to the social group and generalized other schemas. Due to spreading activation, an activation of the concept of Bob in Alice's cognition may cause her concept of S1 to be activated as well. In order for this to occur, two conditions must be met. First, Alice must have a specific other schema for Bob—if not, we assume that the activation of the concept for Bob is too weak to cause a significant level of spreading activation. Second, Alice must still strongly associate Bob with S1. The likelihood of this association is controlled by an exponentially decaying likelihood function, which determines the length of time after Alice's current schema for Bob is constructed that she associates him with S1. The decay of this function can be modified by the *decategorization parameter* (DP), which controls the "half-life" of the function (i.e., the number of turns after which there is a 50% chance that an activation of Bob will lead to an activation of S1).

If activation of Alice's concept for S1 occurs and its activation level rises above the *group activation threshold* (GAT), Alice's schema for S1 is also instantiated due to her interaction with Bob. Like the IAT, the GAT also controls the activation level at which Alice will "forget" group schemas. If Alice's schema for Bob is to be instantiated and it does not exist, a schema will be constructed for S1. If a schema needs to be constructed, Alice forms a perception of S1 based on the



**Table 1.** A Tabular Description of the Virtual Experiment.

Parameters	Values Taken
Parameters of interest	
Initial knowledge distribution	Random, group based, all same
Initial bias parameter (IBP)	0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1
Other parameters varied	
Group activation threshold (GAT)	-5, -1
Group learning parameter (GLP)	5, 25, 50
Individual activation threshold (IAT)	-1, 0
Constants	
Number of simulation turns	150
Number of agents	1,000
Number of knowledge bits	500
Number of interactions	2
Number of knowledge bits passed per interaction	1
Density of knowledge (percentage of bits set to 1)	0.4
Decategorization parameter (DP)	6
Groups per agent	1
Total number of groups	4
Repetitions	
Number of repetitions	10
Total runs	$3 \times 11 \times 2 \times 3 \times 2 \times 10 = 3,960$

lossy intersection of all the specific other schemas she currently holds of individuals in S1. Thus, Alice uses exemplars to construct her perception. Note that the generalized other schema is never activated and that it also cannot reach an activation level lower than the GAT. We also make the assumption that agents never “forget” schemas of their own social groups, though the associated concepts can still be activated. However, it is important to note that other initial social group schemas are not considered to be activated, but rather are temporary structures used in annealing the simulation model to a more stable interaction structure. Thus, unless activated via interaction within the first 10 turns of the simulation, these initial group schemas will be forgotten.

Finally, if they have not already hardened, Alice’s schema of S1 or the generalized other can be updated with the belief that members of S1 (or society) are likely to know fact  $k$ . Only one of these two will be updated on any given interaction. This selection is based probabilistically on the level of activation of these schemas. Once a schema is selected, the *group learning parameter* (GLP) dictates the likelihood that the selected schema has hardened. The GLP specifies the number of turns since the initialization of a social group schema or a generalized other schema where agents are likely to accept new information about it. Note that if an agent “forgets” a group schema, the agent will again be able to learn about the group if the schema is reinitialized at a later time. The GLP is associated with the same mathematical function as the DP, and thus has an equivalent functional form (though, obviously, a different function).

## Experiment

Our virtual experiment is designed to understand how ethnocentric stereotypes and initial distributions of knowledge affect social connections across groups. Our focus is thus on how modifications to the initial distribution of knowledge to agents and changes to the IBP affect intergroup relations. Table 1 provides a concise representation of our virtual experiment. We hold constant the total number of social groups in the model, the number of groups per agent, the DP, and the density of

knowledge in addition to various other more standard model parameters. However, we do vary a number of parameters in order to test the resiliency of our findings to different assumptions about how agents learn and retain group schemas.

First, we vary the IAT, which in essence allows us to see how results differ when agents are able to retain more significant other schemas. We also vary the GAT. Increasing the GAT increases the level of activation required for a group concept to cause an instantiation of a group schema. Intuitively then, as the GAT increases, the salience of social groups in the society decreases—it takes more of a cognitive effort for agents to create and hold distinctions across different social groups, and thus agents are slightly more likely to engage in interaction across groups. We also modify the GLP. As the GLP increases, the number of simulation turns after construction before a social group schema or the generalized other schema hardens increases.

While interesting effects were observed due to modifications of the GLP and the GAT, conclusions for the primary questions of interest did not differ qualitatively across the knowledge conditions for different values of these parameters. We therefore chose, due to space constraints, not to discuss these results and focus instead only on the effects of the IBP and the different knowledge conditions, described below. We thus merged runs from all other parameters together for analysis.

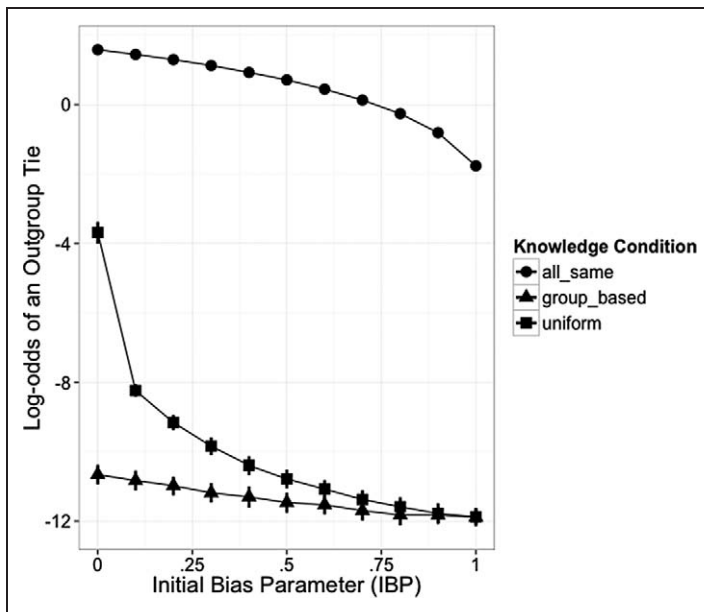
### *Initial Knowledge Distribution and the IBP*

We initialize the knowledge of agents in three different ways to understand the extent to which the underlying knowledge structure of the artificial society can serve to induce or prevent social relationships across groups. In the all same condition, all agents have exactly the same knowledge. Here, agents will, with no ethnocentrism, begin with the perception that all social groups and the generalized other have exactly the same knowledge as they themselves do. Agents therefore begin with an equal likelihood of interacting with anyone. In the second condition (the *random* condition), agent knowledge is distributed randomly—there is no alignment with other agents or within social groups. In this case, the agent will perceive the knowledge of social groups and society to be virtually null, and consequently find they have almost nothing in common with anyone.<sup>3</sup> In the random case, agents with no ethnocentric bias are therefore initially *unlikely* to interact with anyone.

When agents are assumed to have ethnocentric stereotypes (or equivalently, as the IBP is increased), however, initial interaction likelihoods change. In the all same condition, only the negative effect of the IBP on an agent's perception of groups he is not in will affect interaction patterns, as the agent will already perceive his own social group to have the same knowledge he has. In contrast, both positive stereotypes toward the agent's own groups and negative stereotypes toward other groups will affect interaction patterns in the random case.

In both the all same and uniform conditions, the social group an agent is in is merely a label—there is no connection between group and the underlying knowledge structure. These two conditions therefore test, in two different assumptions of underlying knowledge structure, how group labeling and ethnocentrism based on social groups alone lead to social structure. In order to provide a contrast, we use a final condition (the *group-based* condition) in which the true underlying knowledge structure is indeed aligned with social group structure. More specifically, the knowledge set is segmented into four, and each group is associated with an equal-sized segment. Agents are given 90% of the knowledge that their own social group is associated with. The dense concentration of knowledge to these social groups means that both learning and ethnocentrism reinforce stereotypical differences. Therefore, there should be little effect of increasing ethnocentric views with the group-based distribution, as any interactions will only serve to reinforce bias.

In all cases, the density of knowledge is held constant to ensure there is no bias from knowledge saturation. The density selected ensured that an interesting level of interaction occurred in all cases, but also implies that in the group-based condition, agents do hold 77 knowledge bits of the 375



**Figure 4.** The x-axis represents the 10 different IBP conditions and the 3 different shapes of points represent knowledge conditions. The y-axis gives the log odds of an out-group tie, and lines connect the mean outcomes across the different conditions. Ninety-five percent (95%) bootstrapped confidence intervals are drawn at each initial bias parameter (IBP) condition.

possible bits outside their social group. We hope to develop mechanisms to remove the need for this inconsistency in the future.

### Outcome Metric

Our outcome metric describes the extent to which social relations exist across social groups. Importantly, our interest is thus on social relations and not just on interactions between agents, because relations are thought to reflect more stable aspects of social structure. In order to determine a relation, however, we must obviously decide what constitutes a relationship. A common approach in both simulation (e.g., Hirshman et al., 2011) and empirical studies (e.g., Johnson, Kovács, & Vicsek, 2012) is to use a cutoff of interaction counts, where only agents having more than  $N$  interactions between them are considered to have a social relation. In the present work, we choose  $N = 2$  (the same used by Hirshman et al., 2011), providing us with social networks having a mean density of around .016 (around 16 ties per agent).<sup>4</sup>

Having determined what constitutes a social relation, we can now define our metric. We consider the ratio of relations formed between members of different social groups as compared to the number formed between members of the same social group. In order to obtain a value on a useful scale, we take the logarithm of this ratio, adding one to both the numerator and the denominator to avoid undefined values. Thus, the outcome statistic of interest in each simulation is calculated as  $\log_2 \left( \frac{\# \text{ relations connecting two agents in different groups} + 1}{\# \text{ relations connecting two agents in the same group} + 1} \right)$ , which defines the log odds of a relationship with a member of the out-group. We ignore data from the first 25 turns of the simulation, so that the model can anneal to a relatively stable interaction structure.

## Results

Figure 4 shows results across the three knowledge conditions for each value of the IBP. As one would expect, the odds of a relationship across social groups are several orders of magnitude larger when agents all hold the same knowledge. In fact, under this condition, the odds of a tie across groups are, on average, actually greater than the odds of a relationship within a group at all IBP values below 0.8. This is due in part to knowledge similarity being too strong to be entirely reduced by ethnocentrism. However, even at an IBP of 1, when agents believe they shared no knowledge at all with members of other social groups, odds of a relationship across groups are still significantly higher than in the other two knowledge conditions.

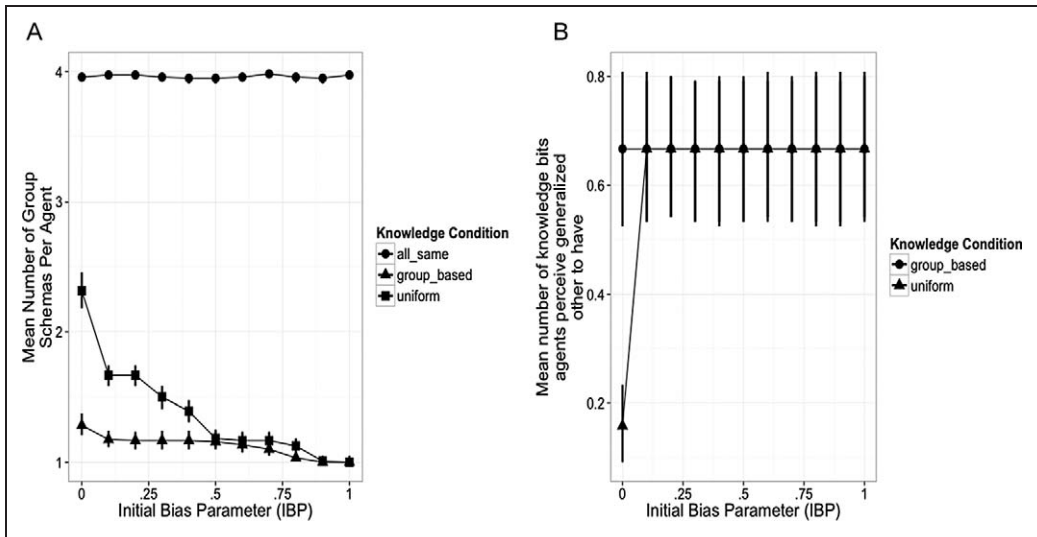
This is due to the agents' belief that the generalized other hold exactly the same knowledge as themselves. An agent uses his schema for the generalized other when determining the knowledge of agents whose social group he does not know. This amounts to any agent outside of his own social group he has not already interacted with, meaning there is a consistently large pool of other agents in different social groups with whom he believes he shares the exact same cultural form.

However, given that agents always have such a large pool of others outside their own social group to interact with regardless of the IBP, it is at first surprising that Figure 4 shows a strong effect of the IBP on the odds of an intergroup relationship in the all same condition. In fact, unlike the other two knowledge conditions, the effect of the IBP on intergroup relations actually increases as the IBP increases.

In order to understand this, we must recall that a single social interaction does not constitute a social relationship and that upon an initial interaction with an alter outside his social group, an agent learns the social group of the alter. Thus, after an initial interaction, an agent will use his biased social group schema to construct his specific other schema for the alter and to directly determine the knowledge of this alter if he forgets this specific other schema in the future. While an agent's generalized other schema pushes him toward interaction with other groups, ethnocentric perceptions therefore push him away from multiple interactions with the same person outside his group and, consequently, away from forming social relationships with them. Though multiple interactions with the same individual outside an agent's group are not impossible, increasing ethnocentrism leads agents to increasingly prefer interacting either with someone in their own social group or with someone whose social group they do not know.

Agents therefore consistently *interact* with those outside their groups even at higher levels of ethnocentric bias, but rarely form social relationships with them. This has two important effects on intergroup relations. First, because agents consistently interact with new individuals in different social groups, group concepts are continuously activated. Consequently, agents are never able to forget initial ethnocentrically biased schemas and thus are not able to reconstruct a more favorable stereotype of other social groups from exemplars. Here, then, interaction with agents in other groups actually promotes negative stereotypes through fleeting intergroup interactions. Second, spurious intergroup interactions actually detract from the number of within-group relationships that agents form. This occurs because the number of interactions agents have per turn is fixed, and agents "spend" many of these interactions in one-shot interactions with members of other social groups. These compounding effects are what likely lead to the increasing effect of ethnocentrism as the IBP increases.

In contrast to the all same condition, the effect of increasing the IBP remains relatively stable and subtle in the group-based condition. As noted previously, because knowledge is already aligned across groups, initial social group schemas in the group-based condition do not change significantly with increased ethnocentrism. This means agents hold a consistently strong preference toward interaction within their own social group regardless of the value of the IBP. This differs from the all same condition, where increasing the IBP gave agents an increasingly incorrect, negative stereotype of



**Figure 5.** (A) The mean number of group schemas that agents held across all conditions. (B) The mean number of knowledge bits that the generalized other schema had set to 1 across all agents in the group-based and uniform knowledge conditions only. Error bars are 95% bootstrapped confidence intervals.

other social groups. It also differs from the random condition, where increasing the IBP gave agent both the perception that their own group is more like them than it really is and the perception that other social groups are more unlike them than they really are.

In the random condition, even small levels of bias therefore lead to a strong and misplaced preference for social relationships within the agent’s own group. This explains why the effect of increasing the IBP actually slows in the random condition—after even a small amount of ethnocentrism is inflicted into initial group schemas, agents are pushed to interact almost exclusively with those in their own social group. Indeed, as Figure 4 shows, by an IBP of 0.6, the average likelihood of an intergroup relationship when knowledge is distributed completely randomly is within the 95% confidence interval of the case where knowledge is actually distributed according to social groups.

Agent distaste for interacting across social groups in the group-based and random conditions eventually leads them to forget most social group schemas aside from their own (which, as noted above, they are not able to forget). This is shown in Figure 5A, which gives the mean number of social group schemas held by any agent. The random and group-based conditions show that as bias increases, agents become less likely to have any group schemas except for their own. Figure 5A also supports our claim that agents rarely dropped group schemas in the all same condition.

Interaction with other social groups thus can only continue in the group-based and random conditions if agents perceive the generalized other to share something in common with them. Figure 5B shows the average number of knowledge bits agents believe the generalized other holds across these two conditions. As a point of comparison, agents in the all same condition associate a minimum of 200 knowledge bits to the generalized other. Figure 5B shows that agents in the group-based and random conditions hold generalized other schemas that lead them to believe “society is unlike me.” Because agents cannot construct a favorable view of society as a whole and can only learn new things about the generalized other for a set period of time before their stereotypes harden, interaction patterns quickly stabilize almost entirely to within-group interactions in these two knowledge conditions.

## Conclusion

In the present work, we develop an ABM that binds individual cognition to the dyad with the principle of activation via interaction and binds cognition to social structure via cognitive schemas that represent stereotypes of social groups and the generalized other. The model is situated within constructivism in the assumption that interaction drives the diffusion of knowledge and in the assumption that shared knowledge is the means by which homophily informs the likelihood of interaction between actors. We use this model to run a virtual experiment that examines how ethnocentrism based on social group affiliation and the true underlying distribution of knowledge affect the likelihood of social relationships across social groups.

We find that ethnocentrism and the underlying cultural structure of our artificial society have an interactive effect on intergroup relations, and consequently that one cannot hope to understand one without the other. Our model suggests that even under moderate conditions of ethnocentrism, the presence of a unifying cultural form may give rise to a highly integrated society. However, when an underlying cultural consensus exists, ethnocentric stereotypes are actually reinforced by intergroup interaction, which leads to a decrease in intergroup relations. When cultural consensus was manifest within social groups but not outside of them, ethnocentric stereotypes were in alignment with the true knowledge structure of society. Consequently, decreasing the level of ethnocentrism did little to increase the level of social relations across groups. Finally, when no cultural forms existed within groups or across society, we observed that even small levels of bias resulted in an exponential decrease in the level of intergroup relationships.

Our findings suggest two broad claims of sociological interest in the context of interventions promoting intergroup relations. First, we show that when cultural structure aligns with group structure, interventions aimed at altering social stereotypes alone will fail; rather, it is necessary to take aim at the dynamic, cultural forms within the society. In contrast, when a unifying cultural form (such as nationalism) already exists but is muted by ethnocentrism (e.g., via race), simply increasing the spread of cultural forms between groups is not always enough to mitigate ethnocentric stereotypes. Instead, intergroup relationships can only be built via the breakdown of ethnocentric stereotypes. Future work hopes to solidify these findings and to provide a stronger connection to similar empirical threads of research in the social psychology literature, most notably those stemming from contact theory (Allport, 1979).

While we believe the work we present to be of interest both practically and theoretically, there are a host of assumptions that may bias results. These include the assumption that the agent's perception of himself is perfect, that agents must interact with another before learning their social group, that agents and their environment are void of any affective or contextual cues that may instantiate particular schema, and that, once learned, agents never forget the social groups of others. In addition, we greatly simplify the complexities involved with schema theory and activation theory, simplifications that will hopefully be reduced in future work. We also simplify the decision of what constitutes a social connection to a simple "cutoff" value, which may not be appropriate to answer certain questions of sociological interest. Finally, we make assumptions and approximations of various functional forms that require further testing. In particular, the functional form of the two annealing functions requires a further discussion of their empirical derivation.

Though our work is set inside the empirically validated framework of the social simulation tool Construct (Carley & Hill, 2001; Schreiber & Carley, 2013), which has been used in a variety of settings itself (e.g., group mobilization; Carley, 1991; the impact of the printing press; Kaufer & Carley, 1993) and the theory of constructivism, the extensions provided here thus espouse new avenues of data collection and validation that are necessary. Beyond verification and validation, however, much interesting work remains to understand how more complex sociocultural structures lead to cross-cutting homogeneity (Blau, 1977). Modifications to the number of social groups in the



model, the number of social groups each agent is placed into, the ability of agents to retain specific other schemas, and the structure and transmission mechanisms of knowledge may all have interesting effects on the development of intergroup relationships. As such, we believe the present efforts to be only the beginning of an interesting avenue of computational research that can further our understanding of the unique interplays between interactions across the dyad, stereotypes of higher order social structures, and intergroup relationships in large social systems.

### Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

### Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported in part by the Air Force Office of Sponsored Research (AFOSR) with grant number FA9550-11-1-0179. Additional support was provided by CASOS—the center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University.

AQ4

### Notes

1. Though there have been attempts to apply adaptive control of thought–rational (ACT-R) to social situations (e.g., Kaulakis et al., 2012), these models are so computationally intensive that to utilize pure ACT-R agents would be infeasible if one wished to study even moderately sized artificial societies.
2. Additionally, for all code and executables used in the present work, we refer the reader to [https://github.com/kennyjoseph/sscr\\_14](https://github.com/kennyjoseph/sscr_14).
3. Recall that via the lossy intersection process, greater than half of the agents in a given group must know a given fact in order for the agent to associate the group schema as “knowing” that fact. Under the uniform knowledge condition, the number of agents knowing any given knowledge bit is binomially distributed. Thus, we can calculate the likelihood of a group of size  $N$  having more than half of its members know any one knowledge bit as. The odds of a group schema having any knowledge bit set to 1 at a group size of 250 are around 0.09%.
4. A cutoff of  $N = 1$  provided networks with a mean density greater than .1 (each agent has an average of 100 neighbors), which we considered to be too dense to find interesting structural properties. At  $N = 3$ , almost no social connections existed at all. In addition, we note that there were a small number of cases where no social connections occurred when  $N = 2$  at low biases in the random knowledge condition. We removed these conditions from analyses, but qualitative conclusions did not change when including them.

### References

AQ5

- Allport, G. W. (1979). *The nature of prejudice*. Basic Books.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford, England: Oxford University Press.
- Anderson, J. R., Matessa, M., & Lebiere, C. (1997). ACT-R: A theory of higher level cognition and its relation to visual attention. *Human-Computer Interaction, 12*, 439–462.
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological science, 2*, 396–408.
- Axelrod, R. (1997). The dissemination of culture a model with local convergence and global polarization. *Journal of Conflict Resolution, 41*, 203–226.
- Blau, P. (1977). *Inequality and heterogeneity: A primitive theory of social structure*. New York, NY: Free Press.
- Carley, K. (1991). A theory of group stability. *American Sociological Review, 56*, 331–354.

- Carley, K. M., & Hill, V. (2001). Structural change and learning within organizations. In A. Lomi (Ed.), *Dynamics of organizations: Computational modeling and organizational theories* (pp. 63–92). Live Oak, CA: MIT Press/AAAI Press.
- Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology*, *113*, 702–734.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological review*, *82*, 407.
- Dovidio, J. F., & Gaertner, S. L. (2010). Intergroup bias. In S. T. Fiske, D. Gilbert, & G. Lindzey (Eds.), *Handbook of social psychology* (pp. 1084–1121). New York, NY: John Wiley.
- Flanche, A., & Macy, M. W. (2011). Small worlds and cultural polarization. *The Journal of Mathematical Sociology*, *35*, 146–176.
- Garcia-Marques, L., & Mackie, D. M. (1999). The impact of stereotype-incongruent information on perceived group variability and stereotype change. *Journal of Personality and Social Psychology*, *77*, 979–990.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, *102*, 4–27.
- Gregg, A. P., Seibt, B., & Banaji, M. R. (2006). Easier done than undone: Asymmetry in the malleability of implicit preferences. *Journal of personality and social psychology*, *90*, 1.
- Hales, D. (1998). Stereotyping, groups and cultural evolution: A case of “second order emergence”? In J. S. Sichman, R. Conte, & N. Gilbert (Eds.), *Multi-agent systems and agent-based simulation* (pp. 140–155). Berlin, Germany: Springer-Verlag.
- Hamilton, D. L., & Mackie, D. M. (1990). Specificity and generality in the nature and use of stereotypes. *Advances in social cognition*, *3*, 99–110.
- Hartshorn, M., Kaznatcheev, A., & Shultz, T. (2012). The evolutionary dominance of ethnocentric cooperation. *Journal of Artificial Societies and Social Simulation*, *16*, 7.
- Hilton, J. L., & von Hippel, W. (1996). Stereotypes. *Annual Review of Psychology*, *47*, 237–271.
- Hirshman, B. R., Charles, J. S., & Carley, K. M. (2011). Leaving us in tiers: Can homophily be used to generate tiering effects? *Computational and Mathematical Organization Theory*, *17*, 318–343.
- Johnson, R., Kovács, B., & Vicsek, A. (2012). A comparison of email networks and off-line social networks: A study of a medium-sized bank. *Social Networks*, *34*, 462–469.
- Kaufer, D. S., & Carley, K. M. (1993). *Communication at a distance the influence of print on sociocultural organization and change*. Hillsdale, NJ: Lawrence Erlbaum.
- Kaulakis, R., Zhao, C., Morgan, J. H., Hiam, J. W., Sanford, J. P., & Ritter, F. E. (2012). Defining factors of interest for large-scale socio-cognitive simulations. In N. Rußwinkel, U. Drewitz, & H. van Rijn (Eds.), *Proceedings of ICCM-2012-Eleventh International Conference on Cognitive Modeling* (pp. 117–122). Berlin, Germany: Universitaetsverlag der TU Berlin.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial intelligence*, *33*, 1–64.
- Lazarsfeld, P., & Merton, R. (1954). Friendship as a social process: A substantive and methodological analysis. In M. Berger, T. Abel, & H. Charles (Eds.), *Freedom and control in modern society* (pp. 18–66). New York, NY: Van Nostrand.
- Lizardo, O. (2006). How cultural tastes shape personal networks. *American Sociological Review*, *71*, 778–807.
- McPherson, M., Lovin, L., & Cook, J. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*, 415–444.
- Mead, G. H. (1925). The genesis of the self and social control. *International Journal of Ethics*, *35*, 251–277.
- Morgan, G. P., Joseph, K., & Carley, K. M. (n.d.). Multi-level construct. *In submission*.
- Pachucki, M. A., & Breiger, R. L. (2010). Cultural holes: Beyond relationality in social networks and culture. *Annual Review of Sociology*, *36*, 205–224.
- Rumelhart, D. E. (1978). *Schemata: The building blocks of cognition*. San Diego: Center for Human Information Processing, University of California.

- Schaller, M., & Latané, B. (1996). Dynamic social impact and the evolution of social representations: A natural history of stereotypes. *Journal of Communication, 46*, 64–77.
- Schreiber, C., & Carley, K. M. (2013). Validating agent interactions in construct against empirical communication networks using the calibrated grounding technique. *IEEE Transactions on Systems, Man, and Cybernetics: Systems, 43*, 208–214.
- Smith, E. R., & Zárate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review, 99*, 3–21.
- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W Austin & S. Worche (Eds.), *The social psychology of intergroup relations* (pp. 33–47). Monterey, CA: Brooks/Cole.
- Turner, J. C., Hogg, M. A., Oakes, P. J., Reicher, S. D., & Wetherell, M. S. (1987). *Rediscovering the social group: A self-categorization theory*. Basil Blackwell.
- Wilder, D. A. (1986). Social categorization: Implications for creation and reduction of intergroup bias. *Advances in Experimental Social Psychology, 19*, 291–355.

AQ7

### Author Biographies

**Kenneth Joseph** is a graduate student in the Computation, Organization and Society program at Carnegie Mellon University. He focuses on agent-based modeling, social cognition, and social media; e-mail: [kjoseph@cs.cmu.edu](mailto:kjoseph@cs.cmu.edu).

**Geoffrey Morgan** is a graduate student in the Computation, Organization and Society program at Carnegie Mellon University. He focuses on agent-based models of organizational systems and organizational research; e-mail: [gmorgan@cs.cmu.edu](mailto:gmorgan@cs.cmu.edu).

**Michael K. Martin** is a project scientist at the Center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University. He received his PhD in cognitive and experimental psychology from the University of Kansas; e-mail: [mkmartin@andrew.cmu.edu](mailto:mkmartin@andrew.cmu.edu).

**Kathleen M. Carley** is the director of the Center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University. She has published on a plethora of topics related to agent-based modeling, cognition, and the dynamics of social networks; e-mail: [kathleen.carley@cs.cmu.edu](mailto:kathleen.carley@cs.cmu.edu).