Modeling the emergence of shared attitudes from group dynamics using an agent-based model of social comparison theory.

Keywords: agent-based modeling, connectionism, social comparison theory, group dynamics, polarization.

ABSTRACT

We propose a novel agent-based implementation of Festingers' Social Comparison Theory (SCT). The Social Comparison Model (SCM) consists of connectionist networks that simulate agent-level social comparison processes. Agent networks are combined into a adaptive network structure that is shaped by social comparisons between individual agents. Simulations show how the SCM produces behavior consistent with the empirical literature on group dynamics. In addition, experimental results are reported that show how the SCM can simulate how critical and conformist norms affect interpersonal processes and emergent attitudes. We conclude that the coupling of simulations and experiments, and the use of psychologically plausible agent models within adaptive network structures, can provide new impetus to the development of models of individual and social cognition. An integrated framework such as the SCM allows investigating key theoretical predictions around the origin and maintenance of socially shared information through social comparisons in fundamentally novel ways.

1. Introduction

An extensive literature in social psychology and sociology has established that shared norms emerge from interpersonal interactions, during which individuals reduce uncertainty and conflict by adapting their individual attitudes and opinions. Two key figures in this research are Sherif (Sherif, 1935) and Festinger (Festinger, 1954). Sherif demonstrated empirically how small groups of individuals tend to develop shared norms about the features of ambiguous phenomena. These studies inspired Festinger to develop a theory of norm formation, which lead to the development of Social Comparison Theory (SCT). The theory postulates that individuals constantly compare their opinions, attitudes and believes to those of others. These social comparisons are seen as a pervasive feature of social interaction, and central to the understanding of how groups develop socially validated information. Social comparison theory thus provides a theoretical account of how interpersonal interactions produce shared attitudes and consensus.

In this paper, we introduce a multi-agent model based on SCT. We take from the theory the guiding principle behind our model, which is that people judge the appropriateness of their behaviors or opinions by comparing them to those of other people around them. We will illustrate how a computational model can aid in the further development of social comparison theory by taking into account both individual-level psychological processes and interpersonal processes. First, we start by briefly describing social comparison theory. We then introduce the model, and illustrate how the implementation of SCT agent principles produces social group behavior that is consistent with empirical findings. We then present the results of an experiment in which a number of key hypotheses of the model were tested. We conclude by comparing the model to other simulation work, and by discussing future directions for research.

2. Social comparison theory

Social comparison theory is one of the earliest and most influential theories of how socially shared knowledge is developed and maintained through social interaction (Festinger, 1954). It has been widely acknowledged that the process of social comparison is a central feature of human social life (Goethals & Darley, 1977), and it has been shown to underlie, for instance, the phenomenon of group polarization (Axelrod, 1997). It has even been argued that the need to compare the self with others has evolved in other species as well, as an general adaptive mechanism for sizing up one's competitors (Gilbert, Price, & Allan, 1995).

The main principle behind the original theory is that people use social reality - in the form of opinions, behavior and beliefs of other people around them - to judge the appropriateness of their own behaviors or opinions, and this particularly when physical reality fails to provide sufficient cues. Probably the most well known example of this is the work by Sherif (1935) on the development of social norms. Using the auto-kinetic effect (perceived motion of a stationary light in a dark room), Sherif showed that participants' judgments about how far a light had moved, converged over time. Although very simple, his experiments illustrate a key axiom of SCT - People tend to *assimilate*: "The drive for self evaluation is a force acting on persons to belong to groups, to associate with others. People, then, tend to move into groups which, in their judgment, hold opinions which agree with their own" (Festinger, 1954). And "The existence of a discrepancy in a group with respect to opinions or abilities will lead to action on the part of members of that group to reduce the discrepancy" (Festinger, 1954). These behavioral principles lead to the development of homogeneous groups, like in the Sherif experiment – groups that organize around a shared norm.

Subsequent research established that the tendency to assimilate decreases as the difference between opinions increases. In fact, when the difference exceeds a certain threshold, individuals will strive to further distinguish themselves from others - this is referred to as *contrast*. A variety of socio-cognitive factors have been shown to affect the position of the threshold at which a drive towards assimilation is replaced with one towards contrast, from motivational factors (i.e. norms) to factors related to personal and social identity processes (for an overview, see Mussweiler, 2003).

Social comparison theory can potentially provide a parsimonious theoretical framework to study the socio-cognitive processes through which shared information is produced and maintained in social groups. However, many aspects of the theory are couched in ambiguous verbal descriptions. This, in combination with the inherent complexity of group dynamics, has made it difficult to directly test the predicted impact of repeated, individual social comparisons in groups, and particularly how they contribute to group dynamics and the validation of shared information. We will now describe a connectionist, agent-based model that aims to overcome this, by providing an algorithmic implementation of the key principles underlying social comparison theory.

3. The multi-agent system

Agent-based models (ABM) build social structures from the "bottom-up", by simulating individuals with virtual agents and stipulating rules that govern interactions among these agents. They express in clear mathematical and computational terms, how complex social structures emerge from interactions of individual agents at various distinct levels, allowing the analysis of properties of individual agents (e.g. their attributes and interactions), and emergent group-level behavior. However, human social groups change not only through structural adaptations (i.e. social organization), but also by guiding and restructuring the behaviors and cognitions of the individuals that form them. To that extent, several modelers (Resnick, 1994; Sallach, 2003) have argued that ABM need to incorporate relatively sophisticated models of individual agents, to allow them to adapt and change their behavior over time.

In this paper, an ABM is introduced that aims at accomplishing this by implementing social comparison processes in a connectionist agent model. Connectionism is an approach in the fields of artificial intelligence, psychology, neuroscience and philosophy of mind, that models mental and behavioral phenomena as the emergent processes of interconnected networks of simple units. Connectionist architectures and processing mechanisms are based on analogies with properties of the human brain, in which learning is conceptualized as a process of on-line adaptation of existing knowledge to novel information provided by the environment. The focus in this paper will be on the recurrent auto-associator (McClelland & Rumelhart, 1988; McClelland & Rumelhart, 1985), a model that has been applied successfully to group biases, causal attribution, person and group impression in social psychology (Smith & Conrey, 2007; Smith & DeCoster, 1998; Van Rooy, Van Overwalle, Vanhoomissen, Labiouse, & French, 2003).

3.1 Agents

The recurrent network used to simulate an agent has three distinctive features (see Figure 1). First, all units within an individual agent network are interconnected, such that all units send out and receive activation. Second, an agents' state is represented as levels of activation of its units. Information arriving at an agent is represented by *external activation* of its units. This is automatically spread among the interconnected units within an agent in proportion to the weights of their connections, resulting in an *internal activation*. These are summed to give the *net activation*, which reflects the short-term memory of the network. The third distinctive feature is that short-term activations are stored in long-term changes in the strength (weight) of the connections. Typically, activations and weights have lower and upper bounds of approximately –1 and +1.

Insert Figure 1: Recurrent agent network

In a recurrent network, processing information takes place in two phases. During the first phase, the activation of each unit in the network is determined. Each unit receives a stimulus from external sources, termed ext_i, which generates an external activation ext_a_i in proportion to an excitation parameter E, or

$$ext\_a_i = E \times ext_i \tag{1}$$

This activation subsequently spreads through the auto-associative network, meaning every unit i receives internal activation int_a_i, which is the sum of the activation from the other units j in proportion to the weight of their connection to unit i, or

$$int\_a_i = \sum_j (w_{ji} \times ext\_a_j) \tag{2}$$

External and internal activation are then summed to the net activation, or

$$net\_a_i = E \times (ext\_a_i + int\_a_i) \tag{3}$$

The updating of activation at each cycle is governed by the following equation:

$$a_i = net\_a_i + (1 - D) \times a_{i-previous} \tag{4}$$

where D reflects a memory decay term. As in previous simulations (D. Van Rooy et al., 2003), parameter values were set to $D = E = 1$. Hence, the final activation of a unit equals the sum of the external and internal activation, or

$$a_i = net\_a_i = ext\_a_i + int\_a_i \tag{5}$$

After activation has been determined, the recurrent model enters the second, learning phase in which the short-term activations are stored in long-term weight changes of the connections. Basically, these weight changes are driven by the difference between the internal activation received from other units in an agents network, and the external activation received from outside sources. This difference, also called the error, is reduced in proportion to the learning rate that determines how fast the network changes its weights and learns. This error-reducing mechanism is known as the delta algorithm (McClelland & Rumelhart, 1988; McClelland & Rumelhart, 1985). In mathematical terms, the delta algorithm strives to match the internal predictions of the network int_a_i as closely as possible to the actual state of the external environment ext_a_i and stores this information in the connection weights. This error-reducing process is formally expressed as

$$\Delta w_{ji} = \varepsilon \times (ext\_a_i - int\_a_i) \times a_j \tag{6}$$

where Δw_{ji} is the change in the weight of the connection from unit j to i, and ε is a learning rate that determines how fast the network learns (McClelland & Rumelhart, 1985). An implication of this learning algorithm is that when an object and its features co-occur frequently, then their connection weight gradually increases to eventually reach an asymptotic value of +1.

3.2. Interaction

A number of authors have illustrated how auto-associative networks can be naturally extended to allow interaction or communication between them (Hazlehurst & Hutchins, 1998). It basically involves creating an agent-based model such that individual recurrent networks or agents are linked in an adaptive network structure. Any agent can (in principle) interact with any other agent, but the impact of the interaction will adapt to experience (Hazlehurst & Hutchins, 1998; Van Overwalle & Heylighen, 2006). In the current simulation, communication involves the transmission of information from one agent network to another, along connections whose adaptive weights reflect the mutual social influence between agents. During a simulated interaction, receiving agents *compare* their information (as represented by internal activation of their own network) with the information from sending agents (represented by the external activation received from sending agents). Consistent with the social comparison framework, we will refer to the receiving agent as the *Comparer*, and the sending agent as the *Target* (see Figure 2). The comparer agent *c* sums all information on each issue *i* from all other target agents *t* in proportion to the inter-agent weights, and then processes this information internally (according to the standard recurrent approach). Or in mathematical terms, for each feature *i,*

$$ext\_a_{ci} = \sum_j w_{ti \to ci} * a_{ti} \tag{7}$$

where ext_a_{ci} represents the external activation received by comparer agent *c* on feature *i*; $w_{ti \to ci}$ is the inter-agent weight from Target agent *i* to the Comparer agent *c*; and a_{ti} denotes the final activation (which combines the external and internal activation received) expressed by the Target agent *t*.

INSERT Figure 2: Interaction between Comparer and Target agent.

3.3 Social comparison process

A number of Festingers' axioms were implemented directly into the agent model (Festinger, 1954; Fridman & Kaminka, 2011). The first principle is that agents compare their state to that of other agents. Second, when a communicated position falls within an agent's latitude of acceptance, an agent will strive to reduce the social difference with the object of comparison (*assimilation*). As pointed out earlier, research has also shown that, under many circumstances, individuals that disagree on a particular issue, or have very different attitudes, will seek to distance themselves. This is reflected in the third principle: When a communicated position falls outside an agent's latitude of acceptance, an agent will strive to increase the social difference with the object of comparison (*contrast*).

These principles can be implemented as natural extensions of the connectionist framework introduced above. Any social comparison process involves a Comparing and Target agent. The Target→Comparer weights are updated driven by the error between the external information provided by the Target, representing the attitude of the Target on a particular feature *i*, and the internal activation of the Comparer, representing the attitude of the comparing agent on the relevant feature:

$$\delta_i = ext\_a_{ci} - int\_a_{ci} \tag{8}$$

where ext_a_{ci} is the final activation send out by the Target to the Comparer agent, and int_a_{ci} is the internal activation generated by the Comparer agent. When the discrepancy between these 2 values falls within the Comparer agent's latitude of acceptance, the agent will strive to reduce the social difference with the Target (*assimilation*). When it exceeds this threshold, the Comparer will strive to increase the social difference with the Target (*contrast*). This is expressed mathematically as:

$$\text{If } |\delta_i| < Latitude$$
$$\text{then } \Delta w_{ti \to ci} = \eta \times (1 - w_{ti \to ci}) \times |a_{ti}| \text{ [Assimilation]}$$
$$\text{else } \Delta w_{ti \to ci} = \eta \times (0 - w_{ti \to ci}) \times |a_{ti}| \text{ [Contrast]} \tag{9}$$

where η is the rate by which the weights are adjusted, and $|a_{ti}|$ is the absolute value of the activation on issue *i* by the Target agent *t*. Formula 9 describes the corrective action of the comparing agent on the selected feature. If the discrepancy falls within the latitude threshold, this action involves assimilation, otherwise contrast.

This constitutes an adaptive social process, in which agents learn from interacting with each other: Agents that consistently express attitudes that fall within each other's latitude, will gradually develop stronger mutual links. The social experience acquired in this way is represented in a distributed manner, in patterns of weighted links across the whole network. In combination with formula 7, this means that agents that share similar attitudes consistently will gradually have greater influence on each other. As such, a group of agents functions as an adaptive network in which information and knowledge are distributed among and propagated in function of the social influence between different individual networks.

4. Model hypotheses

The use of a connectionist agent model is one of the features that distinguishes the current approach from previous ABM. The connectionist learning and processing algorithms underlying the agent model are based on analogies with properties of the human brain (Van Rooy et al. 2003; Smith & DeCoster, 1998), and simulate learning as a process of online adaptation of existing knowledge to novel information provided by the environment. As such, the agent model provides a degree of psychological plausibility that is often missing from previous ABM.

In the following simulations, we will explore the basic dynamics of the model in terms of agent learning and social interaction. We will use a very simple design, based on classic group perception experiments that have received both extensive modeling (Van Rooy et al., 2003) and empirical attention (Van Rooy, Vanhoomissen, & Van Overwalle, 2013). Remarkably, most of this literature has focused almost entirely on the *individual* process of attitude formation, treating human perceivers as essentially socially isolated. We will illustrate how the SCM can be used to generate predictions both about individual, or agent learning, and the impact of social interaction. The outcome of these simulations, and the values of key model parameters are subsequently further validated in a small group experiment with human participants.

4.1 Simulation 1: Agent behavior

In typical group perception experiments, participants read a series of positive and negative behavioral statements about members of a fictional social group (Hamilton & Gifford, 1976; Van Rooy et al., 2003; Van Rooy et al., 2013), such as: "John, member of group A, helps an old lady across the street", or "Peter, member of Group A, was caught littering". After receiving a list of such statements, participants are then usually asked to report their attitudes towards the group.

To illustrate the basic behavior of the agent model, we start with simulating a simplified version of this design involving one agent network, represented in Table 1 and Figure 3. The network contains a group node, representing the social group ("A"), and two valence nodes, representing positive ("+") and negative ("-") attributes of the group. This is a so-called localist encoding scheme, in which each piece of information (or concept) is represented by a single node (for a similar approach, see Van Rooy et al., 2003; 2013).

Two separate unitary valence nodes were taken, rather than a bipolar attribute node, as research shows that evaluations about groups are cognitively represented as mixed and complex constructs including both positive and negative instances of the attribute, rather than a single point on a one-dimensional construct (Wittenbrink, Judd, & Park, 2001).

INSERT TABLE 1 HERE

An agent is trained by presenting it with a series of patterns in which the group unit is activated together with either the positive or negative valence unit. Table 1 shows a simulated learning history: Each row represents a pattern of external activation at a trial that corresponds to a statement presented to a participant. Depending on the type of information represented by a pattern, respective units are either turned on (activation level= 1) or turned off (activation level= 0). For instance, the first row {A+} could represent the statement "John, member of group A, helps an old lady across the street". The # indicates the number of times each pattern of activation is presented. For instance, this Agent was trained on 9 {A+} and 4 {A-} patterns, which simulates a condition in an experiment in which a group of participants is presented with 9 positive and 4 negative statements about members of group A.

Insert Figure 3: Learning histories for 5 agents (top left), resulting in different simulated attitudes towards "Group A". The Figure shows the development of the simulated attitudes in function of amount of learning.

Figure 3 shows the learning process of a number of different agent networks. Their learning regime is indicated in the top left of the Figure – for instance, Agent 5 was trained on 9 {A+} and 4 {A-} patterns. A series of gradual, trial-by-trial adjustments by the learning algorithm, results in a certain configuration of connection weights in the network. This configuration determines how activation flows through the network and activates related concepts. To test the knowledge embedded in these connections, we apply a procedure analogous to when participants are cued with questions on the experimental stimulus material learned previously. More particularly, some concepts in the network serve as a cue to activate related material in the network. In this case, the group label served as a cue to estimate group attributes (positive or negative valence), by turning the activation of the cue on to +1. The difference between the final activation between the 2 attribute nodes is then taken as the agent's attitude towards Group A. Figure 3 shows how different learning histories result in different attitudes.

Although a relatively simple set up, this agent model has been shown to be capable of simulating a number of key social psychological phenomena from group perception (Smith & DeCoster, 1998; Van Rooy et al., 2003). This is due to a number of its learning properties, like it's sensitivity to the ratio of positive to negative information, which is apparent from the figure above. In the next simulations, we will embed this agent model into an extended network and explore the ability of the model to simulate inter-personal processes.

4.2 Simulation 2: Social interaction.

As mentioned, Social Comparison Theory deals with shared norm formation. It argues that individuals develop socially shared knowledge through repeated social comparisons, during which privately held opinions are adapted and validated. In the following simulation, we explored the group dynamics that emerge from the repeated application of the social comparison process implemented in our model (formula 8-9). More particularly, we simulated 10 agents, using the same agent network as in simulation 1. To simulate a certain degree of diversity, there are 3 different "cliques" of agents. Agents 1-5 received patterns that were strongly positive, agents 6-7 received neutral information (i.e. a mix of positive and negative patterns), and finally agents 8-10 received information containing mostly negative patterns. Interaction step 1 in Figure 4 shows the resulting attitudes with which these agents start.

Subsequently, agents interact and exchange their impressions with all other agents. Interaction is simulated in a 2-step process. First, a target agent is made to express it's opinions about group A – similar to the training process in Simulation 1, the target agent is presented with stimuli (external activation) {A+} and {A-} in turn. For each such stimulus, comparer agents receive external activation at their corresponding nodes ({A+} or {A-} respectively). Internal weights for comparer agents and external weights from target to comparer for {A+} or {A-} nodes as appropriate are updated, however target agents internal weights are not yet updated at this stage. At each interaction step in Figure 4, each agent interacted twice (ie: twice {A+} and twice {A-}) with every other agent. Agent attitudes were measured after each step.

The left panel in Figure 4 shows that, after a few interaction steps, agents start to self-organize around a "majority" impression of the target group. This pattern further strengthens over the following interaction rounds, as the 3 agents with the "minority" impression of the group gradually shift towards the other agents that hold the "majority" impression (i.e. they assimilate). In fact, agents not only organize into a homogenous group, they reinforce each other's information and end up with stronger attitudes towards the group than the ones with which they started. This corresponds to group polarization, whereby members of a group on average shift their opinion toward a more extreme position during group discussion. For this simulation, the latitude parameter was set to a value of .01. Increasing the value of this parameter leads to less polarization, and instead all agents in the group converge to the middle position (see right panel of Figure 4). This is line with the literature: Polarization, or a shift of the minority to a majority position, is more likely in a setting where individuals show little tolerance for deviance (i.e. low latitude) (Isenberg, 1986).

INSERT Figure 4: Development of attitudes towards Group A of 10 agents in function of interaction steps. Latitude parameter set to .01. (left panel a) and 0.4 (right panel). At each step, each agent interacted twice with all other agents.

This simulation illustrates an important group dynamic through which real social groups create, validate and maintain socially shared knowledge (Goethals & Darley, 1977; Turner, 1987): Agents organize themselves around a shared norm or attitude, and with every interaction, the connection weights between them increase, reflecting increased social influence. The gradually strengthening links between agents act as positive feedback loops that further reinforce attitudes. This produces the polarization effect, in which all agents end up with more extreme opinions after the interaction. However, it is obvious that groups of individuals not always assimilate, and group dynamics do not always produce polarization. When agents show no latitude whatsoever, opinions are obviously not adapted and remain as they are. Perhaps more interestingly, we can also explore the impact of the amount of information exchange between the different "cliques" of agents.

Figure 5: Development of attitudes towards Group A of 10 agents in function of amount of interaction. Within each clique, agents talk to each other three times, before talking to all other agents once. Latitude is low (L= 0.01).

Figure 5 shows that if we reduce the exchange of information between the cliques, while increasing the communication within them, the amount of polarization and attitude change is greatly reduced. This is consistent with previous simulation work (Axelrod, 1997) and the empirical literature (for an overview, see Isenberg, 1986) that shows that the amount of attitude change, and the willingness to adopt or conform to another point of view, is driven largely by the number of arguments that are exchanged between opinion groups during interaction. Interestingly, and in line with Festingers' theory, the different cliques self-organize around their shared attitude ("ingroups" versus "outgroups"), and maintain their social distance from each other.

5. Small group experiment

As mentioned in the introduction, a number of factors have been identified that affect the amount of latitude individuals will display towards information from others. A good example of this are norm manipulations that amplify the motivation among group members to deliberately and systematically process the information that becomes available during group discussion. For instance, Postmes and colleagues (Postmes, Spears, & Cihangir, 2001) conducted a number of studies in which they induced a "criticality" group norm in some groups, whereas other groups were encouraged to reach consensus. Groups with a criticality norm took more consideration of unshared information and reached more high-quality decisions. The same result has been found under a variety of experimental conditions (Kelly & Karau, 1999; Scholten, van Knippenberg, Nijstad, & De Dreu, 2007). In our own work, we have demonstrated how a similar norm manipulation encouraged a group of software engineers to be more accepting of information coming from others in a critical compared to conformist mindset (Teh, Baniassad, Van Rooy, & Boughton, 2012). In sum, this work shows that, perhaps somewhat paradoxically, a critical mindset encourages individuals to show more latitude towards the information coming from others, and to take information that might be inconsistent with their own views more into consideration than a conformist mindset.

We designed an experiment to specifically test whether the Latitude parameter can model the impact of norms on interpersonal processes. As mentioned, the Latitude parameter determines the discrepancy agents allow when comparing their information (see Formula 9). Based on the literature and our own research, our hypothesis is that low latitude would correspond to a conformist mindset, where participants are not motivated to process information that might be inconsistent with their own. High latitude would then correspond to the critical norm condition, where participants are motivated to process information from other participants, even it conflicts with their own.

5.1 Design and procedure

Eighty psychology undergraduate students (32 men, 48 women; mean age = 21.3) participated in the study. Participants arrived in the lab and were informed that they would be receiving information about individuals who belonged to a particular social group, labeled "Group A". They were asked to form an impression of this group, and told that they would afterwards be able to share their impressions with other participants.

The procedure largely followed a standard group perception design (Hamilton & Gifford, 1976; Van Rooy et al., 2013), in which participants receive information about a fictional social group. In Phase 1 of the experiment, individual participants were exposed to sentences describing fictitious people engaging in various behaviors. Some of these behaviors were positive in nature ("J., member of group A, helped an old lady across the street."), others negative (e.g., "refused to assist an old person who was lost"). These sentences were developed as part of larger study (Van Rooy et al. 2013), and tested in a pilot study to make sure that they elicited similar levels of affective responding. Not all participants received the same information. As Figure 6 shows, the majority of participants ("Majority") received information indicating that the group was largely positive. A minority of participants ("Minority") received information that portrayed the group in a more ambiguous light. Participants read this information individually and then indicated their impressions of the groups on a range of dependent measures.

Figure 6: Experimental design of the small group study.

For the second phase of the experiment, participants were organized into groups of three and assigned to either a Critical or Conformist norm condition: Ten groups (30 participants) participated in the Critical norm condition, and 10 groups (30 participants) participated in the Conformist norm condition. Consistent with previous studies, participants in the Critical Norm condition were encouraged to carefully process and critically evaluate the information provided by other participants; in the Conformist Norm condition, participants were encouraged to form a consensus (Postmes et al., 2001; Teh et al., 2012). Importantly, each group of three contained 2 majority and one minority participant. One of the key aims of the experiment was to discover how participants would combine their initial attitudes towards Group A, and particularly the impact the "minority view" would have.

An experimenter led the group session. Each trial involved the experimenter reading out either positive or negative traits (e.g. "Good", "Lazy", "Intelligent", …) and then asking each participant to indicate on a 11-point rating scale the degree to which they considered each statement to be representative of members of group A. Importantly, they were asked to voice their judgment by reading aloud the number they assigned (i.e. "Ten" to indicate it was very representative). The order in which participants answered was randomized across trials. On finishing the study, participants were debriefed and thanked for their participation.

5.2 Model hypotheses

The basic idea behind this methodology was to encourage individuals to learn from the classification behavior of others. As in Sherifs' studies, we created an experimental situation in which participants voice their judgments about the characteristics of an ambiguous phenomenon, in the hope that they would converge and that agreement would emerge from interpersonal interaction.

Using the SCM, we can simulate the experiment and generate predictions, both about the initial attitudes with which participants start, and about the shared attitudes that develop over the course of the group phase. Figure 7 shows how agent networks adapt as they accumulate information on the co-occurrences of the target groups and their attributes, by changing the internal weights of the connections between the target group and its attributes. The model predicts that when Latitude is high (left panel), the three agents converge to a more neutral attitude towards Group A. With lower latitude, there is less convergence between majority and minority agents, and the shared attitude towards Group A is slightly more positive as compared to higher Latitude condition. Essentially, none of the agents are willing to adjust their attitude quite as much, and the majority position remains more positive. This results in a more positive shared attitude.

INSERT Figure 7: Simulated attitude towards Group A for 2 majority and one minority agent in the High Latitude (left panel, L=.5) and Low Latitude simulation (Right panel, L= 0.1). Each step on the X-axis corresponds to one interaction.

So essentially, the model predicts that in the Critical condition (high latitude), participants will be more likely to adjust their attitudes and converge on a more neutral attitude, which incorporates both majority and minority positions. In the Conformist condition (low latitude), participants will adjust their attitudes less, the majority will remain more positive, and the shared attitudes will reflect this.

5.3 Results

As mentioned, only a few studies have investigated the actual development of shared attitudes within small groups. This is partly due to the failure of existing theories and models to provide clear predictions of how individual cognitions and group dynamics will interact and shape each other. As Figure 7 shows, the SCM overcomes this limitation and makes a number of particular predictions regarding the development of attitudes towards the target group. First, the model predicts that, regardless of the norm (or Latitude), participants will converge to a more neutral attitude as the experiment progresses. Second, the model predicts that this convergence will proceed faster in the Critical condition, compared to the Conformist condition. A first look at Figure 8 suggests that these predictions have both been confirmed.

To monitor their development over time, attitudes towards the target were averaged across the 3 participants for each group, and this after each block of 2 positive and 2 negative adjectives. This corresponds to a step in Figure 8. To test the first hypothesis, a difference score was calculated between Majority and Minority attitudes at step 1 and step 10. Across the two Norm conditions, this difference was significantly smaller at Step 10 (M= 0.455) compared to step 1 (M= 1.35), $t(1, 9)= 6.71$, $p< 0.05$), clearly confirming the first hypothesis. The model also predicted that this convergence effect would be more outspoken in the Critical (corresponding to High Latitude) compared to the Conformist (Low Latitude) condition, which was also supported by the data: The difference between Majority and Minority attitude at step 10 was larger in the Conformist (M=.85) compared to the Critical condition (M=.06), $t(1, 9)= 6.71$, $p< 0.05$.

Figure 8: Average attitudes towards Group A for majority and minority participants in the Conformist (Low latitude) and Critical (High Latitude) conditions. Each step on the X-axis corresponds to a block of 4 trials (i.e. 2 positive and negative traits).

Both of these results confirm model predictions that in the Critical condition (high Latitude), participants would be more likely to adjust their attitudes and converge on a more neutral attitude, incorporating both majority and minority positions. Figure 8 shows that in the Critical condition (Left panel), the majority (slope gradient, or b_1= -.052) and minority (b_1= .058) moved closer to each other, developing a more neutral attitude towards Group A. In the Conformist condition (Right panel), both majority (b_1= -.022) and minority (b_1= .031) participants moved closer at a *slower* speed, and were more inclined to maintain their initial attitude position as compared to the critical condition. This is confirmed by a statistical comparison: Although the Majority Attitude became more neutral in both conditions, it did so significantly faster in the Critical condition (b_1= -.052) compared to the Conformist condition (b_1= -.022), $t=2.97$, $p < .01$. Similarly, the Minority Attitude became more neutral significantly faster in the Critical condition (b_1= .058) compared to the Conformist condition (b_1= .031), $t=-2.613$, $p < .02$. In all, the predictions generated by the model regarding how individual cognition (i.e. the development of attitudes) and social influence (the impact of the norm manipulation) would shape each other were very much confirmed.

5.4 Discussion

In our experiment, participants were provided with information about individuals belonging to a fictional social group (labeled "Group A") and were asked to form an impression of this group. They were told that they would be able to share their impressions with other participants afterwards in groups of 3. Half of these groups were encouraged to carefully process and critically evaluate the information provided by other pants ("Critical norm"), the other half was encouraged to form a consensus ("Conformist Norm") (see Postmes et al., 2001; Teh et al., 2012). At the start of the group phase, each group had one participant with a "minority" attitude and two participants with a "majority" attitude towards Group A. One of the key aims of the experiment was to test SCM predictions on how participants would combine these initial attitudes, and particularly the impact the "minority view" would have on the emerging shared attitude or stereotype.

Overall, there was strong support for SCM predictions: Participants converged to a more neutral attitude as the experiment progresses, and this regardless of the norm condition. In addition, this convergence proceeded faster in the Critical compared to the Conformist condition. And as predicted by the SCM, this was due to the fact that participants in the Critical condition were more likely to incorporate both majority and minority positions, giving rise to a more neutral shared attitude towards Group A .

These results provide further support for the notion that social comparison processes play an important role in the formation and maintenance of stereotypes. Previous studies had shown that individuals use the degree of consensus within a group as a measure of subjective validity of their opinions (Festinger, 1954; Sherif, 1935; Turner, 1987) and stereotypical representations of a target group (Haslam, Oakes, Reynolds, & Turner, 1999). Our study extends this research by providing a more detailed insight into the underlying socio-cognitive process, and how it interacts with norms. As such, the combination of simulations and empirical data provides a detailed process model of the core social comparison processes through which stereotypes and attitudes appear to be socially validated within groups of individuals.

The direct application of a connectionist ABM to small group experiments is unusual within ABM literature. ABM are typically applied to large, complex systems (i.e. the emergence of cooperation and norms within cultures and economies) and agent behavior within such simulations is mostly governed by simple if-then rules. However, a number of authors have pointed out that these simple if-then rules often result in agent behavior that is psychologically implausible, and as a result they do not necessarily produce meaningful insight into real human interaction (Sallach, 2003; Sun, Coward & Zenzen, 2005). The recurrent connectionist network provides a psychologically plausible, but still relatively simple agent model. In addition, any agent in the SCM can in principle interact with any other agent, and the impact of the interaction will adapt to experience. This is different from previous ABM such as cellular automata, where agents interact in their geometric neighborhood; or social network models, where the strength of the social relations does not change as a function of agent experience. Our model incorporates an adaptive agent within an adaptive network structure, which allows studying the co-evolution of individual cognition and social relations. From a social-psychological perspective, this is an important step in developing better models of individual and social cognition.

However, we also argue that such improvements to simulation models alone do not guarantee meaningful insights into individual and social psychological processes. To achieve this, we suggest a tight coupling between simulations and experimental studies, in which iterative loops of simulations and experiments inform psychologically plausible ranges of novel model parameters. This can provide an impetus to experimental research investigating the development of shared attitudes within real, interacting groups of participants. As mentioned, there is a lack of theoretically driven research in this area, due to the complexity that arises from the interactions between even small groups of individuals, and the fluid nature of attitudes. The use of a model such as the SCM can alleviate these problems, and allows making precise predictions of how individual and interpersonal processes will interact to shape the development of shared attitudes.

The simulations and experiment demonstrate the impact that the level of Latitude can have on group dynamics, and the attitudes that emerge from it. Even small variations in Latitude settings can apparently have significant consequences on the development of attitudes. Variations in other parameters, such as agent learning rate or social learning rate, do not have similar impacts on group dynamics. The results from this experiment, and the accompanying simulations, suggest that reduced Latitude results in less adoption of a majority position, and to more divergent opinions. As such, it appears to mimic the impact of reduced communication between cliques (Figure 5). The SCM thus suggests at least two factors that influence the critical group dynamics underlying these results: (1) the amount of communication across cliques, or opinion groups and (b) the amount of agent latitude. This issue, and the way that communication and latitude interact, can be further explored both by more fine-grained experimental set-ups, but also through further simulations, using larger simulated collectives. This will allow us to explore under which precise lab-conditions different cliques show assimilation or contrast, and also how these conditions hold up in larger collectives, using more cliques and more differentiated beliefs. This will give us different views on the important issue of why, and how, different cultures, or subcultures, sometimes assimilate, and other times seek further contrast, leading to more polarization and potentially undesirable societal tension and conflict.

6. Conclusion

The simulations show that the SCM can potentially model the impact of a variety of psychological determinants of Latitude. This has a number of advantages: Psychological theories are often built around verbal axioms that do not lend themselves well to algorithmic implementation. The SCM provides an account in terms of simple, but powerful algorithms that mimic real memory processes, and allows exploring macro level-consequences of the repeated application of these processes, in parallel by many agents, within an artificial social system. Such an integrated framework will allow investigating the impact of psychological latitude determinants, ranging from the extremity of the standard (Herr, Sherman, & Fazio, 1983; Herr, 1986), the ambiguity of the target (Herr et al., 1983), or category membership (Mussweiler, 2003), in an unprecedented way.

This approach can also help develop the existing empirical literature in novel directions. For instance, the empirical approach in this paper was very much inspired by social psychological research into how interpersonal communication *verifies* stereotypes (Festinger, 1954; Sherif, 1935; Turner, 1987). However, it is obvious that there are conditions under which stereotypes will not be accepted by members of a social group. Surprisingly, this topic has received relatively little attention in social psychological research, but there are a number of simulation models that have explored it. For instance, Huet, Deffuant & Jager (2008) demonstrated the impact of multi-dimensional attitudes on the development of conformity, using a model very similar to the SCM: Both models are built around a rejection mechanism determined by a tolerance parameter. The experimental design introduced in this paper could be adapted, for instance by providing information to participants that is inherently contradictory (i.e., Group A is both aggressive and peaceful). This would allow testing the predictions of the SCM, and to formally compare it with other simulations.

The SCM can promote theoretical integration, by examining the impact of a variety of psychological determinants of Latitude, which have often been studied in different fields of research. In addition, the model can also generate novel and testable predictions, and as such support a process of probing and prediction. In this process, simulations provide guidance for empirical research as well as sufficient depth to support interactive modification of the underlying theory. Using a computational model, we can create large sets of simulated social groups, each set having its own characteristics, which can then be put into a wide range of conditions. This represents a very efficient way of testing theoretical predictions in simulated collectives.

By creating a community of networks, new parameters appear that are not present when we only consider an individual network. It is the nature of these community parameters, and the nature of their relationship with features of real social groups, that are the focus of our empirical program. Small group studies, like the one described above, allow us to validate and inform the most psychologically plausible values of these parameters. Although space limitations prevented us from elaborating on this in detail, we also explored the robustness of the reported simulations by using different parameter ranges and variations on the architecture, including an exhaustive search of the novel community parameters. This did not alter the basic simulated patterns in a significant way. Similarly, using a distributed representation rather than a localist representation (i.e. each concept is represented by a pattern of activation across a set of micro-units) does not alter the fundamental simulated patterns.

The SCM shares similarities with other simulation work that integrates assimilative and contrastive social influence (Hegselmann & Krause, 2002; Van Overwalle & Heylighen, 2006). For instance, so-called "seceder" models postulate that people seek distinctiveness when in the company of others. In these models, individuals are assigned particular attitude positions on a continuous scale, and then look at randomly chosen others to select the one whose attitude is the most different from the mean attitude within a particular comparison group. Although there are distinct architectural differences, many of these models are very consistent with each other, as they all predict that individuals will only move to the position of a deviant, nonconforming position when in the company of at least a few similar others (i.e. as part of a clique). As such they combine an assimilative process, i.e. moving toward a chosen individual, with a contrastive process, moving toward the extreme and away from the mean. Although formal simulations are needed, it is likely that the predictions of these models of how groups can end up converging to a common opinion or split into several subgroups holding differing opinions, depending on the initial attitude distribution and the threshold for influence from others, are very similar.

In this paper, we tried to illustrate how a computational model, build around a valid, social psychological theory of agent heuristics can contribute to a better understanding of social complex phenomena. We introduced an agent-based implementation of Festingers' Social Comparison Theory (SCT), consisting of connectionist networks that simulate agent-level social comparison processes. By creating an extended network, or a "community of networks", these agents can exchange information with each other. In comparison to other ABM, the current model allows investigating individual and social adaptive behavior in more detail. The model generates predictions regarding how individual cognition (i.e. the development of attitudes) and social influence shape each other that can be tested in empirical studies. At the same time, simulation 2 (Figures 4 and 5) illustrates that the model can be "scaled up" to simulate group dynamics involving larger number of agents. Although beyond the scope of the current article, it shows that the current model can also be applied to larger, societal phenomena.

By directly applying simulations to a classic small group experiment, we aimed to illustrate how the SCM generates behavior consistent with a number of core group dynamical processes, and how it can also be used to generate precise hypotheses that can be empirically tested. Through an iterative process of prediction, testing and model building, the architecture and its parameters can be further empirically validated. Such an integrated framework will allow investigating some of the key theoretical predictions around the origin and maintenance of socially shared information through social comparisons, in ways that previous generations of social scientists, including Festinger and Sherif, could not.

References

Axelrod, R. (1997). The Dissemination of Culture: A Model with Local Convergence and Global Polarization. *Journal of Conflict Resolution*, *41*(2), 203–226.

Festinger, L. (1954). A Theory of Social Comparison Processes. *Human Relations*, *7*(2), 117–140.

Fridman, N., & Kaminka, G. A. (2011). Towards a computational model of social comparison: Some implications for the cognitive architecture. *Cognitive Systems Research*, *12*(2), 186–197.

Gilbert, P., Price, J., & Allan, S. (1995). Social comparison, social attractiveness and evolution: How might they be related? *New Ideas in Psychology*, *13*(2), 149–165.

Goethals, G., & Darley, J. (1977). Social comparison theory: An attributional approach. In J. Suls & R. Miller (Eds.), *Social comparison processes: Theoretical and empirical perspectives* (pp. 259–278). Hemisphere.

Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, *12*(4), 392–407.

Haslam, S. A., Oakes, P. J., Reynolds, K. J., & Turner, J. C. (1999). Social identity salience and the emergence of stereotype consensus. *Personality and Social Psychology Bulletin*, *25*(7), 809–818.

Hazlehurst, B., & Hutchins, E. (1998). The Emergence of Propositions from the Coordination of Talk and Action in a Shared World. *Language and Cognitive Processes*, *13*(2-3), 373–424.

Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis and simulation. *JASSS*, *5*(3).

Herr, P. M. (1986). Consequences of priming: Judgment and behavior. *Journal of Personality and Social Psychology*, *51*(6), 1106–1115.

Herr, P. M., Sherman, S. J., & Fazio, R. H. (1983). On the consequences of priming: Assimilation and contrast effects. *Journal of Experimental Social Psychology*, *19*(4), 323–340.

Huet, S., Deffuant, G., & Jager, W. (2008). A rejection mechanism in 2D bounded confidence provides more conformity. Advances in Complex Systems, 11(4), 529–549

Isenberg, D. J. (1986). Group polarization: A critical review and meta-analysis. *Journal of Personality and Social Psychology*, *50*(6), 1141–1151.

Kelly, J. R., & Karau, S. J. (1999). Group Decision Making: The Effects of Initial Preferences and Time Pressure. *Personality and Social Psychology Bulletin*, *25*(11), 1342–1354.

McClelland, J. L, & Rumelhart, D. E. (1988). *Explorations in parallel distributed processing: A handbook of models, programs and exercises.* Cambridge, MA: Bradford.

McClelland, James L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, *114*(2), 159–188.

Mussweiler, T. (2003). Comparison processes in social judgment: Mechanisms and consequences. *Psychological Review*, *110*(3), 472–489.

Postmes, T., Spears, R., & Cihangir, S. (2001). Quality of decision making and group norms. *Journal of Personality and Social Psychology*, *80*(6), 918–930.

Resnick, M. (1994). *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*. MIT Press.

Sallach, D. L. (2003). Social Theory and Agent Architectures: Prospective Issues in Rapid-Discovery Social Science. *Social Science Computer Review*, *21*(2), 179–195.

Scholten, L., van Knippenberg, D., Nijstad, B. A., & De Dreu, C. K. W. (2007). Motivated information processing and group decision-making: Effects of process accountability on information processing and decision quality. *Journal of Experimental Social Psychology*, *43*(4), 539–552.

Sherif, M. (1935). A study of some social factors in perception. *Archives of Psychology (Columbia University)*, *187*, 60.

Smith, E. R., & Conrey, F. R. (2007). Agent-based modeling: A new approach for theory building in social psychology. *Pers. Soc. Psychol. Rev.*, *11*(1), 87–104.

Smith, Eliot R., & DeCoster, J. (1998). Knowledge acquisition, accessibility, and use in person perception and stereotyping: Simulation with a recurrent connectionist network. *Journal of Personality and Social Psychology*, *74*(1), 21–35.

Sun, R., Coward, L. A., & Zenzen, M. J. (2005). On levels of cognitive modeling. Philosophical Psychology, 18(5), 613–637.

Teh, A., Baniassad, E., Van Rooy, D., & Boughton, C. (2012). Social Psychology and Software Teams: Establishing Task-Effective Group Norms. *IEEE Software*, *29*(4), 53–58.

Turner, J. C. (1987). *Rediscovering the Social Group: A Self-Categorization Theory*. Basil Blackwell.

Van Overwalle, F., & Heylighen, F. (2006). Talking nets: A multiagent connectionist approach to communication and trust between individuals. In *7th International Conference on Cognitive Modeling* (pp. 606–627). Trieste, ITALY: Amer Psychological Assoc/Educational Publishing Foundation.

Van Rooy, D., Van Overwalle, F., Vanhoomissen, T., Labiouse, C., & French, R. (2003). A recurrent connectionist model of group biases. *Psychological Review*, *110*(3), 536–563.

Van Rooy, Dirk, Vanhoomissen, T., & Van Overwalle, F. (n.d.). Illusory correlation, group size and memory. *Journal of Experimental Social Psychology*.

Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology*, *81*(5), 815–827.

Figure 1: Recurrent agent network

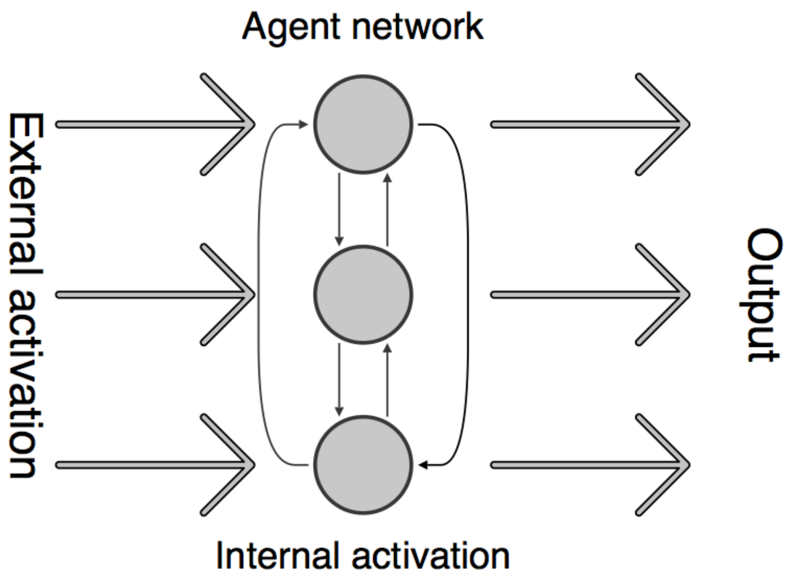

Figure 2: Interaction between Comparer and Target agent.

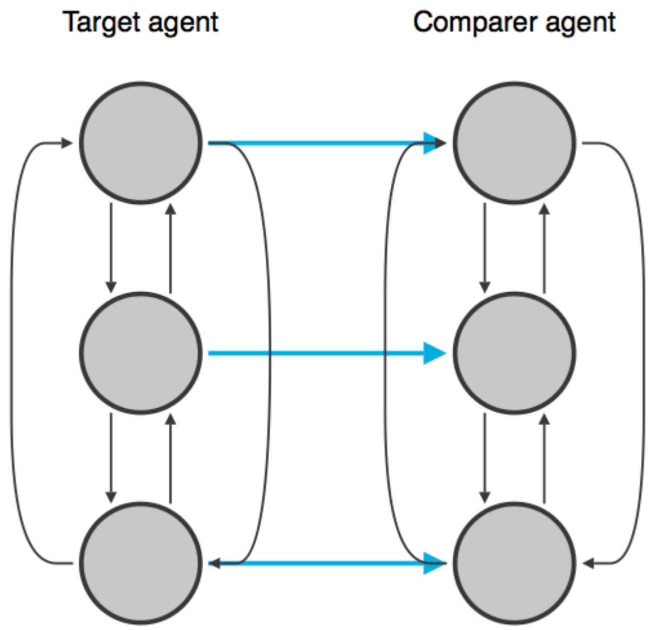

	Agent Network:		
Frequency:	A	+	-
#9 {A+}	1	1	0
#4{A-}	1	0	1

Table 1: Agent network and learning history of 1 agent in a simulation of a group perception experiment. A value of "+1" indicates the unit is activated, while "0" means inactive. Network: A= Group label, += positive attribute, -= Negative attribute.

Figure 3: Learning histories for 5 agents (top left), resulting in different simulated attitudes towards "Group A". Figure shows the development of the simulated attitudes in function of amount of learning.

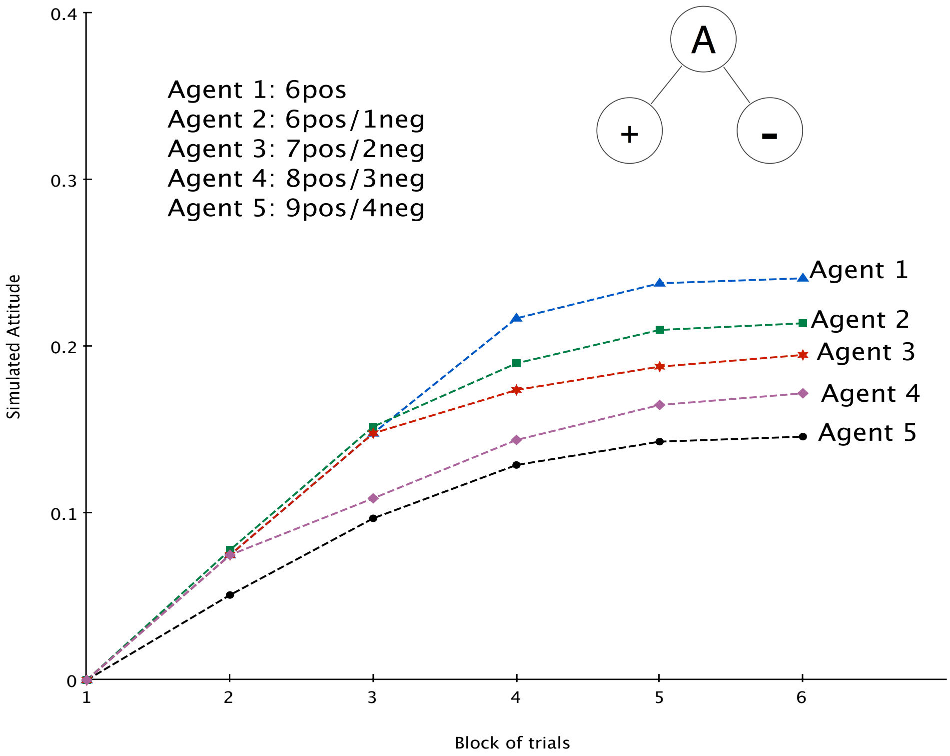

Figure 4: Development of attitudes towards Group A of 10 agents in function of interaction steps. Latitude parameter set to .01. (left panel a) and 0.4 (right panel). At each step, each agent interacted twice with all other agents.

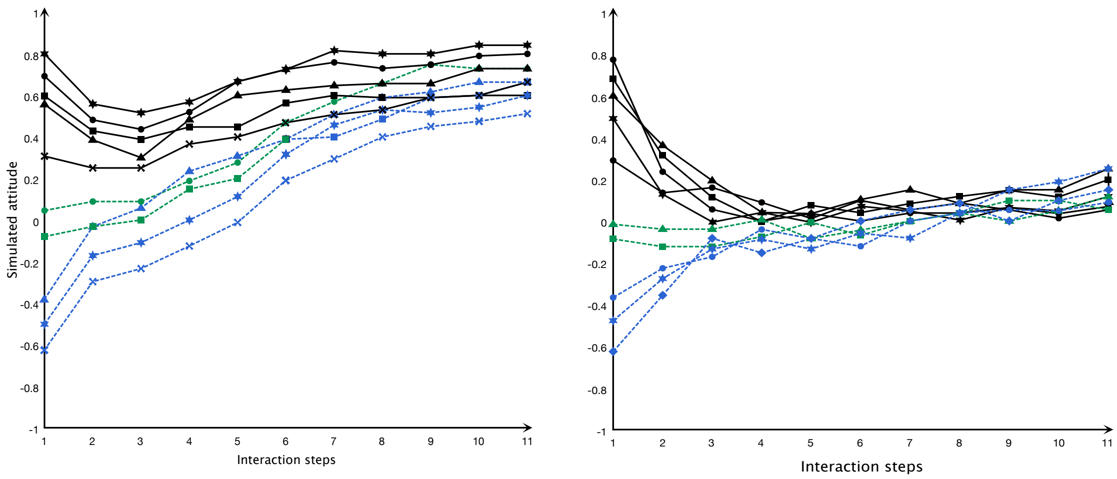

Figure 5: Development of attitudes towards Group A of 10 agents in function of amount of interaction. Within each clique, agents talk to each other three times, before talking to all other agents once. Latitude is low (L= 0.01).

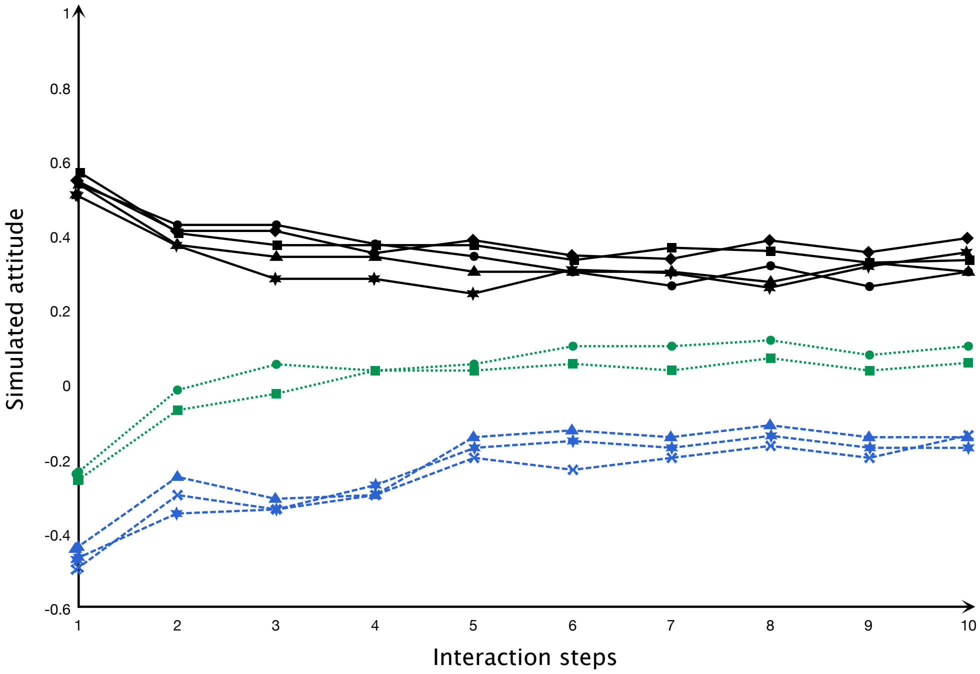

Figure 6: Experimental design of the small group study.

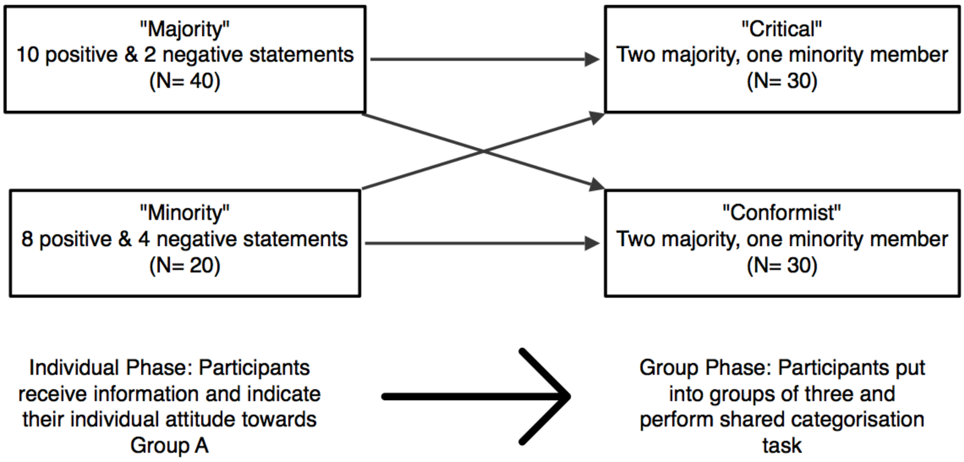

Figure 7: Simulated attitude towards Group A for 2 majority and one minority agent in the High Latitude (left panel, L=.5) and Low Latitude simulation (Right panel, L= 0.1). Each step on the X-axis corresponds to one interaction.

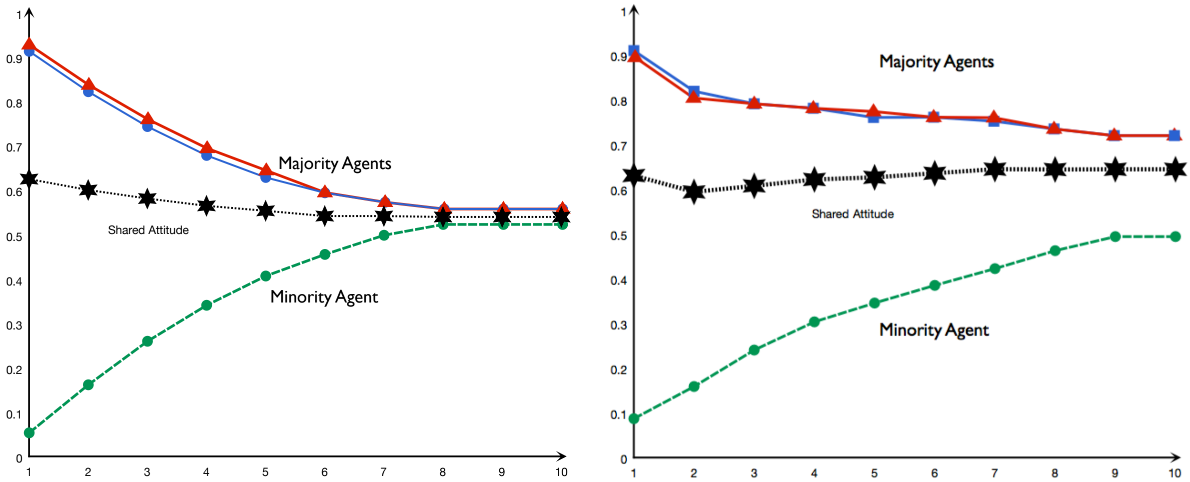

Figure 8: Average attitudes towards Group A for majority and minority agents in the Conformist (Low latitude) and Critical (High Latitude) conditions. Each step on the X-axis corresponds to a block of 4 trials (i.e. 2 positive and negative traits).