

# Norm Diversity and Emergence in Tag-Based Cooperation

Nathan Griffiths<sup>1</sup> and Michael Luck<sup>2</sup>

<sup>1</sup> Department of Computer Science, University of Warwick,  
Coventry, CV4 7AL, UK,  
`nathan@dcs.warwick.ac.uk`

<sup>2</sup> Department of Informatics, King's College London,  
London, WC2R 2LS, UK,  
`michael.luck@kcl.ac.uk`

**Abstract.** In multi-agent systems norms are an important influence that can engender cooperation by constraining actions and binding groups together. A key question is how to establish suitable norms in a decentralised population of self-interested agents, especially where individual agents might not adhere to the rules of the system. It is desirable, in certain situations, to establish multiple co-existing norms within a population to ensure a diversity of norms, for example to give agents alternatives should one norm collapse. In this paper we investigate the problem of norm emergence, and the related issue of group recognition, using tag-based cooperation as the interaction model. We explore characteristics that affect the longevity and adoption of norms in tag-based cooperation, and provide an empirical evaluation of existing techniques for supporting cooperation in the presence of cheaters.

## 1 Introduction

Multi-agent systems often comprise multiple self-interested agents seeking to achieve tasks that they cannot, or not as easily, achieve alone. In a sense, however, this self-interest suggests that without some other constraining influence, cooperation is unlikely to emerge. Norms provide a source of such influence on agent behaviour, by constraining actions and binding a group together so that cooperation naturally arises. In this view, one key question is how to establish a suitable set of norms. While formally established institutional rules offer a means of doing this in a centralised fashion, such centralised control is often not possible in large dynamic environments. Indeed, as has been recognised elsewhere [5, 32], *social norms* are not formal, prescriptive, centrally imposed rules, but often emerge informally through decentralised agent interactions. In this paper, we explore the nature of such social norms and their impact on group formation through empirical analysis, and examine the impact of *cheating* agents: those that fail to comply with norms but seek to enjoy the benefits of the group.

In seeking to investigate these issues, we adopt the *tag-based* approach taken to the problem of *group recognition*, by Riolo, Cohen and Axelrod, who use

observable tags as markings, traits or social cues attached to individuals [23]. Using this approach, Hales and Edmonds have achieved promising results in peer-to-peer settings [14], but these are not resilient when cheaters are introduced, and assume that agents have full control over their links to others, and are able to completely change these links in a single operation. In particular, we need to support cooperation in dynamic environments that may contain cheaters, in which individuals have limited control over their connections.

We view tags as capturing *social norms*: tags are recognised by agents who form groups that share a tag (within their tolerance values). The tag can be seen as an abstraction of a norm that is adopted by the agents who share it, with the group itself being *governed* by that tag, or norm, which binds the group together. This interpretation of norms is rather different from that considered in previous work on norm emergence, which has focused directly on the behaviour of individuals, such that once a behaviour is adopted by a group of agents, the behaviour is said to correspond to a norm. By contrast, in tag-based systems cooperative behaviour is determined by a combination of observable traits, in the form of tags, and tolerance. In changing its tag, an individual alters the nature of its own cooperative actions and the actions of others towards it. The notion of a norm becomes implicit, and results from agents choosing to exhibit similar observable traits, which in turn gives rise to particular behaviour. Individuals may alter their observable appearance in order to influence the manner in which others interact with them, and a particular appearance may come to be correlated with a particular behaviour. In such a setting, adopting a given appearance becomes synonymous with adopting a norm.

Many previous investigations of norm emergence have focused on very small sets of possible actions or states, with the view that the ideal situation is one in which all agents in the population select a common action or state, such as driving on the left. In scenarios such as the Coordination Game, Social Dilemma Game and Prisoner's Dilemma, agents are able to be in one of only two states, and it is with respect to these states that norm emergence is considered [9, 10, 19, 21, 27]. Although other scenarios have been explored that contain a wider range of states, such as the language/vocabulary coordination problem [30], the focus has been on obtaining a single common state in the population [25].

In open societies there are typically many possible actions and states from which many possible norms can emerge, and it may not be desirable, or even possible, to converge on a single norm. Many scenarios exist in which it is desirable to establish a selection of norms. For example, in resource allocation we would like a range of resources to be used rather overusing a single resource, and for role adoption it is desirable for agents to adopt different roles to ensure that all the required roles are taken. Such scenarios are exemplified by the *El Farol Bar Problem* [2] in which patrons must decide whether to go to the El Farol Bar, where the utility of attending is high when the bar is not too overcrowded. A high utility is achieved only when people adopt different norms, and attend the bar on different occasions. Multiple norms are also important in situations where an established norm can collapse, since it is desirable for agents to be able

to adopt an alternative established norm, rather than having to establish a new norm from scratch. For example, in the case of a monopoly failure the impact is lessened if alternative providers are already established. Tags allow us to investigate this multiplicity of norms, since the range of possible tag values is large and corresponds to the range of possible norms; although we desire sufficient commonality of tags to engender cooperation, we do not require a single tag to be adopted by the entire population. Thus, we are able to investigate norm emergence where many norms are present in the population.

In this paper, we examine the problem of supporting cooperation from the perspective of norm emergence, and evaluate the effect of alternative techniques on norm emergence. Specifically, we investigate our recently proposed techniques for coping with cheaters in tag-based systems [12, 13] in terms of the norms that emerge over time. This is a novel interpretation of tag-based cooperation, and allows us to investigate a scenario in which multiple norms co-exist. The key contributions are an evaluation of the characteristics affecting the longevity and size of norm-governed groups in tag-based cooperation, and an increased understanding of the operation and implications of our previously proposed mechanisms for coping with cheaters.

The paper begins with an introduction to tag-based cooperation, followed by the specifics of using *context assessment* and *rewiring* to improve group effectiveness in the presence of cheaters. Then, in Section 4, we present an analysis of our experimental findings, and finally we conclude with a discussion of our results and their more general significance.

## 2 Background

It has been widely argued that *norms* provide a valuable mechanism for regulating behaviour in decentralised societies [3, 11, 32]. Through the ongoing behaviour of individuals, norms can emerge that provide coherence and stability, and support cooperation. A common view is that where a group of agents share a particular strategy, behaviour or characteristic, a norm is established [26]. In this paper we investigate factors influencing norm emergence in a population of agents, each of which has a set of neighbours with which it interacts. This abstract environment reflects the form of many real-world settings, such as ad-hoc communication networks or P2P content sharing. We assume that agents know the identity of their neighbours, rather than assuming anonymity of interactions. Although many P2P researchers assume anonymity, this is often introduced as a way of modelling the rarity of repeat interactions between any two individuals. Our formulation is compatible with this view, since we assume that there are very few repeat interactions, yet at some level the identity or location of others must be known, otherwise it is impossible to interact with them. For example, a network address is needed to share content and an identifier is required to forward packets in a communication network. We also assume that agents have limited observations of others and that there is no direct reciprocity, and so we adopt Riolo, Cohen and Axelrod’s tag-based approach, introduced below.

Tag-based cooperation has been considered for many years by biologists and social scientists investigating how cooperative societies of selfish individuals might evolve through the recognition of cultural artefacts or traits [1, 6, 8, 15, 16]. Simple observable traits, or *tags* [17], can be used as cultural artefacts to engender cooperation without relying on reciprocity [4, 23, 31]. Existing work on tags, however, has given little consideration to the possibility that some members of the population may be *cheaters* who deviate from the rules of the system, by not cooperating when they should. In this paper, our investigation of norm emergence allows for the possibility of cheaters.

Riolo, Cohen and Axelrod (RCA) propose a tag-based approach to cooperation in which an individual's decision to cooperate is based on whether an arbitrary tag (i.e. observable trait) associated with it, is sufficiently similar to that associated with a potential recipient [23]. The approach is illustrated using a simple *donation scenario* in which each agent acts as a potential donor with a number of randomly selected neighbours. Should an agent opt to donate, it incurs a cost  $c$ , and the recipient gains a benefit  $b$  (it is assumed that  $b > c$ ), otherwise both receive nothing. Each agent  $i$  is initially randomly assigned a tag  $\tau_i$  and a tolerance threshold  $T_i$  with a uniform distribution from  $[0, 1]$ . An agent  $A$  will donate to a potential recipient  $B$  if  $B$ 's tag is within  $A$ 's tolerance threshold  $T_A$ , namely  $|\tau_A - \tau_B| \leq T_A$ . Agents are selected to act as potential donors in  $P$  interaction pairings, after which the population is reproduced proportionally to their relative scores, such that more successful agents produce more offspring. Each offspring is subject to mutation, so that with a small probability a new (random) tag is received or noise added to the tolerance. In relation to norms, the key aspect here is that donation rate is an assessment of the effectiveness of the society and the impact of norms: the greater the effectiveness, the higher the donation rate.

RCA have shown that a high cooperation rate can be achieved with this simple approach. They observe cycles in which a cooperative population is established, which is then invaded by a mutant whose tag is similar (and so receives donations) but has a low tolerance (and so does not donate). Such mutants initially do well, leading to them taking over the population and subsequently lowering the overall rate of cooperation, but eventually the mutant tag and tolerance become the most common and cooperation again becomes the norm [23]. It is important to note that cooperation does not require a single tag to be adopted throughout the population and is achieved where there are subgroups of agents in the population that share different tags, corresponding to subgroups having adopted different norms. As these subgroups evolve they mirror the typical life-cycle of norms: initially a small number of agents adopt the norm; the group then expands as the norm is more widely adopted; over time the norm becomes outdated and alternatives emerge conflicting with the established norm; eventually one of these new norms may replace the previously established norm. Tag-based cooperation, therefore, provides us with a scenario in which to consider the emergence and ongoing evolution of multiple norms within a population.

Hales and Edmonds (HE) apply RCA’s approach in a P2P setting, with two main changes [14]. First, a *learning interpretation* of reproduction is adopted, so that each agent compares itself to another in the population at random and adopts the other’s tag and tolerance if the other’s score is higher (subject to potential mutations) [23]. Second, HE interpret a tag as being an agent’s neighbours in the P2P network, i.e. an agent’s links to others. In RCA’s work each agent is connected to each other agent, with no corresponding notion of neighbourhood. In HE’s model, the process of an agent adopting another’s tag is equivalent to dropping all of its own connections, and copying the connections of the other agent (and adding a connection to the other agent itself) [14]. Importantly, in our view, this model reflects the formation of groups based on recognition of tags in group members. Since tags, in HE’s model, are interpreted as corresponding to an agent’s set of neighbours, mutation corresponds to resetting the neighbour list (replacing it with a single neighbour randomly selected from the population). In this paper we use numerical tags, meaning that tag mutation does not affect an individual’s set of neighbours. Instead, inspired by HE, we enable agents to rewire their neighbourhood, and so change the set from which potential donors are selected.

Using simulations, HE have shown this approach to be promising in situations where agents are given free reign to rewire the network and replace all of their connections each reproduction. This rewiring is an *all-or-nothing* operation, in that although an agent can adopt a completely new set of neighbours (*replacing* its existing neighbourhood), it cannot *modify* its existing neighbourhood. Our view is that such extreme rewiring, where the neighbourhood topology might completely change with each new generation, is not practicable in all scenarios. For example, in a communication network this would imply that all existing routes become outdated and need to be re-established, while in a content sharing system an agent would lose all information about the content available in its neighbourhood. In this paper we consider a less extreme situation, in which agents are able to rewire a proportion of their neighbourhood.

Both RCA and HE assume that agents do not deviate from the rules of the system, i.e. they assume that there are no cheaters. A *cheater* is an agent that accepts donations, but will not donate to others, even if the rules of the system dictate that it should. We assume that if a cheater reproduces, then its offspring will also cheat. In this paper we assume that the traits embodied by tags are observable to others, meaning that cheaters cannot falsify their tags. To illustrate this assumption, consider a group of agents in which a tag is analogous to wearing a coloured hat. If an agent selects a red hat, then this is visible to all others, and it cannot falsely claim to be wearing a blue hat. In future work we may consider alternative formulations that allow for the potential falsification of tags, such as scenarios in which agents can present a different appearance to different neighbours. The effect of cheaters in standard tag-based cooperation is catastrophic, with the introduction of even a small proportion of cheaters into the population causing cooperation to collapse [12].

Norm emergence has been considered in several other settings. For example, norms can emerge in a Social Dilemma Game when individuals are repeatedly randomly paired [27], and in certain settings can emerge simply by individuals learning based on their own individual histories [29]. In this paper, however, our focus is on using the interpretation of a shared tag as representing a norm to further our understanding of tag-based cooperation.

### 3 Improving Group Effectiveness

In seeking to examine the impact of cheaters on group formation and norm emergence, we consider a population of agents, each of which has its own tag and a set of connections to  $n$  neighbours, such that agents can only interact with their neighbours (although for reproduction we consider the population as a whole). A *random* network topology is used with degree distribution  $P(k)$ , such that each agent is initially randomly connected to  $k = n$  others. Each agent is selected to act as a potential donor in  $P$  interaction pairings, with the potential recipient randomly selected from its set of neighbours. As discussed in the previous section we assume that agents know the identity of their neighbours. We assume that a proportion of agents are cheaters and will not cooperate with others even when their tags are within the tolerance threshold. The *donation scenario* and parameter values used by RCA are adopted, such that benefit  $b = 1$  and cost  $c = 0.1$  [23]. Each agent  $i$  is initially assigned an arbitrary tag  $\tau_i$  and tolerance  $T_i$  with uniform distribution from  $[0, 1]$ <sup>3</sup>. We investigate norm emergence in relation to RCA’s tag-based approach and two techniques that we have previously proposed for improving cooperation in the presence of cheaters: context assessment [12] and rewiring [13].

#### 3.1 Context assessment

Our first technique, originally proposed in [12], enables agents to assess their neighbourhood, or group, in terms of how cooperative they perceive their neighbours to be. The donation decision is modified so that an agent’s assessment of its neighbourhood context becomes a factor in the decision to donate. Agents are given a fixed length FIFO memory to record the last  $l$  donation behaviours observed for each neighbour. When the neighbour donates, an observation value of +1 is recorded, and when it does not −1 is recorded. This memory is fairly sparse, since the number of interactions is small compared to the number of agents, and so the overhead incurred is relatively small (2 bits per observation for  $n \times P$  observations, where  $n$  is the number of neighbours and  $P$  the number of pairings).

---

<sup>3</sup> We actually use a lower bound on *tolerance* of  $-10^{-6}$  to address Roberts and Sherratt’s concerns regarding agents with identical tags being forced to cooperate [24]. This also allows the population to contain non-cooperative agents of the form considered by Masuda and Ohtsuki [20].

In order to assess its neighbourhood context, an agent considers each of its  $n$  neighbours in turn, and determines the contribution to the context assessment  $c_i$  of neighbour  $i$ , which is simply the proportion of observed interactions in which the neighbour donated, given by:

$$c_i = \begin{cases} \frac{\sum_{j=1}^{l_i} o_i^j}{l_i}, & \text{if } l_i > 0 \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where  $o_i^j$  represents the  $j$ 'th observation of neighbour  $i$ , and  $l_i$  is the number of observations recorded of  $i$ 's donation behaviour ( $l_i < l$ ). By considering each of its  $n$  neighbours, agent  $A$ 's assessment of its current neighbourhood context  $C_A$  is given by:

$$C_A = \frac{\sum_{i=1}^n c_i}{n} \quad (2)$$

This context assessment can be used to influence the donation decision. The intuition is that agents 'expect' that by donating they are more likely to receive a future donation from some other (observing) agent, thus binding a group together. However, since the number of interactions is small compared to the number of agents, this is a *weak* notion of indirect reciprocity. An agent's donation to another is unlikely to be directly repaid or directly observed by a third party, and so there is little direct or indirect reciprocity. Instead, context assessment gives an impression of the donation behaviour in a neighbourhood, indicating the likelihood of receiving future donations. An agent's assessment of its neighbourhood context is incorporated into the model by adapting the decision to donate, such that both tolerance and neighbourhood context are considered. Thus, an agent  $A$  will donate to  $B$  if:

$$|\tau_A - \tau_B| \leq (1 - \gamma) \cdot T_A + \gamma \cdot C_A \quad (3)$$

The parameter  $\gamma$ , called the *context influence*, allows us to tune the technique. The context influence is in the range  $[0, 1]$ , with  $\gamma = 0$  making the technique identical to RCA's method, while  $\gamma = 1$  causes the donation decision to be determined solely by an agent's assessment of its neighbourhood context.

We adopt a learning interpretation of reproduction, such that after a fixed number of interaction pairings  $P$  an agent compares itself to another, randomly selected from the population. If the other agent is more successful, then its tag and tolerance are copied (subject to a small probability of mutation), otherwise the tag and tolerance are unchanged. Note that although an agent's interactions are restricted to its neighbourhood we do not similarly constrain reproduction, and we allow an agent to compare itself against another that is randomly selected from the whole population. Thus, although agents can only donate to (and receive donations from) their immediate neighbours, they can observe the performance of the population at large. If the comparison for reproduction was restricted to an agent's set of neighbours, this would have the effect of structuring the population into (overlapping) sub-groups. In this paper we are concerned

with investigating norm emergence within a population as a whole, and so we do not impose such a restriction.


### 3.2 Rewiring

Our second technique, introduced in [13], enables agents to *rewire* their network neighbourhoods, such that after reproduction an agent removes a proportion  $\lambda$ , called the *rewire proportion*, of connections, and replaces them with connections to new neighbours. This approach is motivated by HE’s results, but unlike HE we do not assume that agents can replace *all* of their connections since, as discussed above, this is likely to be impractical in real-world settings. In our mechanism, after reproduction the  $n \times \lambda$  worst neighbours are removed and the best (non-duplicate) neighbour from each of the agent’s  $n \times \lambda$  best neighbours are added. The neighbours are considered in descending rank order and, for each, the best non-duplicate neighbour is added. Additional randomly selected neighbours are added if necessary to prevent the neighbourhood shrinking due to duplication (agents have at most one connection to another given individual, and duplicate connections are meaningless). Other rewiring strategies are of course possible, with alternative criteria for selecting links to remove and add. In [13] we evaluate a small number of alternatives and show that the strategy described above gives reasonable performance.

Connections to remove are determined by ranking each neighbour  $i$  using the contribution to the context assessment  $c_i$  (defined in Equation 1), with agents having the lowest  $c_i$  values being removed. The contribution to the context assessment is also used to determine which connections to add, with an agent asking each of its  $n \times \lambda$  best neighbours to recommend their best non-duplicate neighbour. If the  $c_i$  values of two or more agents are equal then one is selected arbitrarily. The *rewire proportion* determines the extent to which the network is rewired in each generation. Such rewiring can be thought of as being the result of a simplistic reputation mechanism, since agents update their connections based on the experiences and recommendations of others. However, unlike typical reputation mechanisms, the assessment is based on relatively little information, which is not predicated on repeated encounters or on a notion of (direct or indirect) reciprocity [18, 22]. Clearly there is a chance that an agent might ask a cheating neighbour to recommend a new connection, and so be given false or poor information. However, since the approach is to ask the *best* neighbours, the requested agents will be those that have been observed to be cooperative, and we assume that they are likely to give honest information. We consider this issue further in the discussion in Section 5.

Figure 1 illustrates the alternative rewiring approaches, in which circles represent agents, thin solid lines represent existing connections, and dotted and bold solid lines show dropped and new connections respectively. Agent  $A$ ’s original neighbourhood is shown in (a). The results of applying HE’s rewiring approach is shown in (b) where  $A$  drops all of its connections and adopts those of  $B$ . Our rewiring approach is illustrated in (c). If  $A$ ’s neighbours in order of preference are  $B, C, D, E, F$ , and 2 neighbours are to be replaced, then connections to  $E$





**Fig. 1.** Rewiring showing (a) the original neighbourhood rewired using (b) HE's method and (c) our rewiring approach.


and  $F$  will be dropped. If  $B$ 's neighbours, are  $D, H, I, G, A$  and  $C$ 's neighbours are  $J, K, L, A, I$  in preference order, then  $A$  will add  $H$  from  $B$ 's neighbourhood ( $D$  is already in  $A$ 's neighbourhood and so not added) and  $J$  from  $C$ 's neighbourhood.

## 4 Experimental Analysis

Using the PeerSim P2P simulator<sup>4</sup>, we have built a simulation that allows us to explore norm emergence using RCA's standard approach, context assessment, and rewiring. The quantitative results presented here are averaged over 10 runs using a population of 100 agents, a neighbourhood size of  $n = 10$  (meaning that the population has a degree distribution of  $P(k) = 0.1$ ), with 10 pairings per agent per generation ( $P = 10$ ), and a cheater proportion of 30%. Where context assessment is incorporated a *context influence* of  $\gamma = 0.5$  is used, and similarly where rewiring is incorporated we use a *rewire proportion* of  $\lambda = 0.5$ . After reproduction there is a 0.001 probability of mutating the tolerance of each agent by adding Gaussian noise (with mean 0 and standard deviation 0.01), along with a probability of mutating each agent's tag by selecting a new random value.

We consider two configurations for mutating tags: a *low mutation* rate of 0.001 and a *high mutation* rate of 0.01. The low mutation rate represents a generally stable population in which mutation is simply a small part of the evolutionary process. Conversely, the high mutation rate represents a more dynamic and less stable environment in which there is more significant fluctuation in tags present in the population. Evolutionary approaches, such as tag-based systems, typically use low mutation rates, but in our experiments we use a high rate to simulate a dynamic environment. Since tag mutations are imposed at the beginning of a new generation (with memories of previous interactions having been reset), the mutation of an agent's tag is equivalent to the pre-mutation agent leaving the system and being replaced by a new agent with a new tag, such that the new agent has the same connections and tolerance as the agent being replaced. Thus, using a high tag mutation rate is a simple approximation of a dynamic environment. In Section 5, we discuss the importance of considering a more

<sup>4</sup> <http://peersim.sourceforge.net/>

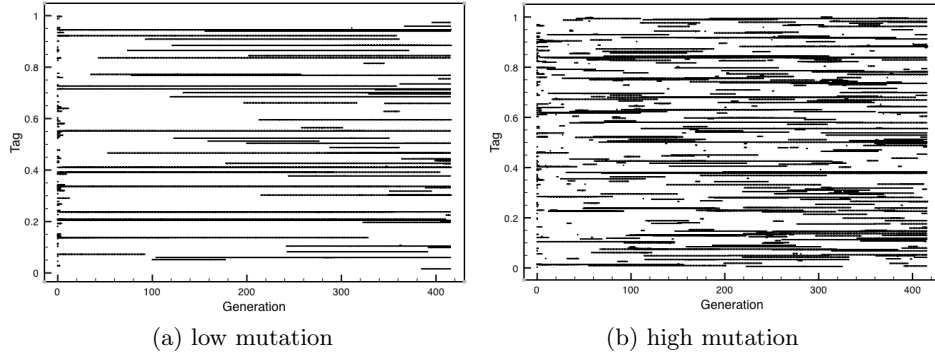


**Fig. 2.** The effect of cheaters using RCA’s standard approach with low and high mutation rates.

accurate representation of *churn* (agents leaving and joining the system), but using a high mutation rate gives an indication of the likely performance of the approach in such a setting.

In the remainder of this section we give an overview of the main findings from our simulations, focusing on two main characteristics. First, we consider the *donation rate* defined as the proportion of interactions resulting in a donation in the final generation of the simulation, averaged across the population and across simulation runs. This indicates the effectiveness of the groups that emerge in complying with the norms that establish those groups and govern their maintenance. Second, we consider the *number of unique tags* present in the final generation, again averaged across runs, which indicates the number of norm-governed groups that have been established. Where a group of agents share a tag, or each others’ tags are within their tolerance values, we interpret this as recognising a norm that is then established, since those agents will cooperate by donating to each other (provided that they are not cheaters). The number of unique tags indicates the number of such norm-governed groups that are formed, since each tag value corresponds to a tag group. However, it is important to note that this is only a coarse metric since (i) some tags may be adopted only by a single agent with no other agents having a tag and tolerance combination that gives rise to mutual cooperation, in which case there is no norm, and (ii) in some cases where a single agent has adopted a particular tag the tolerance values of this and other agents with similar tags may still give rise to mutual cooperation, and so a norm is established. The role of tolerance in the donation decision significantly increases the complexity of identifying for certain whether a norm is established. In our analysis we take a simple approach by focusing on the number of unique tags, and do not attempt to precisely quantify the number of agents that have tag and tolerance combinations that result in mutual cooperation.

Where there is a low number of unique tags, so that there are few groups, the average number of agents adopting each tag is high, and the groups are larger in size, with the respective norms being more widely adopted, and a reduced likelihood of a tag belonging to a single individual. Conversely, as the number



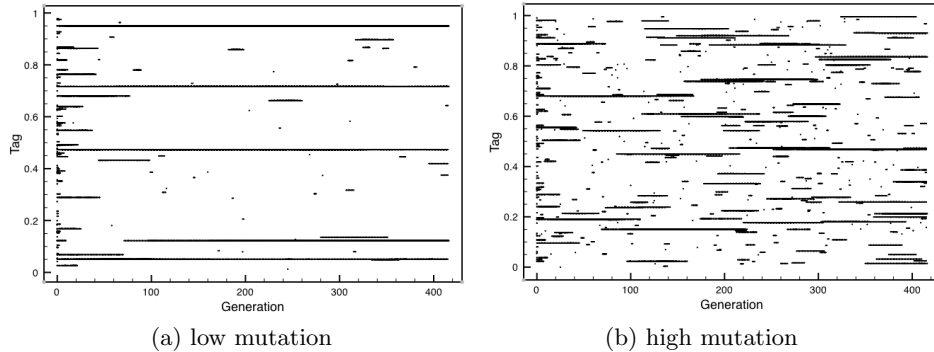
**Fig. 3.** Tags with RCA’s standard approach using low and high mutation rates.

of unique tags (and therefore groups) increases, the average number of agents having adopted each tag (and hence in each group) reduces, so that the corresponding norms are less widely adopted and there is an increased likelihood of a tag being ascribed to a single agent only. Thus, for lower numbers of unique tags there is more significance in them representing groups of agents having adopted those tags as norms. Note that this is an informal notion of norm establishment, in comparison to other approaches that have a group leader [7], or explicit strategies emerging rather than simply shared tags [27].

We use a cheater proportion of 30% in our simulations since in stable populations (with a low tag mutation), this is known to have a significant impact on the level of cooperation [12]. Figure 2 shows how the donation rate declines in a stable population (the solid line) as cheaters are introduced using RCA’s standard approach (with no context assessment or rewiring). With no cheaters we observe an average donation rate of 40.2% which reduces to 23.2% if 30% of the population are cheating agents. In a dynamic environment (the dotted line), represented by the high mutation rate, we can see that the donation rate is significantly reduced (to an average of 2.4%) regardless of the proportion of cheaters.

#### 4.1 Context assessment

The base case for our comparisons is RCA’s standard approach (without context assessment or rewiring), which with a low mutation rate gives a donation rate of 23.2%, and with a high mutation rate gives a donation rate of 1.5%. The increased dynamism of the environment, represented by the increased mutation rate, has a catastrophic effect on the donation rate, and in turn on group effectiveness. For a donation to occur, agents must share a tag (within their tolerance values). With a low mutation rate, RCA’s approach gives an average of 35.8 tags, each shared by 2.8 agents on average. With a high mutation rate, there are 79.2 tags shared by 1.3 agents on average. Thus, with a high mutation rate a large number of tags are adopted by a single agent, as confirmed by the very low

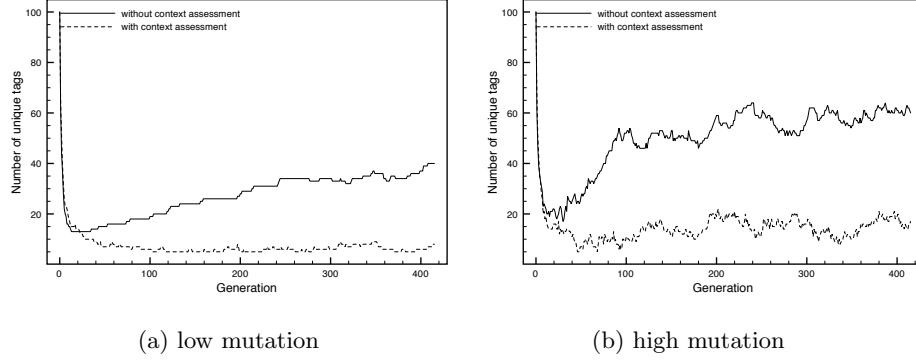


**Fig. 4.** Tags with context assessment using low and high mutation rates.

donation rate observed (3.2%). In relation to norm emergence, this means that in the low mutation case the resulting norms are on average only adopted by 2.8 agents. In the high mutation case the tag groups do not, on average, correspond to norm emergence since less than two agents adopt each tag. It is not our concern in this paper to attempt to define the number of agents needed for norm emergence, but clearly there must be at least two agents involved.

Figure 3 shows the evolution of tags in the population over the duration of a sample simulation run for both low and high mutation settings. Each point represents the presence of a tag in a particular generation, and where a tag persists for several generations the points form a line from the generation in which the tag group is created to the generation in which it collapses. Our numerical results are confirmed by Figure 3 which shows that many more tags are present in the high mutation setting than the low setting. This graphical representation also allows us to observe the formation of norm-based groups. In particular, in the low mutation setting norm-governed groups are established and maintained for many generations, while in the high mutation setting many such groups have very brief durations appearing as points or very short lines. (Note that since the number of unique tags is large in Figure 3(b), many of the points do not represent norm establishment, as discussed above.)

The evolution of tags in the population when using the context assessment technique is shown in Figure 4. Comparing Figures 3 and 4 it is immediately apparent that there are significantly fewer tags present using context assessment than with the standard approach. On average, context assessment in a low mutation setting results in only 3.7 tags (compared to 35.8 with the standard approach) and 12.1 tags (compared to 79.2) for a high mutation rate. Donation rates of 47.5% and 42.9% are obtained for the low and high mutation settings respectively (compared to 20.4% and 3.2%). The reduction in donation rate in the high mutation setting compared to the low mutation setting is less pronounced (approximately 10%) than with the standard approach (where the reduction is approximately 85%). We see this as demonstration that context assessment is more stable in supporting cooperation in dynamic environments than RCA's



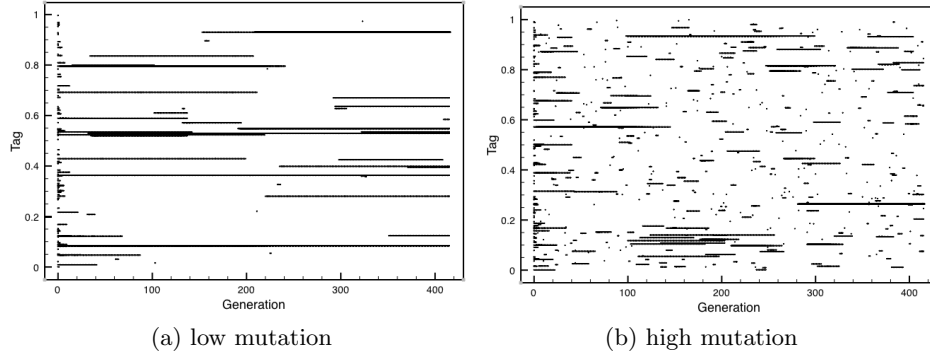
**Fig. 5.** Number of unique tags with context assessment using low and high mutation rates.

standard approach. This tells us that norms resulting from context assessment are more widely adopted than with RCA’s approach (by 27 and 8.3 agents on average for the low and high mutation rates respectively). Given that these norms are more widely adopted, we would expect an increase in the group effectiveness (as indicated by donation rate achieved), which is indeed the result we observe.

The evolution of the number of unique tags during a sample simulation run for low and high mutation rates is illustrated in Figure 5. In both settings the number of tags is initially very high and is (approximately) equal to the population size since agents are randomly allocated tags. During the first few generations the number of unique tags drops significantly as agents begin to copy tags from their more successful neighbours. As the simulation progresses the number of unique tags then stabilises. From Figure 5 we can see that in addition to context assessment resulting in significantly fewer tags than RCA’s approach, the number of tags also stabilises more quickly, meaning that the population reaches a form of convergence. In a stable population (with a low mutation rate) we achieve a quicker convergence using context assessment than without, and the resulting population contains significantly fewer tags. When the mutation rate is high the number of tags increases, as does the extent of the fluctuations over generations (with both approaches having similar fluctuation levels). Thus, dynamism in the population reduces convergence. It is clear that norm-governed groups emerge more quickly using context assessment than with the standard approach, and on average norms are adopted by many more agents. Due to space constraints we do not discuss the effect of memory length in this paper, but in [13] we have shown that it is not a significant factor.

## 4.2 Rewiring

Rewiring gives similar improvements to context assessment, with donation rates of 57.0% and 49.8% for the low and high mutation settings respectively (both of which are higher than with context assessment). The average number of unique

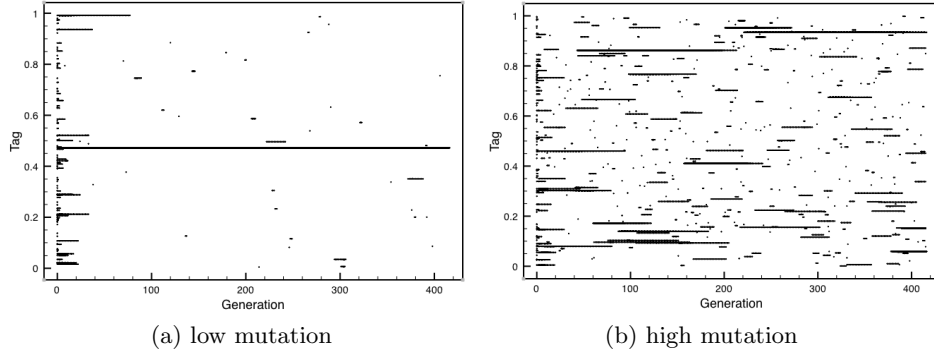


**Fig. 6.** Evolution of tags in a population using rewiring with low and high mutation rates.

tags is higher than for context assessment, with 11.9 and 16.8 for low and high mutation rates, but is significantly lower than with RCA’s approach. The norms that are established are adopted by fewer agents using the rewiring technique than with context assessment, 8.4 (rather than 27) and 6.0 (rather than 8.3) for the low and high mutation rates, although they are still adopted by many more agents than with the standard approach. This result is unexpected, since we intuitively expect that a higher donation rate (and therefore higher group effectiveness) would be achieved only when norms are more widely adopted.

The discovery that a higher donation rate can be obtained with smaller groups is potentially very powerful, since in many situations we would like to balance the desire to have widely adopted norms (and few groups) with the desire to ensure that there are several established groups from which agents can select. The widest possible adoption of a norm is where a single norm is adopted by the population in a single group. Conversely, the largest set of norms is obtained where we have a minimum number of agents per tag for norm emergence, giving groups of that size. As mentioned above, we are not concerned with attempting to define a minimum membership, but clearly more than one agent is required, and so the upper bound of the number of norm-governed groups is half the population size. In the donation scenario the motivation behind balancing the level of adoption and number of norms established is that mutations can cause a norm to collapse at any point, and in such cases we would like agents to be able to adopt an alternative norm by joining another tag group. However, this balance is a general issue in many distributed systems, and corresponds to the general view that fostering competition and avoiding monopolies can be beneficial. The question of how many norms or groups is ideal in a particular setting is an open question and requires further investigation.

The evolution of tags in the population using rewiring for sample simulation runs is shown in Figure 6. As is the case with context assessment (Figure 4) there are significantly fewer tags at any point in time than with RCA’s approach (Figure 3). As noted above, rewiring results in slightly more tags in any generation



**Fig. 7.** Evolution of tags using context assessment and rewiring with low and high mutation rates.

than context assessment. Figure 6 also allows us to observe that the duration of a given tag group is generally slightly reduced using rewiring in comparison to context assessment, implying that the norms emerging are less long-lived. In the donation scenario this is not a major concern, since there is little cost to changing tag groups, but more generally there may be a higher cost associated with such a change. It is desirable, in general, for norm-governed groups to be of longer duration as we observe with context assessment, but also for there to be a reasonable number of alternative groups as with rewiring.

#### 4.3 Combining context assessment with rewiring

We have previously shown that combining context assessment and rewiring improves the donation rate [13], and here we observe donation rates of 68.6% and 63.1% for the low and high mutation settings respectively. In terms of norm emergence, a key question is which, if any, of the properties of context assessment (few norms of longer duration) and rewiring (more norms of shorter duration) the combined approach gives. We find that the combined approach results in 4.4 and 10.4 tag groups in the low and high mutation settings, corresponding to adoption by 22.7 and 9.6 agents on average respectively. In the low mutation setting this is similar to context assessment, and we would prefer more norms given our desire to have alternative groups established in case of norm collapse, so that an agent can join an alternative group if their current group collapses. With a high mutation rate, the number of norms is similar to that obtained with rewiring, indicating that there are alternative groups in the event of a norm collapsing. However, the duration of norms and groups is also important in ensuring that alternatives are available.

Figure 7 shows the evolution of tags using a combination of the context assessment and rewiring techniques in sample runs, and we can make two important observations. First, the run for the low mutation rate represents a particularly low number of tag groups, with a single dominant group spanning all generations

**Table 1.** Summary of donation rate, number of tags and size of tag group (number of agents per tag) for the low mutation rate.

Approach	Average donation rate	Average number of tags	Average size of tag group
RCA’s mechanism	20.4%	35.8	2.8
Context assessment	47.5%	3.7	27
Rewiring	57.0%	11.9	8.4
Context assessment and rewiring	68.6%	4.4	22.7

**Table 2.** Summary of donation rate, number of tags and size of tag group (number of agents per tag) for the high mutation rate.

Approach	Average donation rate	Average number of tags	Average size of tag group
RCA’s mechanism	3.2%	79.2	1.3
Context assessment	42.9%	12.1	8.3
Rewiring	49.8%	16.8	6.0
Context assessment and rewiring	63.1%	10.4	9.6

and a small number of short duration groups appearing. A similar situation can occur when using context assessment alone (although the run in Figure 4 shows an example where this does not occur). We would like to avoid this by ensuring that there are alternative established norms (i.e. tag groups) at any point in time, in case of norm collapse. Second, with a high mutation rate many of the groups established persist for only very short durations. Thus, although there are many groups, avoiding the problem from the low mutation case, we would like them to be of longer duration. It seems, therefore, that while combining context assessment with rewiring increases the donation rate, it does not result in groups that have the desired characteristics. We would prefer a higher number of norm-governed groups than observed in the low mutation case, to ensure alternative groups are established in case of collapse, and we would prefer more persistence than observed in the high mutation case, to avoid frequent collapses. Thus, although the donation rate is high when using the combined approach, it does not achieve the more generally desired number of norms or persistence.

Context assessment appears to have a significant reduction on the number of norm-governed groups established, and so increases their size, especially with a low mutation rate. Although this results in an increased donation rate (indeed, when combined with rewiring it gives the highest donation rate) suggesting group effectiveness, it also gives reduced diversity in the population. In many settings this is undesirable, since if there is collapse of cooperation in those groups due to the collapse of the respective norms, there are no alternative established groups for agents to join.



#### 4.4 Summary of results

Our simulations show that both context assessment and rewiring improve group effectiveness (through the donation rate) by the formation of groups of agents sharing a particular tag value, which we can interpret as norm establishment. A summary of the quantitative results is shown in Tables 1 and 2. In both the low and high mutation settings the highest donation rate is achieved when combining context assessment and rewiring. Rewiring gives the second highest rate, followed by context assessment, and finally RCA’s approach. In the low mutation setting, context assessment (with and without rewiring) results in a low number of norm-governed groups being established. In the high mutation setting the number of groups established is reduced when context assessment is used in comparison to rewiring alone, but this difference is less significant. A visual analysis of the evolution of tags reveals that norm and group duration is potentially an issue, with each of the techniques discussed resulting in several short duration norm-governed groups being established.

### 5 Discussion and conclusions

Interpreting the formation of groups of agents sharing a tag as the emergence of norms allows us to view RCA’s approach as facilitating norm emergence. Indeed, it is norm emergence that leads to donations, since only when two or more agents share a tag (within their tolerance) will a donation occur. Unlike many previous investigations on norm emergence, in our model multiple norms can successfully co-exist, and we are not endeavouring to establish a single norm throughout the population. This is important for situations in which multiple established norms are desirable, such as resource allocation and role adoption, as discussed in Section 1. Our previously proposed techniques of context assessment and rewiring to cope with cheaters [12, 13] also facilitate norm establishment, and this interpretation enables us to explore their operation. The norms and groups that are formed are more widely adopted (with fewer norms being established) than with RCA’s approach. Using context assessment increases adoption, especially in the low mutation setting, but as a result there may be no alternative norm-governed groups available in the event of the collapse of an established norm.

There are four key areas of future work, which aim to broaden the applicability of viewing tags as an abstraction of norms and of the context assessment and rewiring techniques. In particular we aim to (i) further investigate the impact of mutation, (ii) improve our techniques with respect to the persistence and number of established norms, (iii) explore the effects of neighbourhood connection topologies, and (iv) evaluate the approach in more realistic environments.

An increase in mutation rate leads to fewer long lived and fewer widely adopted norms, so that there are more groups, with fewer members, persisting for fewer generations. Norms also emerge with many tag groups and few agents per tag, but only few such norms are widely adopted and long lived, as seen in Figures 4, 6 and 7, where there are few long duration groups. This is the first area of future investigation, to understand how the tag mutation rate, and

fluctuations in mutation rate over time, relate to the factors that influence the number and duration of norms, and to investigate whether there is an optimal number of groups in a given configuration.

Using tag-based cooperation to investigate norm emergence allows us to observe the effects of various approaches. In particular, it is through norm establishment (i.e. reducing the number of tags) that the context assessment and rewiring techniques improve group effectiveness, and this interpretation may inform improvements to tag-based approaches. We have shown that tag groups have a similar lifecycle to that observed in more traditional notions of norms: establishment of norms, widening of adoption, and eventual collapse or replacement. The second area of future work aims to explore this lifecycle in more detail, with a particular focus on improving our techniques to obtain an appropriate number of norms and to increase their persistence. Importantly, since mutation can cause norm collapse it is important to ensure that there are alternative groups. Context assessment (with or without rewiring) is found to reduce the diversity of tags to a low level, and this suggests that further work is needed to develop techniques for coping with cheaters.

The third area of future investigation, and the most significant, is concerned with the network topology that defines agents' neighbourhoods. In our investigations we used a simple random with degree distribution  $P(k)$ , where  $k$  corresponds to the number of neighbours ( $k = n$ ), and observed that cooperation is successfully established. However, such a topology is unlikely to occur in real-world applications, which are more likely to exhibit scale-free or small-world topologies [10]. Further investigation is needed to evaluate our techniques with respect to such topologies, since network topology has been found in other settings to have a significant impact on the emergence of norms and conventions [9, 19, 25]. Although further investigation is needed, our results show that the tag-based approach has a potentially important difference from other techniques. For example, Pujol *et al.* [21] investigated the impact of network topology on convention emergence, using the example setting of the Highest Cumulative Reward action selection strategy [28] in the Coordination Game. They found that in low clustered environments (such as a random topology like ours) existing conventions are not replaced by newly emergent ones, which differs from our results in the context of tag-based cooperation. We see norms being established, and then being replaced by subsequent norms. More investigation is needed to establish the reason for this difference, but we suspect that it may be connected to our consideration of *multiple* norms, since the Coordination Game only contains two possible states (or norms).

There is a complex relationship between the evolution of tag groups, the number of tags, and the donation rate depending on the particular environment. Our final area of future investigation is to consider more realistic environments. In addition to the issues described above we will investigate techniques to cope with false information provided by malicious neighbours, and will incorporate *churn* into the system, such that agents may leave and join at runtime. By

exploring these areas we aim to develop techniques that ensure norm diversity and prolong norm duration in realistic environments.

## References

1. P. D. Allison. The cultural evolution of beneficent norms. *Social Forces*, 71(2):279–301, 1992.
2. W. B. Arthur. Inductive reasoning and bounded rationality. *The American Economic Review*, 84(2):406–411, 1994.
3. R. Axelrod. An evolutionary approach to norms. *American Political Science Review*, 80(4):1095–1111, 1986.
4. R. Axelrod, R. A. Hammond, and A. Grafen. Altruism via kin-selection strategies that rely on arbitrary tags with which they coevolve. *Evolution*, 58(8):1833–1838, 2004.
5. C. Bicchieri. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press, 2006.
6. R. Boyd and P. J. Richerson. The evolution of indirect reciprocity. *Social Networks*, 11(3):213–236, 1989.
7. J. C. Burguillo-Rial. A memetic framework for describing and simulating spatial prisoner’s dilemma with coalition formation. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-09)*, pages 441–448, 2009.
8. R. Dawkins. *The Selfish Gene*. Oxford University Press, 1976.
9. J. Delgado. Emergence of social conventions in complex network. *Artificial Intelligence*, 141(1-2):171–185, 2002.
10. J. Delgado, J. M. Pujol, and R. Sangüesa. Emergence of coordination in scale-free networks. *Web Intelligence and Agent Systems*, 1(2):131–138, 2003.
11. F. Flentge, D. Polani, and T. Uthmann. Modelling the emergence of possession norms using memes. *Journal of Artificial Societies and Social Simulation*, 4(4), 2001.
12. N. Griffiths. Tags and image scoring for robust cooperation. In *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-08)*, pages 575–582, 2008.
13. N. Griffiths and M. Luck. Changing neighbours: Improving tag-based cooperation. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-10)*, pages 249–256, 2010.
14. D. Hales and B. Edmonds. Applying a socially inspired technique (tags) to improve cooperation in P2P networks. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35(3):385–395, 2005.
15. W. D. Hamilton. The genetical evolution of social behaviour I. *Journal of Theoretical Biology*, 7(1):1–16, 1964.
16. W. D. Hamilton. The genetical evolution of social behaviour II. *Journal of Theoretical Biology*, 7(1):17–52, 1964.
17. J. H. Holland. *Hidden order: How adaptation builds complexity*. Addison-Wesley, 1995.
18. A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2):618–644, 2007.

19. J. E. Kittock. Emergent conventions and the structure of multi-agent systems. In *Lectures in Complex systems: The proceedings of the 1993 Complex systems summer school, Santa Fe Institute Studies in the Sciences of Complexity Lecture Volume VI*, pages 507–521. Addison-Wesley, 1993.
20. N. Masuda and H. Ohtsuki. Tag-based indirect reciprocity by incomplete social information. *Proceedings of the Royal Society B*, 274(1610):689–695, 2007.
21. J. M. Pujol, J. Delgado, R. Sangüesa, and A. Flache. The role of clustering on the emergence of efficient social conventions. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05)*, pages 965–970, 2005.
22. P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.
23. R. Riolo, M. Cohen, and R. Axelrod. Evolution of cooperation without reciprocity. *Nature*, 414:441–443, 2001.
24. G. Roberts and T. N. Sherratt. Does similarity breed cooperation? *Nature*, 418:499–500, 2002.
25. N. Salazar, J. A. Rodríguez-Aguilar, and J. L. Arcos. Robust convention emergence through spreading mechanisms. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-10)*, pages 1431–1432, 2010.
26. B. T. R. Savarimuthu, M. Purvis, and M. Purvis. Social norm emergence in virtual agent societies. In *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-08)*, pages 1521–1524, 2008.
27. S. Sen and S. Airiau. Emergence of norms through social learning. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 1507–1512, 2007.
28. Y. Shoham and M. Tennenholtz. Emergent conventions in multi-agent systems. In *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, pages 225–231, 1992.
29. Y. Shoham and M. Tennenholtz. On the emergence of social conventions: modeling, analysis and simulations. *Artificial Intelligence*, 94(1–2):139–166, 1997.
30. L. Steels. Self-organising vocabularies. In *Artificial Life V*, pages 179–184, 1996.
31. A. Traulsen. Mechanisms for similarity based cooperation. *The European Physical Journal B - Condensed Matter and Complex Systems*, 63(3):363–371, 2008.
32. D. Villatoro and J. Sabater-Mir. Dynamics in the normative group recognition process. In *Proceedings of the Eleventh conference on Congress on Evolutionary Computation*, pages 757–764, 2009.