The Role of Clustering on the Emergence of Efficient Social Conventions

Josep M. Pujol *, Jordi Delgado, Ramon Sangüesa † and Andreas Flache

Polytechnic University of Catalonia
Barcelona, Spain
{jmpujol,delgado,sanguesa}@lsi.upc.edu

University of Groningen
Groningen, The Netherlands
a.flache@ppsw.rug.nl

Abstract

Multiagent models of the emergence of social conventions have demonstrated that global conventions can arise from local coordination processes without a central authority. We further develop and extend previous work to address how and under what conditions emerging conventions are also socially efficient, i.e. better for all agents than potential alternative conventions. We show with computational experiments that the clustering coefficient of the networks within which agents interact is an important condition for efficiency. We also develop an analytical approximation of the simulation model that sheds some light to the original model behavior. Finally, we combine two decision mechanisms, local optimization and imitation, to study the competition between *efficient* and *attractive* actions. Our main result is that in *clustered* networks a society converges to an efficient convention and is stable against invasion of sub-optimal conventions under a much larger range of conditions than in a *non-clustered* network. On the contrary, in *non-clustered* networks the convention finally established heavily depends on its initial support.

1 Introduction

Social conventions, according to [Ullmann-Margalit, 1977], are a special type of norms related to coordination problems, that is, those regularities of behavior which are a result of being a solution of a recurrent coordination problem, which with time, turn normative. The emergence and stabilization of norms, including social conventions, is one of the fundamental problems of social sciences [Bendor and Swistak, 2001]. Intuitively, a social convention might be regarded as any rule of behavior or a behavioral constraint [Walker and Wooldridge, 1995]. A convention simplifies people's decision making problem by dictating how to act in certain situations. Therefore, social conventions help to reduce complexity and uncertainty, particularly when the environment is open and dynamic. From this perspective it is obvious that mechanisms underlying the emergence of social conventions are of great interest for Distributed Artificial Intelligence.

[Shoham and Tennenholtz, 1992] addressed the question whether and how the *emergence* of conventions is possible from decentralized, local interaction, as opposed to conventions designed and enforced a priori by a *central* authority. Shoham demonstrated that self-coordination based on a simple local optimization rule can emerge from initial disorder. Following Shoham's, many researchers have studied and improved his seminal work, proposing other strategies rather than HCR [Walker and Wooldridge, 1995], studying the time needed before the convention is established [Kittock, 1993], adding complex interaction patterns among agents [Delgado, 2002], and many others.

However, an important question remains open, when two potential different conventions exist which one will be established at the end? This question is not applicable to the original work of [Shoham and Tennenholtz, 1992], since the game was a *pure* coordination game, where no action pareto-dominates the other. Nevertheless, the question applies to all coordination games where one action is more efficient than the other. The discussion in social sciences gives no clear answer to why and under what conditions efficient conventions may prevail. As [Bendor and Swistak, 2001] point out, there are at least two conflicting positions in sociology, the *strong functionalism* thesis and the *rational choice* view. *Strong functionalism* claims that norms and conventions exist because they are functional for the group, that is, they yield optimal collective outcomes. According to this perspective, the system should always end up in an efficient convention. However, the *functionalist* approach has been criticized for its lack of a microfoundation. Adherents of the *rational choice* view on norms argue that individuals adopt norms only when it serves their self-interest to do so. This includes the possibility that mechanisms such as social control may stabilize conventions that are an individually efficient response to the given constraints, but are socially inefficient.

In section 2 we describe our model. In section 3 we describe the experimental results and analyze the necessary conditions to converge to the optimal convention. In section 4 we analyze the time efficiency of the emergence of conventions. In section 5 we integrate and analyze experimentally optimization and imitation partially competing decision mecha-

*This work was supported in part by EuropeAid @LIS TechNET project ALA2002/049-055/2209

†This work was supported in part by the Catalan Autonomous Government Project e-Catalunya

nisms in the model. Finally, main results are summed up in section 6.

2 Formal Model

We follow the conceptual framework introduced by [Shoham and Tennenholtz, 1992], and the extension proposed by [Delgado, 2002] that introduced complex networks as the underlying topology. .

2.1 The Coordination Game

A set of *N* agents must choose to play one of two possible actions: either *A* or *B*. Accordingly to its current action, or state, an agent interact with its neighbors receiving an outcome defined in payoff matrix *G*.

		Agent j A	Agent j B
Agent i	A	(x,x)	(u,v)
	B	(v,u)	(y,y)

Figure 1: Payoff Matrix of the Game G

The payoff matrix G defines a 2-person 2-choice symmetric coordination game provided that $x > v$ and $y > u$. The condition on the entries of G makes clear that to play the same action is the best choice. It is trivial to demonstrate that the game G has two nash-equilibrium, both agents playing either A or B. Most previous work focused in the study of *pure coordination games*: where $x = y = +1, v = u = -1$ [Walker and Wooldridge, 1995; Kittock, 1993; Shoham and Tennenholtz, 1997; Delgado, 2002]. In our approach we move beyond this work and allow for coordination games with equilibrium differing in social efficiency (i.e. $x > y$).

For simplicity we assume that coordination on action A is at least as profitable as coordination on action B: $x \geq y$. Thus, game G is defined as follows: $v = u = -1$, $y = +1$, and $x = \alpha$ provided that $\alpha \geq 1$. When $\alpha > 1$, coordination in B is a sub-optimal solution since there exists a pareto-efficient solution that pareto-dominates B, which is, coordination in action A.

2.2 Action Selection Rule and Dynamics

Our MAS is composed of N agents that interact only with its neighbors in the social network, playing the game G once per interaction. Every agent, say the kth, has memory M_k that records the M last interactions of agent k. The value of the position i of the memory M_k is the tuple $\langle a_k^i, p_k^i, t^i \rangle$, where a_k^i stands for the action played by k, p_k^i stands for the payoff received after playing action a_k^i, and t^i denotes the time the interaction took place. The initial action of the agents is set randomly with a probability r_B, which is the density of agents playing action B in the beginning.

Following [Shoham and Tennenholtz, 1992] we will use the *Highest Cumulative Reward* (HCR) action selection rule. Intuitively HCR says: if the accumulated payoff obtained from playing A is bigger than that from playing B then keep on A, otherwise change to action B. The HCR rule is very appropriate since it provides: 1) Locality: the selection function only depends on the agent's personal history. No global knowledge of the system is required, not even the payoff matrix of the game. 2) Adaptability: the agent learns from its experience without assuming cognitive capabilities. These characteristics are very important in MAS.

The dynamics of the system are as follows. At each time step t, an agent k is randomly activated. Once the agent is activated, it plays the game G with an agent randomly chosen from k's neighborhood, say agent l. The result of the interaction is stored into agent k's memory M_k, removing the oldest entry if necessary. Finally, agent k must decide whether to change its action or not. To do so it uses the *Highest Cumulative Reward* rule. Agent k will compute the payoff received for using action S on the last M activations: $P_S^k = \sum_{i:a_k^i=S} p_k^i$, where $S = \{A|B\}$. Agent k will switch to action $\overline{S}$ if $P_{\overline{S}}^k > P_S^k$. Agent l also carries out the memory storage and the action updating while the rest of agents remain still. The system ends once all agents play either action A or B, which means, that a convention on either A or B have been established.

2.3 Underlying Topology

To model qualitatively different interaction structures, we use several graph models that recently have been shown to have profoundly different effects on cooperation and diffusion dynamics in MAS [Watts and Strogatz, 1998; Albert and Barabási, 2002; Delgado, 2002]. The chosen graphs are: a) Random graphs: $R_N^{\langle k \rangle}$, where N is the number of nodes, and $\langle k \rangle$ is the average connectivity, that is, the average size of node's neighborhood. Random graphs have a clustering coefficient that tends to zero. The average path length grows logarithmically in function of N, the number of nodes. b) Regular graphs: C_N^k, regular graphs display an extremely high clustering coefficient, while its average path length and diameter grows linearly. Which means, for big graphs the average path length is very long, which does not agree with empirically studied networks. However, regular graphs display the close-knit property due to its high clustering coefficient, which does agree with empirical studies. c) Small-world graphs: $W_N^{\langle k \rangle, p}$, these are *highly-clustered* graphs (like regular graphs) with small average path length (like random graphs). This is the small-world property. We chose the [Watts and Strogatz, 1998] model as model of small-world graphs, where p is the rewiring probability. d) Scale-free graphs: $S_N^{\langle k \rangle, -\gamma}$, these are graphs with a connectivity distribution $P(k)$ of the form $P(k) \propto k^{-\gamma}$. The connectivity degree, the number of neighbors of a node, decays as a potential law. This favors the so-called *fat-tail* phenomena; few nodes with an extreme high connectivity. We chose the [Barabási and Albert, 1999] model as a model of scale-free graph.

Recent studies on empirical networks show that neither regular nor random graphs appear in nature. Noticing this, [Delgado, 2002] studied the effect of *complex* networks (small-world and scale-free networks) on the emergence of coordination and found that this class of networks were as efficient as the complete graph in terms of time to reach a convention, $O(NlogN)$ compared to $O(N^3)$ from regular

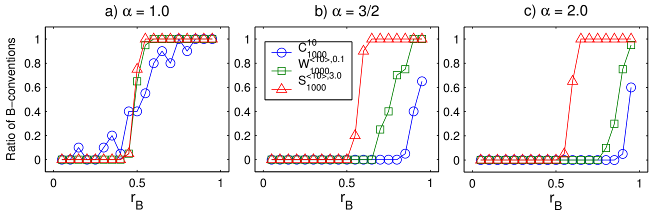

Figure 2: Ratio of conventions to action B. The x-axis is the initial density of agents playing action B: r_B, ranging from .05 to .95. The y-axis is the ratio of B-conventions, that is, the number of replications that ended up in all agents playing B over the total number of replications (20). In sub-figure a) there is no efficient convention, $\alpha = 1$. In sub-figure b) and c) all agents playing action A is the efficient convention, $\alpha = \frac{3}{2}$ and $\alpha = 2$ respectively.

graphs, already shown by [Kittock, 1993]. It is important to stress that many empirical social networks such as the collaboration network among actors, the co-authorship network of researchers on mathematics, the researchers on neuroscience and many others are classified as scale-free networks. However, for these empirical social networks the clustering coefficient found is very high: 0.79, 0.59, 0.76 respectively. Let us show the clustering coefficient of graphs we will use in the experiments: C^{10}_{1000}=0.666, $W^{<10>,0.1}_{1000}$=0.492, $S^{<10>,-3}_{1000}$= 0.0433 and $R^{<10>}_{1000}$=0.0081. Notice that empirical social networks are much more clustered than the scale-free networks yield by the [Barabási and Albert, 1999] model. But we still use this model since scale-free models are focused in reproducing the connectivity degree rather than clustering. This fact must be kept in mind when analyzing the results and conclusions. For a comprehensive survey in complex networks see [Albert and Barabási, 2002].

3 HCR-Model Experimental Results

We use a graph classification based on the clustering coefficient, regular and small-world graphs will be called *highly-clustered* graphs, and scale-free and random graphs will be called *low-clustered* graphs. For the sake of clarity we did not include results on random graphs in the figures, since they behave like scale-free graphs.

Let us stress which are the key manipulations in the parameters of our model. We systematically variate the proportion of agents who initially adopt the convention B (r_B), the efficiency gap α between the two conventions, and finally, the type of network, in particular its clustering coefficient, that defines the agent's neighborhood.

Another thing to mention before getting into the details is that there are two clear phases: 1) all the replications converging to A, and 2) all the replications converging to B. There is as well a space between these two phases where the system converges, with some replications to A and some to B, which we call *transitional* space, that is wider or narrower depending on the underlying topology and α. There is a critical point r^*_B that sets the boundaries between the two phases; and the *transitional* space may be defined with an ε, such that the result is: $r^*_B \pm \varepsilon$.

Now, let us comment the results of the experiments on the HCR model. It is helpful to first consider the case where both conventions are equally efficient ($\alpha = 1$), which is the case of a *pure coordination game*. The results for this case are displayed in sub-figure 2.a. When the initial number of agents playing B is less than the half of the population, that is, $r_B < \frac{1}{2}$, the system ends up establishing the convention on B. And, when $r_B > \frac{1}{2}$ the convention on A is established. The subfigure shows in particular that this result is widely independent of the network topology. It also shows that the transitional space is very narrow except for the regular graphs.

This pattern changes when convention differ in efficiency. Sub-figures 2.b and 2.c show the case where the payoff for action A exceeds the payoff for B ($\alpha > 1$). To understand the results, consider an example. Let $\frac{3}{4}$ of the population follow the action B, and $\frac{1}{4}$ follow the action A. In this case, although B is the initially chosen action for most of the agents, coordination in A is more efficient since it yields a better payoff. Thus, which will be the final convention agreed by the whole population? The answer depends on 1) how much better off, more efficient, is the action A over B, denoted by α. And 2) the underlying topology. If $\alpha = \frac{3}{2}$ (sub-figure 2.b) the final convention will be B when the underlying topology corresponds to a random or scale-free graph, A when having a regular graph, and can be both when having a small-world graph. By increasing the efficiency of coordination on A to $\alpha = 2$ (sub-figure 2.c) both small-world and regular graph converge to convention A, whereas random and scale-free graph still converge to convention B. The explanation behind the result is striking, *low-clustered* graphs seems very sensitive to the initial population density, whereas *highly-clustered* graphs behave in the opposite way, they are more sensitive to the efficiency of a particular action.

The conclusion derived from these results can be summed up as follows: *low-clustered* agent communities where a convention already exists will not be infected by a set of agents who play a new action, even though it is more efficient. Conversely, *highly-clustered* agent communites can be infected by a new action if the new action is more efficient, replacing the current convention for the convention on the efficient action.

Therefore, *highly-clustered* agent communities are more innovative, or adaptive, since a new action can be spread and finally established as a convention. The drawback would be that this community would be unstable due to its receptiveness to new action, and in the transient time needed to reach a stable new convention no coordination will exist. In contrast, *low-clustered* agent communities are very stable since the infection with a new convention is unlikely to happen, but on the other hand they are reluctant to adopt new actions even though they are more efficient, thus, becoming very conservative and static communities. We must remark that studies on empirical social networks have shown that these networks are very clustered [Albert and Barabási, 2002].

We do not provide a proof of the system's convergence. Nevertheless, throughout all the simulation runs, with their corresponding replications, the system has always converged to a convention with an upper bound of $O(N^3)$.

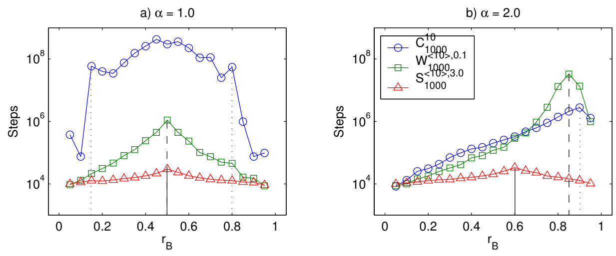

Figure 3: Time to converge to a convention, the value is the mean over 20 replications. The x-axis is the initial density of agents playing action B: r_B, ranging from .05 to .95. The $y-$axis is the number of interactions (steps) before reaching a convention. Note the logarithmic scale. In the left sub-figure both conventions, A or B. In the right sub-figure A is the optimal convention since $\alpha = 2$.

On figure 3 (left), α is set to 1, therefore both conventions are equally efficient. It can we observed that the regular graph takes much longer to converge than the rest of graphs, which is completely consistent with the findings of [Kittock, 1993] and [Delgado, 2002]. Notice the existence of peaks in $r_B = \frac{1}{2}$ in all the graphs except regular. These peaks are close to the critical point $r_B^* = \frac{1}{2}$ where the transition takes place. This is a typical behavior, when the system is close to a transition boundary it takes longer to converge [Yeomans, 1992]. The regular graph, however, is a particular case since it shows a *plateau* rather than a peak. This is due to its wide *transitional* state, in which different replications of the same setting can converge to different conventions.

On figure 3 (right) the convention A is more efficient that convention B, since playing action A yields payoff $\alpha = 2$, when playing action B yields a payoff 1. In this case the *plateau* of the regular graph does not appear because its *transitional* space narrowed as shown in figure 2.c. Thus, we see only the peaks in r_B where the transition takes place $r_B^* = \{.6, .85, .90\}$ for scale-free, small-world, and regular graphs respectively, which correspond to the transitions observed in figure 2.c. Again, we observe that the convergence is more inefficient in time when the initial r_B is close to r_B^*. It is important to remark that the regular graph has improved its efficiency in reaching the convention, that is, it is more efficient than the small-world graph when $\alpha = 2$. As was shown by [Young, 1999], convergence to a risk-dominant equilibrium, which in our case corresponds to the pareto-efficient one, is surprisingly rapid provided a close-knit (clustered) graph. The scale-free graph is always the most efficient graph in terms of convergence no matter α, however, scale-free graphs do not allow convergence to the efficient convention unless the initial number of agents playing A is very big, 40% compared to the 15% needed with a small-world graph, with α set to 2).

4 Analytical Model

First of all, we replace the role of past interactions stored in agent i's memory (M_i) by the current state of agent i's neighbors; notice that in our model we set the memory size to the average connectivity, $M = \langle k \rangle$. Instead of applying the HCR-rule over M_i we will transform it as follows: Let us take k_S as the number of neighbors in the same state as agent i, and $k_{\overline{S}}$ the number of neighbors in the opposite state. Therefore, the update is performed if $k_{\overline{S}} > \phi k_S$, where ϕ is the payoff yield by the payoff matrix G when playing (S, S). Therefore, the switch from playing A to play B is done when $k_B > \alpha k_A$, since $G(A, A) = \alpha$. Similarly, the switch from playing B to play A is done when $k_A > k_B$, since $G(B, B) = 1$. The probability of updating the state is defined by equation 1, we decided to introduce some stochasticity to compensate the fact that now the update is calculated by the current state of the neighbors instead of by the interactions, as the HCR-rule does.

$$f_{\beta,\phi}(k_{\overline{S}}) = \frac{1}{1 + e^{\beta(\phi-(\phi+1)\frac{k_{\overline{S}}}{k_S+k_{\overline{S}}})}} \quad (1)$$

We will use what in physics is called a *mean-field* argument [Pastor-Satorras and Vespignani, 2001]. Let $N_B(t)$ be the number of agents playing action B at time t, and $\rho(t) = \frac{N_A(t)}{N}$ be the density of agents playing B. A first approach is to assume the following *homogeneity* condition: for every agent with k neighbors, the number of neighbors in state B is $k_B(t) \simeq k\rho(t)$. This condition is completely fulfilled for random graphs, and approximately fulfilled for scale-free and small-world graphs (when $p \to 1$). Nevertheless, this *homogeneity* condition is not fulfilled for regular and for small-world graphs for low values of p. What breaks the *homogeneity* is the clustering coefficient, for *low-clustered* graphs this condition holds since the global density of agents in state S corresponds to the proportion of neighbors in state S. Intuitively, the clustering coefficient can be defined as the probability that a node i and a node j have a link provided that node l has a link to both i and j. Thus, when clustering tends to 0 the node's neighborhood is a good sample of the graph. However, when clustering coefficient is high the node's neighborhood is not a sample of the graph, since its neighbors form a *clique*, a *close-knit* group. Therefore, we propose a new *homogeneity* condition that takes clustering (cc) into account. Let us define cc as the clustering coefficient, provided that cc is the probability of agent l's neighbors being also neighbors, $(1 - cc)k$ is the number of neighbors which are not in the l's *clique*, and to whom the previous *homogeneity* condition holds. Therefore, for an agent playing A with k neighbors, the number of neighbors in the opposite state (B) is $k_A \simeq (1 - cc)k\rho(t)$, which is the *cc-biased-homogeneity* condition.

Now we can write an equation for the evolution of $\rho(t)$. First, notice that the variation of $\rho(t)$ after a small time interval Δt is proportional to Δt, that is, $\rho(t + \Delta t) = \rho(t) + \frac{\partial\rho(t)}{\partial t}\Delta t + O(\Delta t^2)$. Then, we can neglect the $O(\Delta t^2)$ term (since we perform a continuum approximation $\Delta t \to 0$) and compute the variation of $\rho(t)$ as the balance between the agents switching from action A to B and the agents switching from action B to A. On one hand, the fraction of agents in A (that is, $1 - \rho(t)$) that change to state B in a time interval Δt is the product $(1 - \rho(t))f(\rho(t))\Delta t$, provided Δt is small enough; on the other hand, the fraction of agents that switch from action B to A in Δt is $\rho(t)f(1 - \rho(t))\Delta t$, also for small

Table 1: Fixed points ρ^* of the Analytical Model: those are the critical point such that $\frac{\partial \rho}{\partial t} = 0$. In brackets the critical points observed experimentally with the HCR model (r_B^*) (figure 2). The parameters that model the graph are the clustering coefficient cc, and the average connectivity $\langle k \rangle$ set to 10.

$\rho^*(r_B^*)$	$\alpha = 1$	$\alpha = \frac{3}{2}$	$\alpha = 2$
cc=0.666	0.5 (0.5)	0.831 (0.90)	$\nexists$ (0.95)
cc=0.492	0.5 (0.5)	0.68 (0.7)	0.827 (0.85)
cc=0.0433	0.5 (0.5)	0.566 (0.55)	0.614 (0.60)
cc=0.0081	0.5 (0.5)	0.560 (0.55)	0.607 (0.60)

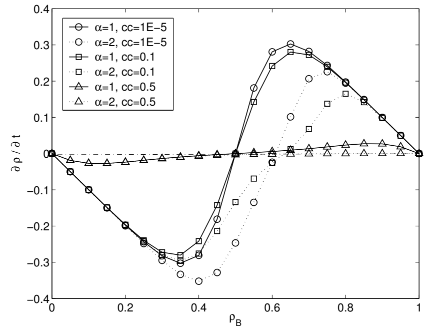

Figure 4: Study of the fixed points of the density equation 3. ρ is the density of agents playing B. Clustering coefficient is set to $cc = \{1E-05, 0.1, 0.5\}$, α is set to $\{1, 2\}$

Δt. Thus after $\Delta t \rightarrow 0$, the mean-field equation for $\rho(t)$ can be written as

$$\frac{\partial \rho(t)}{\partial t} = (1 - \rho(t))f(\rho(t)) - \rho(t)f(1 - \rho(t)) \qquad (2)$$

After substitution of $f_{\beta,\phi}$ to which the *cc-biased-homogeneity* condition has been applied. And setting $\beta = \langle k \rangle$ and $\phi = (S, S)$ the equation reads

$$\frac{\partial \rho}{\partial t} = \frac{1-\rho}{1+e^{\langle k \rangle(\alpha-(\alpha+1)(1-cc)\rho)}} - \frac{\rho}{1+e^{\langle k \rangle(1-2(1-cc)(1-\rho))}} \qquad (3)$$

We want to study the stable fixed points of 3 since these will give us information on the final state of the system. Thus we must find the solutions of $\frac{\partial \rho}{\partial t} = 0$. As we can see in 4, stable fixed points are $\rho^* \sim 0$, and $\rho^* \sim 1$ (these have been computed numerically) and the unstable fixed-point lies in $(0, 1)$. In table 1 we display the unstable fixed points of the density equation, which are possible critical points of the HCR model provided our assumptions (see above) are correct. Furthermore, in table 1 we find a comparison between the analytical unstable fixed points and those critical points coming from the simulation of the HCR model. Notice that the plausibility of the simplifying assumptions behind our analytical model is supported by the agreement between analytical and experimental results (remind that r_B was sampled with a resolution of .05).

Figure 4 shows the variation in ρ, $\triangle\frac{\partial \rho}{\partial t}$, for different α and cc. We can see the effect of α enlarging the basin of attraction of convention A. On the other hand cc has the effect of reducing the amount of variation due to the effects of the *cc-biased-homogeneity* condition. Consequently, the time elapsed to reach a convention will be longer, and fluctuations in initial conditions will have a bigger impact. As a matter of fact, the experimental results on the HCR model show us these two consequences apply for highly-clustered graphs such as the regular graph. The convergence time is much higher compared to non-clustered graphs. And the transitional space, where the system can converge either in B or in A for the same initial parameters, is wider.

5 The Role of Imitation

In this section we will modify the model based of the HCR-rule introducing an imitation propensity i_S, which is the probability that, after a dyadic interaction where at least one agent was playing action S, both agents end up playing action B regardless of the HCR action updating rule. By introducing an imitation propensity we model the effect of having an *attractive* action which more likely to be chosen. Imitation is considered as the key factor of the adoption of norms, and by extension, conventions. How does the existence of an *attractive* action affect the final convention reached by the agents? And what if the *attractive* action happens to be sub-optimal? In our model coordination on action A is most efficient solution (if $\alpha > 1$), however, action B (if $i_B > 0$) might become a better replicator since it can be adopted by imitation as well as adopted by the learning process (HCR-rule).

The effect of i_B heavily depends on the clustering of the underlying topology. For instance, when $\alpha = 2$ and $i_B = .4$ the clustered communities adopts the *attractive* convention (B) over the efficient one (A) regardless of the initial density (r_B), while *non-clustered* communities will still adopt A provided $r_B < .25$. Therefore, it might be derived that *non-clustered* communites are more resilient against *attractive* conventions in favor of efficient conventions. This result would seem to contradict the previous claim, that clustered communites are better off in converging to the efficient convention. However, when i_B is small enough, for example when $i_B = .1$ the opposite effect is observed, clustered communities keep on converging to the *efficient* convention regardless of the initial density of agents. This two-fold behavior is perfectly clear in the case of the regular graph (left column of figure 5). There is a threshold i_B^* under which the system ends up in the efficient convention, and over which the system ends up in the *attractive* convention. For example $i_B^* = .2$ when $\alpha = 2$ and the underlying topology is the regular graph. Notice, that this threshold is not found in the case of the small-world graph whose clustering coefficient is high, although not so high as the case of the regular graph. However, even without the threshold we observe a similar behavior than in the regular graph although more progressive. If we compare it against the behavior of *low-clustered* graphs, we find again that the system is more resilient to the *attractive* action invasion for low values of i_B. For high values, on the contrary, the system is very receptive to an invasion of agents playing the *attractive* action. We must state that the model is interesting for small values of i_B, for high values of i_B the dichotomy of the agent to choose between the efficient or the *attractive* action dissapear and becomes an epidemic spread model [Pastor-Satorras and Vespignani, 2001].

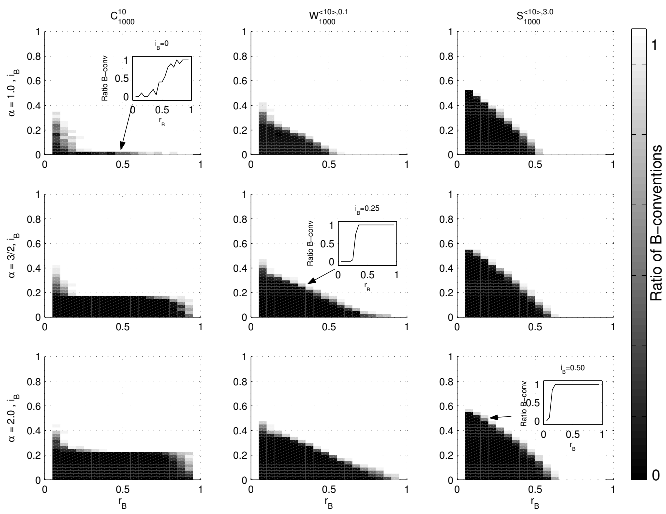

Figure 5: Effect of imitation propensity (i_B) on the ratio of conventions on action B. i_B ranges from $[.0.975]$ in a .0275 resolution, r_B ranges from $[.05.95]$ in a .05 resolution. The sub-figures within the figures are for the sake of clarity: setting the parameter i_B we observe the evolution of B-conventions depending on r_B, which correspond to the way the previous results were shown (when $i_B = 0$).

6 Concluding Remarks

Our research aimed to identify conditions under which local interactions in a multi agent community can give rise to an efficient convention. We have shown that the key factor is the clustering of the underlying agent's social network. When communities are *highly-clustered* the system converges to the pareto-efficient action even though the initial population choosing that action was clearly a minority. This suggests that the efficient convention is a stable convention because it cannot be invaded by a set of agents playing another sub-optimal action. However, a sub-optimal convention can be replaced by a set of agent playing an action that yields a better payoff. Accordingly, when the clustering coefficient is high the system always converges to the most efficient convention and this convention is stable. On the other hand when clustering tends to zero the adopted convention depends solely on the density of agents following an action. If the majority of agents play the sub-optimal action the inefficient convention will be established, and it will be stable. To back up our findings, we provided an analytical approximation that reproduces the results observed in our model based in the HCR-rule. To do so, we had to introduce a new *homogeneity* condition which let us work in clustered graphs, where the classical mean-field *homogeneity* condition is not met.

In accordance with the *strong functionalism* thesis from classical sociology, we found that in certain graphs the agent system was capable to find and maintain the optimum in the stable state. However, this only applies to *highly-clustered* communities, which resemble many empirical social networks. At the same time, our model also corresponds in two respects to the view that rational choice theorists in sociology take on social norms. First, we have shown that global efficiency arises from individual goal oriented actions. Second, we found that under certain conditions optimizing individual actions fail to generate socially efficient outcomes, a problem that is central to the contemporary discussion about the emergence of conventions and norms.

To conclude, our results seem to correspond more with a rational choice on norms than with the strong functionalism thesis. We have shown that socially optimal conventions can arise from individual optimization, but there is no guarantee that this happens. In this sense, our model matches well the ample evidence of examples of suboptimal conventions, for example in market processes. We believe that a part of the explanation for this may lie in the competition between optimizing and imitation that we have addressed with our model. We have shown that imitation processes make it possible for a sub-optimal yet *attractive* action to overthrow the efficient action, and become stable, provided that its *attractiveness* is high enough to be worthy of imitation.

References

[Albert and Barabási, 2002] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Review of Modern Physics*, 74:47–97, 2002.

[Barabási and Albert, 1999] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[Bendor and Swistak, 2001] J. Bendor and P. Swistak. The evolution of norms. *American Journal of Sociology*, 106(6):1493–1545, 2001.

[Delgado, 2002] J. Delgado. Emergence of social conventions in complex networks. *Artificial Intelligence*, 141:171–185, 2002.

[Kittock, 1993] J.E. Kittock. Emergent conventions and the structure of multi-agent systems. In *In Proc. of the 1993 Santa Fe Inst. Complex Systems Summer School*, pages 384–389, 1993.

[Pastor-Satorras and Vespignani, 2001] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Phys. rev. Lett.*, 86:3200–3203, 2001.

[Shoham and Tennenholtz, 1992] Y. Shoham and M. Tennenholtz. Emergent conventions in multi-agent systems. In *In Proceedings of Knowledge Representation and Reasoning*, pages 225–231, 1992.

[Shoham and Tennenholtz, 1997] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: Modeling, analysis and simulations. *Artificial Intelligence*, 94:139–166, 1997.

[Ullmann-Margalit, 1977] E. Ullmann-Margalit. *The Emergence of Norms*. Ofxord Univ. Press., 1977.

[Walker and Wooldridge, 1995] A. Walker and M. Wooldridge. Understanding the emergence of conventions in multi-agent systems. In *In Proceedings of 1st International Conference on Multi-Agent Systems*, pages 384–389, 1995.

[Watts and Strogatz, 1998] D.J. Watts and S.H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.

[Yeomans, 1992] J.M. Yeomans. *Statistical Mechanics of Phase Transitions*. Oxford University Press, 1992.

[Young, 1999] H. Peyton Young. Diffusion in social networks. Technical report, Center on Social and Economics Dynamics, Brookings Institution, 1999.