

Modeling the Emergence and Convergence of Norms

Logan Brooks

University of Tulsa
Tulsa, Oklahoma, USA
logan-brooks@utulsa.edu

Wayne Iba

Westmont College
Santa Barbara, CA, USA
iba@westmont.edu

Sandip Sen

University of Tulsa
Tulsa, Oklahoma, USA
sandip@utulsa.edu

Abstract

In many multi-agent systems, the emergence of norms is the primary factor that determines overall behavior and utility. Agent simulations can be used to predict and study the development of these norms. However, a large number of simulations is usually required to provide an accurate depiction of the agents' behavior, and some rare contingencies may still be overlooked completely. The cost and risk involved with agent simulations can be reduced by analyzing a system theoretically and producing models of its behavior. We use such a theoretical approach to examine the dynamics of a population of agents playing a coordination game to determine all the norms to which the society can converge, and develop a system of linear recurrence relations that predict how frequently each of these norms will be reached, as well as the average convergence time. This analysis produces certain guarantees about system behavior that cannot be provided by a purely empirical approach, and can be used to make predictions about the emergence of norms that numerically match those obtained through large-scale simulations.

1 Introduction

Many aspects of daily life are governed not by laws or moral injunctions but rather by *norms*. Norms or conventions routinely guide the choice of behaviors in human societies and plays a pivotal role in determining social order [Hume, 1978]. Norms serve to simplify and regularize much of the complexity of the world around us [Lewis, 1969]. Norms typically become codified into law, but do not start out that way. An example of a norm might be deciding to always drive on either the left side or the right side of the road. The important characteristic of a norm is that it guides behavior where moral or rational reasoning does not provide clear guidance. Norms may also form in contexts where rational choices cannot be applied because of the myopic behaviors and the limited processing abilities of individuals. Consequently, such norms may prove to be suboptimal.

The systematic study and development of robust mechanisms that facilitate emergence of stable, efficient norms

via learning in agent societies promises to be a productive research area that can improve coordination in agent societies. Correspondingly, there has been a number of recent, mostly empirical, investigations in the multiagent systems literature on norm evolution under different assumptions about agent interaction frameworks, societal topology, and observation capabilities [Delgado *et al.*, 2003; Kittock, 1994; Sen and Airiau, 2007; Shoham and Tennenholtz, 1997]. There is an associated need to develop analytical frameworks that can predict the trajectory and dynamics of emergence and convergence of society-wide behaviors [van Dyke Parunak, 2002].

Toward this end, we mathematically model the emergence of norms in societies of agents who adapt their likelihood of choosing one of a finite set of options based on their experience from repeated one-on-one interactions with other members in the society. The goal is to study both the process of emergence of norms as well as to predict the likely final convention that is going to emerge if agents had preconceived biases for certain options.

2 Related Research

Recent literature in multiagent systems show a significant increase in interest and research on normative systems. Norms or conventions can be substituted as external correlating signals to promote coordination. These aspects of norms or conventions have merited in-depth study of the evolution and economics of norms in social situations [Epstein, 2001; Posch, 1995; Young, 1996].

Norms can also be thought of as focal points evolved through learning [Young, 1996] that reduce disagreement and promote coherent behavior in societies with minimal oversight or centralized control [Coleman, 1987]. Norms can therefore have economic value to agents and help improve their efficiency.

While some research on norm emergence assumed uniform interaction probabilities between agents [Sen and Airiau, 2007; Shoham and Tennenholtz, 1997] others have studied the effect of topologies of agent relationships [Delgado *et al.*, 2003; Kittock, 1994]. This body of research can be further divided into *interaction-based learning* [Sen and Airiau, 2007] versus *observation-based adoption* approaches to norm emergence [Delgado *et al.*, 2003]. In the interaction-based learning mode, agents learn utility estimates of their behavior

choices and over time converge on a particular behavior that becomes the norm in the society. In this mode, agents actually do not even need to observe others' behaviors, and agents over time adopt biased behaviors trial-based on their own experiences of what produces higher utility. In the observation-based learning mode, agents' behaviors must be fully observable and typically there is no direct consideration of utilities. The basic idea is to adopt behaviors of successful individuals.

3 Agent Interaction Model

We first characterize our model of agent interactions that drives the emergence of conventions in the society. Consider a population of n agents, each of which is capable of performing two distinct actions. Pairs of agents are randomly selected to interact. In each interaction, each agent selects one of its two available actions based on their current biases. The goal for the agents is to coordinate by selecting identical actions, but they cannot communicate or negotiate beforehand to coordinate their action choice. Based on the outcome of their interaction (coordination for identical actions or conflict otherwise), agents adjust their biases for their selected actions. For the purposes of our formal analysis, we can consider each bilateral agent interaction to consist of three phases: *pairing of agents*, *action selection*, and *bias revision*.

In our model, an agent i contains a single value, p_i , representing the bias or probability of selecting action 1 (this model can be expanded to include a vector of size $m - 1$ for if agents could choose from m different actions). Agent i selects the other action with the complementary probability, $(1 - p_i)$. At each time step, l agents are randomly selected to interact with each other; based on the outcome of the interaction, each agent i chosen to interact will change its p value by no more than some constant x using a predefined update rule.

Consider the scenario of a society of agents choosing over time to adopt a convention of driving on the right or left side of the road. Each agent can choose one of two actions: an action value of $+1$ indicates that the agent will choose to drive on the right side of the road, while a value of -1 corresponds to driving on the left side. Let p_i^t represent the bias or likelihood that agent i will choose action $+1$ at time t . We assume that all agents have the same initial bias: $(\forall i)(p_i^0 = \bar{p}_0)$. At each time step, two distinct agents are randomly selected with uniform probability from the population. Each agent selected calculates a random number r_i from $U[0, 1]$ and chooses an action value :

$$act_i^t = \begin{cases} +1, & r_i^t < p_i^t \\ -1, & r_i^t \geq p_i^t \end{cases}.$$

If both agents in the pair choose the same action, both will positively reinforce the action that they chose and play it with greater frequency in the future. If their actions did not coordinate, then each agent reduces the frequency with which it plays the action that it chose at that time step. Mathematically, this can be expressed by:

$$p_i^{t+1} = \max\{0, \min\{1, p_i^t + act_i^t \cdot act_j^t \cdot \Delta_i(act_i)\}\},$$

where act_j^t is the action taken by the other agent in the pair, and $\Delta_i(act_i)$ is an update function. We have used the two following update functions.

- Constant update rule: $\Delta_i(act_i^t) = x \cdot act_i^t$, where $x \in (0, 1)$ is a constant. This rule assumes that each agent will update its preferences in constant increments.
- Proportional update rule:

$$\Delta_i(act_i^t) = x \cdot \begin{cases} +\beta + (1 - \beta)p_i^t, & act_i^t = +1 \\ -\beta - (1 - \beta)(1 - p_i^t), & act_i^t = -1 \end{cases},$$

where $\beta \in (0, 1)$ is a constant. This rule assumes that the agents that are less biased towards either action will be less willing to make large changes in their biases.

To aid our analysis and provide better insight into the norm emergence process, we will make some simplifications to the model. If $1/x$ is an integer and an agent is initialized with a p value that is a multiple of x , then we find that that agent's p value will always remain a multiple of x . If all the n agents have bias values constrained this way, then the population average, $\bar{p} = \frac{1}{n} \sum_{i=1}^n p_i$, can only assume values that are multiples of $\frac{x}{n}$, or $\frac{n}{x} + 1$ distinct values.

We can then predict the average convergence time and the final convergence value for a given population bias, $\bar{p}$. Let $P(\bar{p})$ represent the estimated average convergence value for any population with a current average bias of $\bar{p}$, and $T(\bar{p})$ be the expected number of time steps before converging. We ignore the corrections for values that fall below 0 or above 1. Consequently, we can express the value of $P(\bar{p})$ as a weighted average of the P values for all distributions that could be reached at the next time step:

$$P(\bar{p}) = \sum_{j=0}^{n/x} [F(\bar{p} \cdot n/x, j) \cdot P(j \cdot x/n)],$$

where $F(i, j)$ is the expected probability that a population with mean bias $i \cdot x/n$ will change its mean to $j \cdot x/n$ due to a single interaction. For the constant update rule defined above, this is equivalent to

$$P(\bar{p}) = (1 - \bar{p})^2 P\left(\bar{p} - 2\frac{x}{n}\right) + 2(1 - \bar{p})\bar{p}P(\bar{p}) + \bar{p}^2 P\left(\bar{p} + 2\frac{x}{n}\right).$$

In this expression, $(1 - \bar{p})^2$ is the probability that two randomly selected agents will both choose to drive on the left side of the road ($act_i^t = act_j^t = -1$), $\bar{p} - 2\frac{x}{n}$ is the resulting $\bar{p}$ value at the next time step, and $P(\bar{p} - 2\frac{x}{n})$ is the average value to which $\bar{p}$ will converge in this case. The second term, $2(1 - \bar{p})\bar{p}P(\bar{p})$, deals with the case in which the two agents do not coordinate and $\bar{p}$ does not change, since the increase in the bias value of one agent is exactly offset by the decrease in the bias value of the other agent. The last term corresponds to when the two agents coordinate and increase their p values.

A similar expression can be used for $T(\bar{p})$, with an additional term of 1 to represent the current time step.

$$\begin{aligned} T(\bar{p}) &= 1 + \sum_{j=0}^{n/x} [F(\bar{p} \cdot n/x, j) \cdot T(j \cdot x/n)] \\ &= 1 + (1 - \bar{p})^2 T(\bar{p} - 2x/n) \\ &\quad + 2(1 - \bar{p})\bar{p}T(\bar{p}) + \bar{p}^2 T(\bar{p} + 2x/n). \end{aligned}$$

We need some values of P and T to solve the system. Since $\bar{p}$ values of 0 or 1 indicate that the population has converged,

we have definite values of P and T at these points: $P(0) = 0$, $P(1) = 1$ (everyone chooses the same action hereafter), and $T(0) = T(1) = 0$ (population already converged). The above equations for P and T form nearly-diagonal linear systems of equations, which can be solved in $O(n/x)$ time and space to give a close approximation of the average convergence time and values obtained in the simulations. For the proportional update rule given above, discretization of the distribution of individual agent p values poses a significantly more difficult problem. However, a fairly accurate model can be obtained by discretizing the possible $\bar{p}$ values into $m \cdot n/x + 1$ intervals, where m is some positive integer constant. The average value over all agents of $\Delta_i(act_i^t)$ when $act_i^t = +1$ is given by:

$$\begin{aligned}\Delta^+ &= \frac{1}{n} \sum_{i=1}^n x[\beta + (1 - \beta)p_i] \\ &= x \frac{n\beta + (1 - \beta)n\bar{p}}{n} \\ &= x(\beta + (1 - \beta)\bar{p}).\end{aligned}$$

Similarly, the average value of $\Delta_i(act_i^t)$ for $act_i^t = -1$ is $\Delta^- = -x(\beta + (1 - \beta)(1 - \bar{p}))$. If we use this average value of Δ_i for all agents, then we can write a set of relationships similar to those for the constant update rule:

$$\begin{aligned}P(\bar{p}) &= (1 - \bar{p})^2 P\left(\bar{p} + \frac{2\Delta^-}{n}\right) \\ &\quad + 2(1 - \bar{p})\bar{p}P\left(\bar{p} + \frac{\Delta^+ + \Delta^-}{n}\right) + \bar{p}^2 P\left(\bar{p} + \frac{2\Delta^+}{n}\right), \\ T(\bar{p}) &= 1 + (1 - \bar{p})^2 T\left(\bar{p} + \frac{2\Delta^-}{n}\right) \\ &\quad + 2(1 - \bar{p})\bar{p}T\left(\bar{p} + \frac{\Delta^+ + \Delta^-}{n}\right) + \bar{p}^2 T\left(\bar{p} + \frac{2\Delta^+}{n}\right).\end{aligned}$$

As with the previous expressions, $P(0) = 0$, $P(1) = 1$, $T(0) = 0$, and $T(1) = 0$. These relationships can also be transformed into nearly diagonal matrices and solved in $O(m \frac{n}{x})$ time and space.

The expressions that we have defined for P and T have neglected any clipping effects for p values outside the interval $[0, 1]$. While these expressions provide a good approximation when all agents begin with the same p value, they lose accuracy when the initial distribution is more spread out. In order to incorporate the clipping effect into the model, additional terms must be added to the expressions for P and T . The modified expression for $T(\bar{p})$ is as follows:

$$T(\bar{p}) = \begin{cases} 0, & \bar{p} \leq 0 \\ \begin{aligned} &1 + L_{\bar{p}}^2 T\left(\bar{p} + \frac{2\Delta^-}{n}\right) \\ &+ 2L_{\bar{p}}Z_{\bar{p}}T\left(\bar{p} + \frac{\Delta^-}{n}\right) \\ &+ 2R_{\bar{p}}Z_{\bar{p}}T\left(\bar{p} + \frac{\Delta^- + \beta/n}{n}\right) \\ &+ (O_{\bar{p}}^2 + 2O_{\bar{p}}Z_{\bar{p}} + Z_{\bar{p}}^2)T(\bar{p}) \\ &+ 2L_{\bar{p}}O_{\bar{p}}T\left(\bar{p} + \frac{\Delta^+ - \beta/n}{n}\right) \\ &+ 2R_{\bar{p}}O_{\bar{p}}T\left(\bar{p} + \frac{\Delta^+}{n}\right) \\ &+ 2L_{\bar{p}}R_{\bar{p}}T\left(\bar{p} + \frac{\Delta^+ + \Delta^-}{n}\right) \\ &+ R_{\bar{p}}^2 T\left(\bar{p} + \frac{2\Delta^+}{n}\right) \end{aligned}, & 0 < \bar{p} < 1 \\ 0, & \bar{p} \geq 1 \end{cases},$$

where $Z_{\bar{p}}$ represents the proportion of agents with p values of 0, $O_{\bar{p}}$ the proportion of agents with p values of 1, $R_{\bar{p}} =$

$\bar{p} - O_{\bar{p}}$ the probability that an agent with a p value other than 0 or 1 will choose to drive on the right side of the road, and $L_{\bar{p}} = (1 - Z_{\bar{p}} - \bar{p})$ the counterpart for the left side of the road; $Z_{\bar{p}} + O_{\bar{p}} + R_{\bar{p}} + L_{\bar{p}} = 1$.

We also modeled situations where the initial agent distribution is symmetrically distributed about $\bar{p}_0$, e.g., half of the population has initial bias of $\bar{p}_0 + \sigma$ and the other half has the initial bias of $\bar{p}_0 - \sigma$. When a population is distributed symmetrically about its mean, we note that $Z_{\bar{p}} + O_{\bar{p}}$ increases when $|\bar{p} - 0.5|$ is larger (given a constant σ). Additionally, $Z_1 = 0$ and $O_0 = 0$. Given these facts, we estimate $Z_{\bar{p}}$ and $O_{\bar{p}}$ using the cumulative distribution function Φ of the standard normal distribution:

$$\begin{aligned}Z_{\bar{p}} &\approx \Phi\left(\frac{0 - \bar{p}}{\max\{\sigma, \sigma + 2(\frac{1}{3} - \sigma)|\bar{p} - 0.5|\}}\right), \\ O_{\bar{p}} &\approx 1 - \Phi\left(\frac{1 - \bar{p}}{\max\{\sigma, \sigma + 2(\frac{1}{3} - \sigma)|\bar{p} - 0.5|\}}\right).\end{aligned}$$

The standard deviation of distributions with $\bar{p}$ values near zero and one is assumed to approach $\sigma + 2(\frac{1}{3} - \sigma)|\bar{p} - 0.5|$ (where σ is the initial population standard deviation). This is a somewhat arbitrary choice from many viable candidates; almost any scaling factor that would yield $Z_1 \approx 0$ and $O_0 \approx 0$ would have been appropriate.

4 Empirical Validation of Model Predictions

A mathematical analysis provides conceptual value only when it accurately models the underlying process in question. We now examine how well our analysis of the two-agent interaction problem corresponds to simulated interactions of agents a simulated population.




Figure 1: Constant Update Average Convergence Value.

Figure 1 presents a comparison of our analytical computations and empirical observations of the average value to which a population will converge as a function of the initial mean bias. In these simulations, all agents were initialized with bias values identical to the mean, and the constant update amount is $x = 0.01$. Each data point represents the average over 100,000 simulations or a prediction from the model.

Note that the mean always converges to 0 or 1. Inspection of the graph suggests that populations tend to converge to 0 when they have an initial $\bar{p}$ value less than 0.5, and to 1 when their initial bias is greater than 0.5, regardless of the other characteristics of the initial distribution. It is heartening to note that the theoretical predictions from our model almost perfectly match the empirical observations.




Figure 2: Constant Update Average Convergence Time.

Figure 2 displays the average number of time steps required for convergence by populations with the given initial bias. Both the simulations and the model indicate that populations with initial means near 0.5 take the longest to converge. The model predictions for convergence times, though not as accurate as those for convergence values, are still effective approximations of simulation results.




Figure 3: Proportional Update Average Convergence Value, $\beta = 0.6$.

Figures 3 and 4 compare the results for the model and simulation for the proportional update rule. In these figures we have zoomed in on the interesting transition region of initial bias in the range $[0.49, 0.51]$. For populations with initial means below 0.49, the average convergence value is approximately 0, and the average convergence time increases linearly with the initial population mean. Symmetrical results hold

for populations with initial mean biases above 0.51, with the population always converging to 1. Like the simulations and model using the constant update rule, the maximum average convergence time is observed when the initial mean bias is 0.5. It is interesting to note that the range of initial $\bar{p}$ values for which a population is observed to converge both to 0 and to 1 in at least about 1% of simulations is $[0.4, 0.6]$ for the constant update rule, but only $[0.49, 0.51]$ for the proportional update rule. The fact that the model has accurately predicted the decrease in size of this interval provides further validation of the model's fidelity.




Figure 4: Proportional Update Average Convergence Time, $\beta = 0.6$.

Next, we present results where initial population biases are symmetrically distributed about the initial population mean. Increasing the standard deviation of the initial distribution of population p values does not significantly impact the average final convergence value. It can, however, result in large increases in average convergence time for both of the update rules discussed. Figure 5 shows the average convergence time for an initial population of $n = 100$ agents, 50 having a initial p values of $0.5 - \sigma$, and the other 50 starting with $p = 0.5 + \sigma$ when the proportional update rule is used.

4.1 Proof of Convergence

The models discussed above accurately predict the frequency with which the populations converge to 0 and 1, and the time that it takes to get there. However, we want to ensure that the model does not overlook any other possible convergence values. We now present a proof that the population bias always converges to 0 or 1, which means that the model does consider all possible cases.

A *distribution* D represents a set of bias values for all members of a population. We call a set of distributions, $C \in \mathcal{P}([0, 1]^n)$, a “convergence set” iff $C \neq \emptyset$ and $\forall D \in C, S(D) \subseteq C$, where $S(D)$ is the set of all distributions that can possibly be reached at time $t + 1$ if a distribution D was present at time t . C is a “minimal” convergence set iff $\forall C' \subset C, C'$ is not a convergence set. Let D_0 be the distribution in which all agents have p values of 0, and D_1 be the corresponding distribution for p values of 1. The goal is




Figure 5: Effects of Initial Standard Deviation on Average Convergence Time, Proportional Update Rule.

to prove that $\{D_0\}$ and $\{D_1\}$ are the only minimal convergence sets. During this discussion, a population of n agents will be represented by the multiset $\biguplus_i^n \{p_i\}$, where p_i is the i th agent's bias value.

Proposition 1: $\{D_0\}$ and $\{D_1\}$ are minimal convergence sets. *Proof:* For distribution D_0 , every agent has a p value of 0; at each time step, both agents selected will always drive on the left side of the road and will always coordinate. Both agents will attempt to lower their p values, but since 0 is the lowest possible, the only configuration that can be reached is D_0 . Thus $S(D_0) = \{D_0\}$ and $\{D_0\}$ is a convergence set. The only proper subset of $\{D_0\}$ is the empty set $\emptyset$, which, by definition, cannot be a convergence set. Hence, $\{D_0\}$ is a minimal convergence set. We can also show that $\{D_1\}$ is a minimal convergence using similar logic. No proper super set of $\{D_0\}$ or $\{D_1\}$ can also be a minimal convergence set, by definition. Thus, no minimal convergence sets contain D_0 or D_1 except $\{D_0\}$ and $\{D_1\}$.

Proposition 2: There are no minimal convergence sets other than $\{D_0\}$ and $\{D_1\}$.

Proof: Suppose C is a convergence set such that $D_0, D_1 \notin C$. Since D_0 is the only distribution in C that contains all zero elements, every distribution in C must contain at least one non-zero p value. Similarly, from the absence of D_1 we can conclude that there is at least one non-one p value in each distribution.

Let $C_{low} = \{D \in C : (\forall D' \in C)(\bar{p}_D \leq \bar{p}_{D'})\}$. If C_{low} contains a distribution $D_{01,\bar{p}}$ composed of all zeroes and ones, it must contain at least one of each: $D_{01,\bar{p}} = \{0, 1\} \uplus D_{01,\text{rest}}$, where $D_{01,\text{rest}} \in \{0, 1\}^{n-2}$. Thus $D_{01,\bar{p}} - \{0, 1\} + \{\beta x, 1 - \beta x\}$ is in $S(D_{01,\bar{p}})$, C_{low} , and the next category of distributions. If C_{low} contains a distribution D with a non-one p value and another non-one, non-zero value, then $S(D) \ni D_{\bar{p}-\frac{x}{n}}$, a distribution with a lower $\bar{p}$ value than D that could result when the agents from D with these two p values are selected to interact with each other. However, C does not have any distributions with a lower mean than those in C_{low} , which contradicts the initial assumption that C was a convergence set. Thus, there are no minimal convergence sets except $\{D_0\}$ and $\{D_1\}$.

4.2 Further empirical results

The following trends were observed from simulation runs, but were not approximated by our model predictions.

Consider a population in which the majority of agents usually drive on the left side of the road while a minority always drives on the right. This type of skewed distribution exhibits different behavior than other distributions with the same mean. When the initial population mean is not close to 0.5, the skewness of the distribution does not significantly impact convergence statistics. However, when the minority subgroup has a strong bias towards one action while the majority subgroup has a weaker bias towards the other action such that the mean of the entire population is near 0.5, the skewness produces significant effects on both the average convergence value and time. Increasing the standard deviation of the initial distribution will magnify these effects. Figures 6 and 7 show the convergence statistics for an initial population of 100 agents, m of them with initial p values of $0.5 + 0.5 \frac{100-m}{100}$, and the other $100 - m$ starting at $0.5 - 0.5 \frac{m}{100}$. These results can be approximated by our model using $\bar{P}$ and T , but only if detailed information about the shape of the distribution at various times can be provided. For initially symmetric populations, we were able to approximate the distribution at any time with a normal distribution, and achieved good results. However, no such simple estimation can be made for skewed distributions.




Figure 6: Effects of Initial Skewness on Average Convergence Value.

When the constant update rule is used, $\bar{p}$ will generally follow an exponential relationship with time. However, when the population nears convergence, the rate of change of $\bar{p}$ decreases. This can be attributed to the increased frequency of events in which agents must adjust the changes in their p values so that they still lie in the interval $[0, 1]$.

When the proportional update rule is used, the average p value is characterized by short intervals of rapid change alternating with long intervals with only minimal oscillations in $\bar{p}$ values. During these long intervals, when $\bar{p}$ is almost constant, the individual agent p values are polarized to either endpoint of their domain, $[0, 1]$, and, like the average, only exhibit small oscillations. This phenomenon can be explained




Figure 7: Effects of Initial Skewness on Average Convergence Time.




Figure 8: A Typical Simulation Run Using the Proportional Update Rule.

by the fact that, if an agent's p value is near 0, several increases in its p value can be offset by only a small number of decreases, due to the proportionate nature of $\Delta_i(\text{act}_i)$. Figure 8 shows the average population bias over time for a typical simulation using the proportional update when all agents are initialized with biases of 0.5. Our model is unable to predict when these sharp drops in p values will occur.

5 Conclusions and Future Work

We developed a formal model to predict the convergent norms and the time taken for such norms to emerge in a population of agents adapting their biases for one of two choices based on success or failure of coordination in bilateral interactions. This analysis covered several variations of the initial state of the society and the bias update rule used by the agents. Our model predictions for both constant and proportional bias update schemes were validated by empirical results from a large number of simulations. In addition, we proved that a population using one of these two update rules will almost surely (that is, with 100% probability) converge on one of a small set of norms.

The analysis presented in this paper establishes a broad foundation for several types of subsequent work. While the matrices used to predict convergence statistics were nearly diagonal and thus could be scaled to model much larger populations, it would be ideal if closed-form expressions could be produced to express convergence statistics in terms of initial population characteristics.

We would also like to develop a better understanding of how increasing diversity in the population impacts convergence time. In a similar vein, skewness in the population could have unanticipated effects on convergence and we would prefer to gain insights into those effects through a more expansive analysis. We would also like to extend our analyses to other norm-emergence scenarios. Apart from different interaction schemes, we would also explore larger action sets, more complicated l -way interactions, and different network topologies.

References

- [Coleman, 1987] James S. Coleman. Norms as social capital. In G. Radnitzky and P. Bernholz, editors, *Economic Imperialism: The Economic Approach Applied Outside the Field of Economics*. Paragon House, New York, NY, 1987.
- [Delgado et al., 2003] Jordi Delgado, Josep M. Pujol, and Ramon Sanguesa. Emergence of coordination in scale-free networks. *Web Intelligence and Agent Systems*, 1:131–138, 2003.
- [Epstein, 2001] Joshua M. Epstein. Learning to be thoughtless: Social norms and individual computation. *Computational Economics*, 18:9–24, 2001.
- [Hume, 1978] Hume. *A Treatise of Human Nature*. Oxford University Press, 1978.
- [Kittock, 1994] James E. Kittock. The impact of locality and authority on emergent conventions: Initial observations. In *AAAI*, pages 420–425, 1994.
- [Lewis, 1969] David Lewis. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, MA, 1969.
- [Posch, 1995] Martin Posch. Evolution of equilibria in the long run: A general theory and applications. *Journal of Economic Theory*, 65:383–415, 1995.
- [Sen and Airiau, 2007] Sandip Sen and Stephane Airiau. Emergence of norms through social learning. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pages 1507–1512, 2007.
- [Shoham and Tennenholtz, 1997] Yoav Shoham and Moshet Tennenholtz. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, 94:139–166, 1997.
- [van Dyke Parunak, 2002] H. van Dyke Parunak. A mathematical analysis of collective cognitive convergence. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 248–255, New York: NY, 2002. ACM Press.
- [Young, 1996] Peyton H. Young. The economics of convention. *The Journal of Economic Perspectives*, 10(2):105–122, Spring 1996.