Simple models of distributed coordination

Frédéric Kaplan
Sony CSL - Paris - 6 Rue Amyot, 75005 Paris
E-mail: kaplan@csl.sony.fr

Abstract

Distributed coordination is the result of dynamical processes enabling independent agents to coordinate their actions without the need of a central coordinator. In the past years, several computational models have illustrated the role played by such dynamics for self-organizing communication systems. In particular, it has been shown that agents could bootstrap shared convention systems based on simple local adaptation rules. Such models have played a pivotal role for our understanding of emergent language processes. However, only few formal or theoretical results were published about such systems. This article discusses deliberately simple computational models in order to make progress in understanding the underlying dynamics responsible for distributed coordination and the scaling laws of such systems. In particular, the article focuses on explaining the convergence speed of those models, a largely underinvestigated issue. Conjectures obtained through empirical and qualitative studies of these simple models are compared with results of more complex simulations and discussed in relation with theoretical models formalized using Markov chains, game theory and Polya processes.

1 Introduction

> "Suppose you and I are rowing a boat together. If we row in rhythm, the boat goes smoothly forward; otherwise the boat goes slowly and erratically, we waste effort, and we risk hitting things. We are always choosing whether to row faster or slower; it matters little to either of us at what rate we row, provided we row in rhythm. So each is constantly adjusting his rate to match the rate he expects the other to maintain"
>
> David Lewis, Convention [24]

Linguistic dynamics involve many instances of coordination problems like agreeing on sound repertoires, word-meaning mappings or on the use of particular grammar constructions. In the 60s, Lewis made important steps in clarifying the processes underlying conventional aspects of language and meaning, suggesting to rephrase them in a game theoretical framework [24]. Understanding the role played by coordination dynamics in the context of language formation and evolution has been a crucial issue since then. In the mid-90s first models of self-organizing lexicons (e.g. [15, 33]) showed that agents could collectively agree on a shared mapping between words and meaning provided that they follow some well-chosen production and adaptation rules. Building on these pioneering approaches, self-organized communication systems have been successfully bootstrapped in increasingly complex systems including phonological simulations [9, 29] and population of autonomous embodied agents [36, 38]. However, despite an increased interest for these kinds of processes and a large amount of empirical studies, only few formal approaches or theoretical results have been published about such systems so far.

The sparseness of theoretical results about coordination dynamics for communication systems is probably related to the complexity of the models studied so far. Simple simulations of self-organization lexicons are for instance often already too complex to be studied formally (one interesting exception is [10]). Other computational approaches to language modeling can be considered to have been more successful in that respect (see [6, 23] for general overviews of the field). *Generational models* have led to interesting formal investigations (e.g. [32]). They are based on the simplifying assumption that language transmission is a unilateral process that goes from one generation to the next (with no generation overlap). In a similar manner, models based on *evolutionary algorithms* have also been studied in a relatively well defined framework. Their dynamics rely on a fitness criteria stating that agents that communicate best have a higher survival chance, leaving more offspring that can learn the language of their parents (e.g. [7, 26]). From another perspective, progress have also been made for issues related to *Zipf's power law* and *the least effort principle* (e.g. [13, 39]). Unfortunately, these different approaches do not address directly the central issues of coordination dynamics.

We will call *distributed coordination* the result of dynamical processes enabling independent agents to coordinate their actions without the need of a central coordinator. During such processes, the behavior of each agent is only the result of the history of its interaction. In particular, agents have no direct access to global properties of the population. Nevertheless, coordination arises as a result of collective dynamics depending on the adaptation rules used by the agents, in a distributed self-organized manner.

Distributed coordination in itself is not specific to emergent communication systems. The study of these dynamics is central to many disciplines like economy, physics, chemistry, ethology or sociology. This is particularly true for systems with self-reinforcing dynamics like auto-catalytic reactions, spin-glass systems, competition of norms, stigmergetic effects in ant colonies, opinion dynamics, etc. Successful theoretical approaches of such systems are usually based on abstract simplified models. Results obtained in these simple contexts can then be empirically extended to describe more complex instances of the problems studied. Despite apparent similarities between problems considered in the various disciplines, great care must be taken before transferring results from one context to another. Assumptions underlying each model are often specific to the field considered and may reveal not to be relevant anymore for another discipline. Models may generally deal with the same processes, but differ in the details of dynamics.

In this article, we discuss simple models of distributed coordination. Our objective is to progress in understanding (1) the dynamics underlying distributed coordination in the context of emergent communication systems and (2) the scaling laws of such systems regarding the number of agents involved in the coordination. We deliberately study models much simpler than most systems traditionally considered in this field. We believe that progress in understanding the formal properties of self-organizing lexicons will be difficult without a finer characterization of the dynamics involved in simpler situations of competitions between conventions. The next section presents an empirical study of three related models, and focuses on explaining the convergence times of those models. Each model illustrates a particular dynamics of distributed coordination. Experimental results show that only the first two lead to actual convergence towards the use of a unique convention. The first one ensures a slow convergence, whereas the second one permits to reach high coherence in a faster way. This study suggests that fast convergence is in $N \cdot log(N)$ (where N is the number of agents). A qualitative interpretation of this dependency is provided. Section 3 discusses various theoretical frameworks for interpreting the empirical findings of section 2, including Markov chains, models based on stochastic games and Polya processes. Finally, section 4 studies a classic model of lexicon self-organization, showing that the conjectures about convergence times resulting from simple models can scale to more complex ones.

2 Three simple models

Let us consider a population of N agents where each agent can choose a particular conventional name among a convention set $\mathcal{C} = \{c_1, c_2, .., c_{\|\mathcal{C}\|}\}$, where $\|\mathcal{C}\|$ is the cardinal of $\mathcal{C}$. In this section we will restrict ourselves to the particular case of a set containing only two elements $\mathcal{C} = \{c_1, c_2\}$. Each agent a is characterized by a preference vector V_a, which components are different depending on the models. The preference vector V_a of an agent cannot be inspected by another agent. At each time step, two agents are randomly chosen. Agent a_1 produces a convention c_k according to a production rule $\mathcal{P}(V_{a_1}) = c_k$ and agent a_2 updates its vector V_{a_2} with an update rule $\mathcal{U}$. Let $N_1(t)$ be the number of agents producing convention c_1 and $N_2(t) = N(t) - N_1(t)$ the number of agents producing convention c_2. We can define the *coherence level* at time t as:

$$CL(t) = \frac{max(N_1(t), N_2(t))}{N} \tag{1}$$

Coordination is said to be *complete* when $CL = 1$. This means that all the agents of the population have converged to a consensus.

This section discusses successively three simple models, an imitation-based model (Model A) and two frequency-based models (Model B and C). They are representative of many more complex ones studied in the field. Each model is defined as a couple of production and update rules $(\mathcal{P}, \mathcal{U})$. The rules used are always based on local interaction and are function of the agent's personal history. They can be interpreted intuitively as different strategies of production and interpretation during interaction between agents. In model A, the speaker simply produces the convention he heard last as a listener. In model B, The speaker produces the convention that he has heard most frequently as a listener. In model C, The speaker produces a convention with a probability proportional to the frequency that he has heard as a listener. These intuitive interpretations are summarized in table 1. However, it should be noted that given the simplicity of the models, other types of interpretations can be considered.

Table 1: Intuitive interpretation of the three models

Model	Intuitive interpretation in terms of communication interaction
Model A	Imitation-based model: The speaker simply produces the convention he heard last as a listener
Model B	Frequency-based model: The speaker produces the convention that he has heard most frequently as a listener
Model C	Frequency-based model: The speaker produces a convention with a probability proportional to the frequency that he has heard as a listener

2.1 Imitation-based model A

Model A. In this first model, V_a can only have two values: V_1 and V_2. Agent a_1 produces convention c using the following $\mathcal{P}_A$ rule:

$$\mathcal{P}_A(V_{a_1}) = c_k = \left\{ \begin{array}{l} c1 \text{ if } V_{a_1} = V_1 \\ c_2 \text{ if } V_{a_1} = V_2 \end{array} \right\} \tag{2}$$

Agent a_2 updates its vector by adopting immediately the convention use of a_1, using the following rule:

$$\mathcal{U}_A : \left\{ \begin{array}{l} V_{a_2} = V_1 \text{ if } c_k = c_1 \\ V_{a_2} = V_2 \text{ if } c_k = c_2 \end{array} \right\} \tag{3}$$

Starting with $N_1(0) = \frac{N}{2}$ (agents with $V_a = V_1$) and $N_2(0) = \frac{N}{2}$ (agents with $V_a = V_2$), what kind of evolution will be observed?

Exp A.a ($N = 100$, $N_1(0) = \frac{N}{2}$, $N_2(0) = \frac{N}{2}$, End criteria : CL=1, 4 runs) Fig 1 shows four sample evolutions for 100 agents. The population eventually converges to a state of complete coordination ($CL = 1$). However, convergence happens only after a long series of oscillations.

Exp A.b ($N = 100$, $N_1(0) = \frac{N}{8}$, $N_2(0) = \frac{7N}{8}$, End criteria : CL=1, 4 runs) Fig 2 shows four sample evolutions for 100 agents for a different initial configuration. In all the cases, the population eventually converges to a state of complete coordination ($CL = 1$), but not necessarily towards to convention initially preferred.

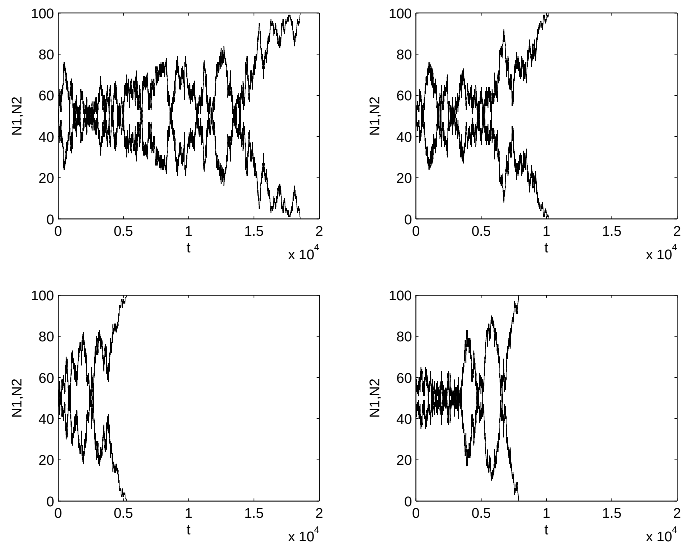

Figure 1: Competition between two conventions c_1 and c_2 in a population of 100 agents. Initially, 50 agents choose c_1 and 50 other agents choose c_2. Several oscillations are observed before convergence (Exp A.a.)

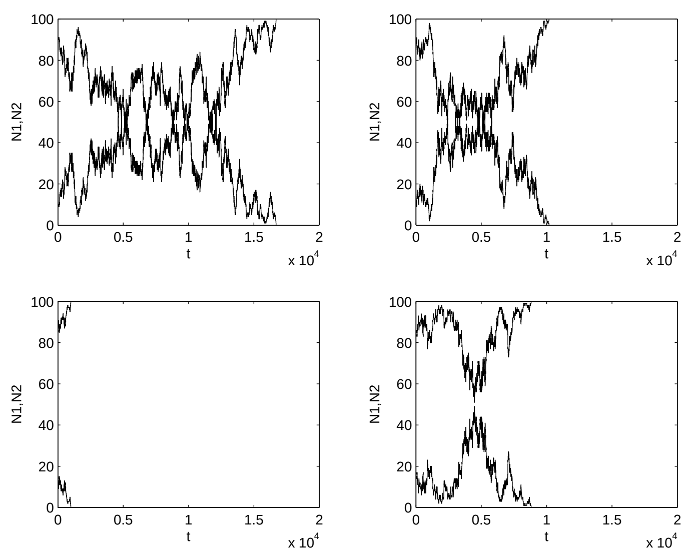

Figure 2: Competition between two conventions c_1 and c_2 in a population of 100 agents in a biased initial configuration. The population eventually converges to a state of complete coordination, but not necessarily towards to convention initially preferred (Exp A.b.)

The dynamics associated with this model A can be better understood if we consider the different probabilities of evolution at time t.

- Probability to choose an agent using convention c_1: $p_1(t) = \frac{N_1(t)}{N}$
- Probability to choose an agent using convention c_2: $p_2(t) = \frac{N_2(t)}{N}$
- Probability that an agent using $c1$ is chosen as agent 1, and an agent using $c2$ is chosen as agent 2 (and therefore adopts convention $c1$): $p_1(t) \cdot p_2(t)$
- Probability that an agent using $c2$ is chosen as agent 1, and an agent using $c1$ is chosen as agent 2 (and therefore adopts convention $c2$): $p_2(t) \cdot p_1(t)$
- Probability that an agent interacts with an agent using the same convention $p_1^2(t) + p_2^2(t)$

With this model, at any time t it is *equally probable* that $N_1(t)$ or $N_2(t)$ increases. This means that no dynamics drive the population towards coordination. However, after some time, convergence occurs and the population ends up in using only $c1$ or $c2$. How is this possible? This situation is similar to a *random walk* or *brownian movement*. A random walk corresponds to the path of someone that would choose randomly at each step whether to go forward or backward. Such a walker would on the average oscillate around its starting position but from time to time it would get away from it. During a random walk, the quadratic average distance of the walker is $\sigma = \sqrt{n_{step}}$ where n_{step} is the number of steps taken by the walker. This means that as the walker takes more steps, the probability of being far from the center increases (figure 3). Suppose that we want to be sure at 99% that the walker has at least been once at a certain distance d from the starting position. This should be true if σ is sufficiently big compared to d (in a ratio that remains to be defined). To get the same certainty for a distance $4 \cdot d$, we would have to wait 16 times longer.

One difference between the dynamics of model A and the ones of a random walk is that the probability of evolution in model A is a factor of p_1 and p_2 (whereas it is fixed in a classic random walk). The expression $p_1^2 + p_2^2$ reaches its minimum $1/2$ for $p_1 = p_2 = 1/2$. This means that $N_1(t)$ and $N_2(t)$ change more rapidly when $N_1(t)$ is close to $N_2(t)$ than when they are more different (Figure 4).

Despite this difference, can we make hypotheses about the scaling law of model A based on its analogy with a random walk? To enter in a state of complete coordination, the random walk must reach distance $d = \frac{N}{2}$ (converting the other half of the population). This means that convergence time T_c should increase in N^2. The following experiments permit to verify this conjecture for model A.

Exp A.c (different N, $N_1(0) = \frac{N}{2}$, $N_2(0) = \frac{N}{2}$, End criteria : $CL = 1$) Figure 5 shows a log-log plot of simulation results for various population sizes N. Each point corresponds to the number of time steps necessary to reach complete coordination. The slope of the curve obtained by linear regression is 2,02. This is an experimental verification of the expected quadratic dependency.

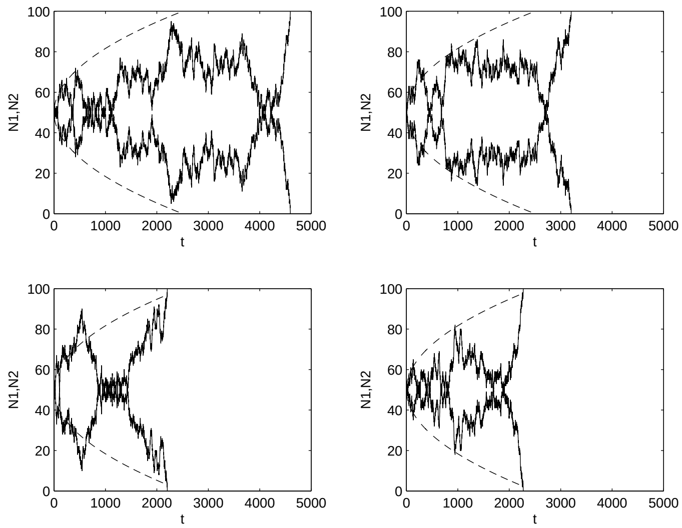

Figure 3: Sample evolution for four random walks and associated values of the theoretical average distance $\sigma = \sqrt{n_{step}}$ in the same initial condition than for experiment A.a.

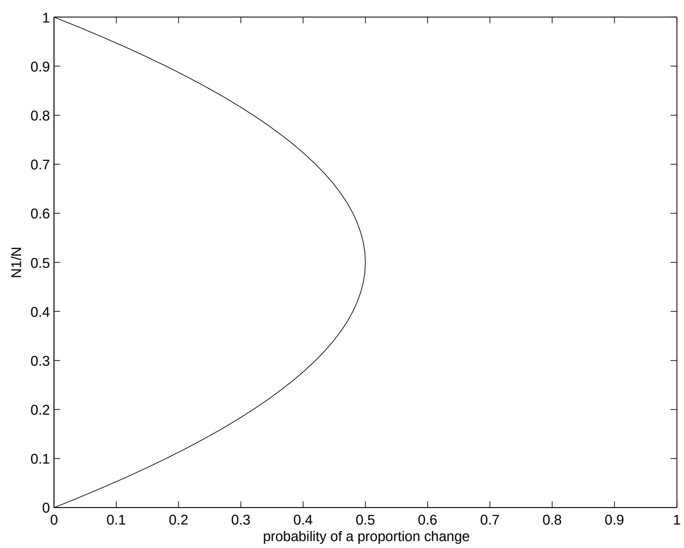

Figure 4: Probability of a proportion change in the population. This means that, in model A, $N_1(t)$ and $N_2(t)$ change more rapidly when $N_1(t)$ is close to $N_2(t)$ than when they are more different.

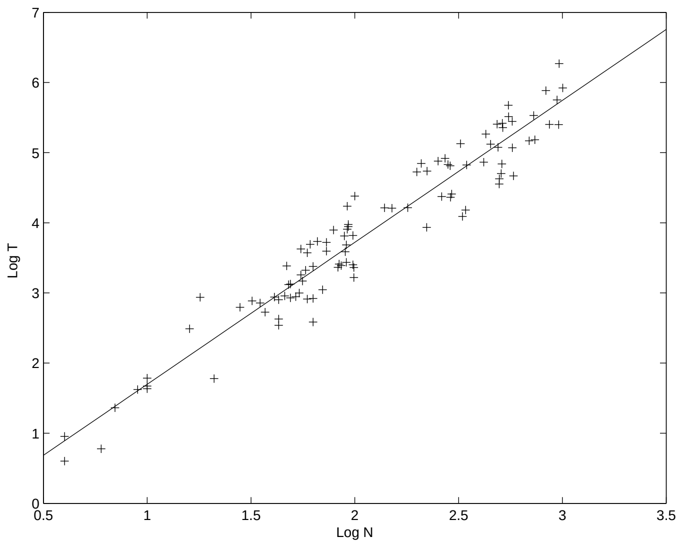

Figure 5: Log-log diagram comparing time of convergence T_c for different population sizes N. The slope obtained by linear regression is 2,02. This suggests a quadratic dependency (Exp A.c).

2.2 Frequency-based model B

Model B. In this model, each agent a is characterized by a preference vector V_a of size 2 where each convention c_i of $\mathcal{C}$ is associated with a score $v_{a,i}$.

$$V_a = \left\{ \begin{array}{c} v_{a,1} \\ v_{a,2} \end{array} \right\} \tag{4}$$

Agent a_1 produces convention c_k using the following $\mathcal{P}_B$ rule:

$$\mathcal{P}_B(V_{a_1}) = c_k = c_{argmax_i(v_{a_1,i})} = \left\{ \begin{array}{c} c_1 \text{ if } v_{a_1,1} > v_{a_1,2} \\ c_2 \text{ if } v_{a_1,2} > v_{a_1,1} \\ \text{random if } v_{a_1,1} = v_{a_1,2} \end{array} \right\} \tag{5}$$

Agent a_2 updates its vector by increasing the score associated with the convention c_k:

$$\mathcal{U}_B : \left\{ \begin{array}{c} v_{a_2,1} \leftarrow v_{a_2,1} + \delta \text{ if } c_k = c_1 \\ v_{a_2,2} \leftarrow v_{a_2,2} + \delta \text{ if } c_k = c_2 \end{array} \right\} \tag{6}$$

At the beginning of the experiments, $\frac{N}{2}$ are initialized with $(\delta, 0)$ and the other half with $(0, \delta)$.

Exp B.a ($N = 100$, $N_1(0) = \frac{N}{2}$ et $N_2(0) = \frac{N}{2}$, End Criteria: CL=1, 4 runs) Four sample evolutions for 100 agents is presented on figure 6. The oscillations observed with model A are much smaller. As soon as a convention spreads more in the population than the other, its domination seems to amplify even more over time.

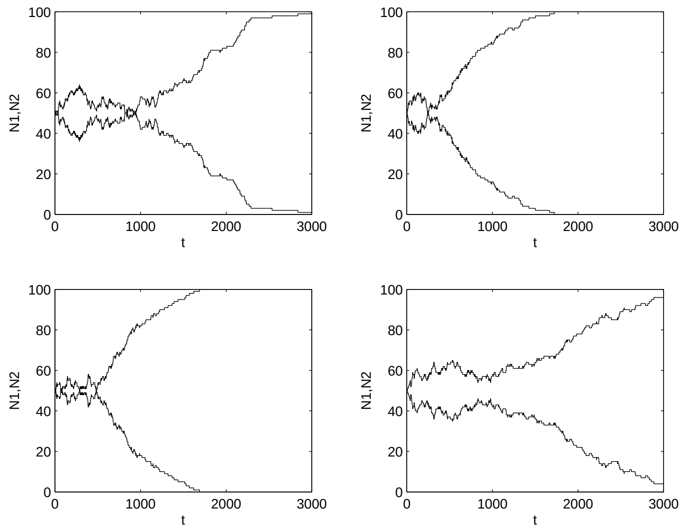

Figure 6: Competition between two conventions c_1 and c_2 in a population of 100 agents. Initially, 50 agents choose c_1 and 50 other agents choose c_2. Dominance of one convention tends to increase over time (Exp B.a)

There is a crucial difference between model B and model A. In model B interactions between agents already producing the same convention c_k strengthen the tendency to produce c_k in the future. In model A, such kind of interactions had no effects. This self-reinforcing dynamics result in a *positive feedback loop*: as soon as a convention starts to spread more than the other in the population, the probability that it wins the competition increases. The update rule $\mathcal{U}_B$ performs a form of statistical induction about the diffusion of the each convention in the population. With this interpretation, production rule $\mathcal{P}_B$ consists in choosing the most diffused convention from the point of view of the agent.

Exp B.b (different values of N, $N_1(0) = \frac{N}{2}$, $N_2(0) = \frac{N}{2}$, End Criteria: CL = 1) Figure 7 presents a log-log diagram of the time of convergence T_c for different population size N. The slope of the linear regression is 1.30. As expected, convergence is much faster than for model A. The value 1.30 being close to 1, we can test a $N \cdot log(N)$ law. Figure 8 plots the average convergence time divided by the population size on a logarithmic scale. The number of steps necessary to reach complete consensus ($CL = 1$) and partial consensus ($CL = 0.8$) are represented. Although the data is dispersed, a linear fit is possible suggesting a $N \cdot log(N)$ law.

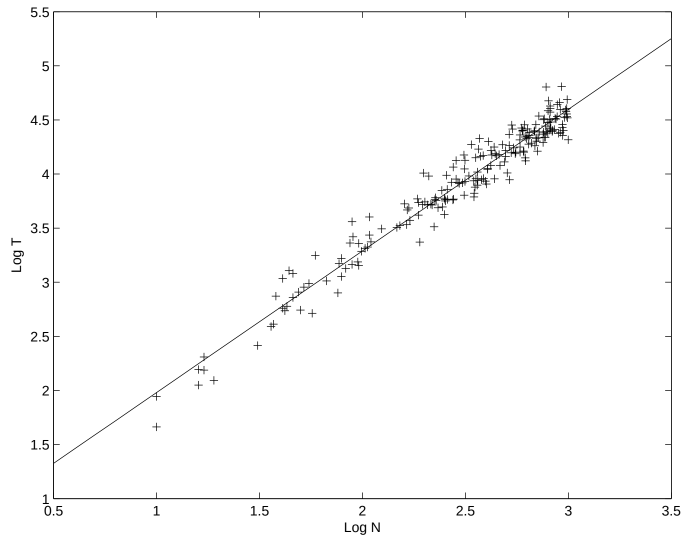

Figure 7: Log-log diagram comparing time of convergence T_c for different population sizes N. The slope obtained by linear regression is 1,30 (Exp B.b)

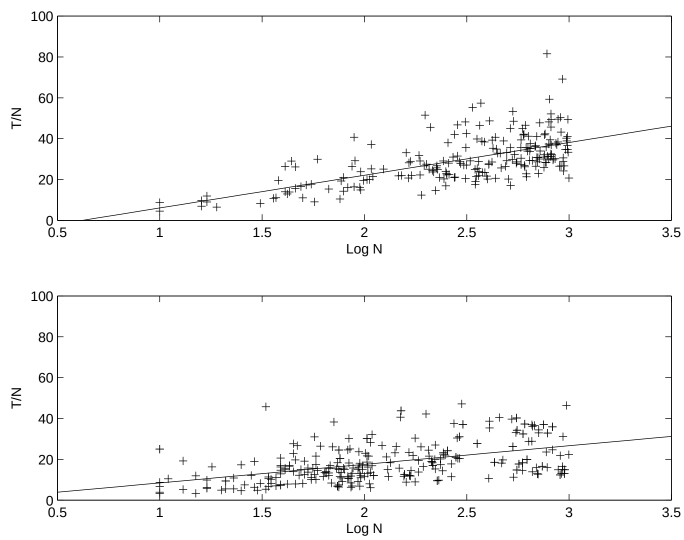

Figure 8: Ratio between convergence time T_c and population size N for different population sizes plotted on a logarithmic x-axis. Cases of partial and complete consensus are considered. Although the data is dispersed, a linear fit is possible suggesting a $N \cdot log(N)$ law. Slopes obtained by linear regression are 16.0 (complete consensus) and 9.1 (partial consensus) (Exp B.b).

We will now present a qualitative reasoning in order to interpret the $N \cdot log(N)$ convergence empirically observed with model B. Consider a population of size N, where $N_1(t)$ and $N_2(t)$ are respectively the number of agents using c_1 or c_2 after t iterations. We will assume that during the first N iterations, the positive feedback loop does not yet have an important effect and that the system is comparable to a random walk. At iteration $t = N$, given that the agents are randomly picked, the number of agents using convention c_1 and c_2 should have slightly changed so that, for instance, $N_1(t)$ is a bit more important than $N_2(t)$. Let's define ϵ so that:

$$\frac{N_1(N)}{N_2(N)} = 1 + \epsilon \tag{7}$$

A typical value of ϵ is $\epsilon = \frac{\sigma}{N}$ where $\sigma = \sqrt{N}$ is the quadratic deviation of a random walk. As a consequence, $\epsilon = \frac{\sqrt{N}}{N} = \frac{1}{\sqrt{N}}$.

During the next cycle of N iterations, the evolution will not be a pure random walk anymore but biased towards convention c_1. The positive feedback loop starts to have an effect. After $2N$ iterations, on average, $1 + \epsilon$ more agents using c_1 have been selected.

$$\frac{N_1(2N)}{N_2(2N)} = (1 + \epsilon)\frac{N_1(N)}{N_2(N)} = (1 + \epsilon)^2 \tag{8}$$

After $3N$ iterations, on average, $(1 + \epsilon)^2$ more agents using c_1 have been picked.

$$\frac{N_1(3N)}{N_2(3N)} = (1 + \epsilon)^2\frac{N_1(2N)}{N_2(2N)} = (1 + \epsilon)^4 \tag{9}$$

Therefore, in general after the mN first interations

$$\frac{N_1(mN)}{N_2(mN)} = (1 + \epsilon)^{2^m} \tag{10}$$

Note that equation 10 is supposed to be valid only at the beginning of the evolution, but may not be true anymore at the end of the experiment, as the rate of increase of the ratio should slow down as fewer and fewer agents producing the least frequent convention are chosen during the random selection process.

A partial consensus is obtained when $\frac{N_1(mN)}{N_2(mN)} > A$, for A sufficiently big. So for $(1 + \epsilon)^{2^m} = A$. Using logarithms, this is equivalent to:

$$2^m \cdot log(1 + \epsilon) = logA \tag{11}$$

For N sufficiently big, $log(1 + \frac{1}{\sqrt{N}}) \approx \frac{1}{\sqrt{N}}$. As a consequence:

$$\frac{2^m}{\sqrt{N}} = logA = K \tag{12}$$

Taking the logarithm, this gives:

$$m = log_2(K.\sqrt{N}) = log_2(K) + \frac{1}{2}log_2(N) \propto log(logA)) + \frac{1}{2}log(N) \tag{13}$$

When N is sufficiently big, the first term can be neglected. For instance to reach a 90% consensus ($CL = 0.9$), $A = 9$ and $log(log9) = -0,022$. For $N = 100$, $log(N)$ is a hundred times bigger. This means that if N and A are sufficiently big m is proportional to $log(N)$:

$$m \propto log(N) \tag{14}$$

The most important part of the convergence is achieved in $N.m$ iteration, so for $T_c(A)$the number of iteration necessary to reach a partial convergence defined by A:

$$t_c(A) \propto N.LogN \tag{15}$$

We have experimentally observed (slopes of figure 8) that the ratio between the time to reach a partial convergence at 80% and a complete convergence at 100% stays constant for the different population sizes we considered. Our result can therefore be extrapolated to the case of complete convergence;

$$T_c \propto N.LogN \tag{16}$$

2.3 Frequency-based model C

Model C. This model is closely similar to model B apart from the production rule $\mathcal{P}_C$, which corresponds now to a probabilistic choice. The probability of choosing c_k is proportional to the relative score of this convention compared to the other.

$$\mathcal{P}_C(V_{a_1}) = c_k : \left\{ \begin{array}{c} P(c_1) = \frac{v_{a_1,1}}{v_{a_1,1}+v_{a_1,2}} \\ P(c_2) = 1 - P(c_1) = \frac{v_{a_1,2}}{v_{a_1,1}+v_{a_1,2}} \end{array} \right\} \tag{17}$$

Agent a_2 updates its vector following rule $\mathcal{U}_B$. Changing the production rule from a greedy winner-take-all strategy to a probabilistic one has an important effect on the dynamics. We can draw from the following experimental results that complete coordination cannot be obtained with such a production rule.

Exp C.a ($N = 100$, $N_1(0) = \frac{N}{2}$ et $N_2(0) = \frac{N}{2}$, End criteria: $T = 600$, 4 runs) This experiment starts with the same initial conditions than the one considered for model B: $\frac{N}{2}$ agents are initialized with $(\delta, 0)$ and the other half with $(0, \delta)$. Figure 9 presents four sample evolutions. After an initial drift, dynamics tend to maintain the distribution of c_1 and c_2 over time. The production rule $\mathcal{P}_C$ reinforces the relative distribution of the two conventions as they are induced using the update rule. The system is stationary.

2.4 Conjectures

Two conjectures can be drawn based on the experiments conducted in this section with simple models.

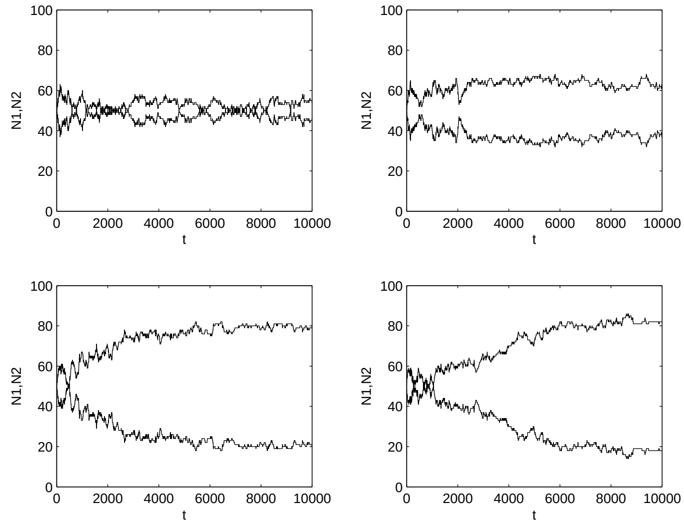

Figure 9: Competition between two conventions c_1 and c_2 in a population of 100 agents. Initially, 50 agents have a bias toward c_1 and 50 other a bias toward c_2. After an initial drift period, the distribution tends to be maintained (Exp C.a).

- Conjecture 1: Among the three models studied, only model B (self-reinforcing dynamics) permits a fast coordination of the entire population towards the use of a single convention. Model A is approximatively similar to a random walk, converging in quadratic time. On the contrary, dynamics of model C tends to maintain the distribution of the convention at a fixed level.

- Conjecture 2: Experimental results and qualitative interpretations suggest that self-reinforcing dynamics of model B converge in $N \cdot log(N)$, where N is the population size.

These results are summarized in table 2. In section 3 we will discuss several theoretical framework to interpret conjectures 1 and 2. In section 4 we present results that corroborate the $N \cdot log(N)$ conjecture for more complex models.

Table 2: Conjectures based on empirical results with simple models

Model	Distributed coordination	Convergence time
Model A	convergence towards a single convention	N^2
Model B	convergence towards a single convention	$N \cdot log(N)$
Model C	stabilization of the current distribution	∞

3 Theoretical frameworks

Can the empirical results of the models studied in the previous section be studied from a more theoretical point of view? Phenomena related to distributed coordination have been studied in many disciplines under various frameworks ranging from mathematical economics to statistical physics. In various contexts, global coordination emerge out of a set of simple elements (e.g. particles, individuals, agents, cells) which undergo simple repetitive local changes. However, not all these framework are adapted to the interpretation of the models that interest us. For instance, in physics, Ising models(which can be considered as a particular case of Markov random fields [22]) are concerned with set of spins that can take binary states $-1, 1$ a situation that bears some resemblance with the models of competition described in the last section. Such kind of models have been used to study spontaneous magnetization of spins but have also been extended to more abstract cases involving the dynamics of consensus in quantitative sociology [40] and computational ecology [14]. However, as most of these models focus on the dynamics of particular statistics over the population rather than on the particular udpate and production rules used by the agent, results obtained in such frameworks cannot be easily adapted to our own. Other types of formal modelling are more promising. In this section we will review successively the relative advantage of formalism based on Markov chains, stochastic games and Polya processes to progress in the understanding of the dynamics of models A, B and C.

3.1 Interpretation of model A with Markov chains

Ke, Minett, Au and Wang have conducted interesting research concerning the use of a Markov chain formalism to study emergent communication systems [20]. The dynamics of model A can be studied in such a framework. Each state of the Markov chain corresponds to a particular proportion of agents using convention c_1. At any time t, their is a certain probability that the population changes to an adjacent state where the population of agents using convention c_1 would have either increased or decreased by one. In model A, this probability only depends on the current proportion of agents using the convention, thus respecting the Markov property:

$$Pr(X_{t+1} = k | X_0 = h, ..., X_t = j) = Pr(X_{t+1} = k | X_t = j) \tag{18}$$

Therefore, the dynamics can be captured using a single transition matrix $\mathbf{P}$ of size $(N+1) \cdot (N+1)$. Here is an example of such a matrix for $N = 6$:

$$\mathbf{P} = \left\{ \begin{array}{ccccccc} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ c(1) & d(1) & c(1) & 0 & 0 & 0 & 0 \\ 0 & c(2) & d(2) & c(2) & 0 & 0 & 0 \\ 0 & 0 & c(3) & d(3) & c(3) & 0 & 0 \\ 0 & 0 & 0 & c(4) & d(4) & c(4) & 0 \\ 0 & 0 & 0 & 0 & c(5) & d(5) & c(5) \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array} \right\} \tag{19}$$

For model A, $c(j)$ and $d(j)$ are defined as:

$$c(j) = c(N-j) = p_1 \cdot p_2 = \frac{j \cdot (N-j)}{N^2} \tag{20}$$

$$d(j) = d(N-j) = p_1^2 + p_2^2 = \frac{j^2 + (N-j)^2}{N^2} \tag{21}$$

So, for $N = 6$

$$\mathbf{P} = \left\{ \begin{array}{ccccccc} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{5}{36} & \frac{26}{36} & \frac{5}{36} & 0 & 0 & 0 & 0 \\ 0 & \frac{8}{36} & \frac{20}{36} & \frac{8}{36} & 0 & 0 & 0 \\ 0 & 0 & \frac{9}{36} & \frac{18}{36} & \frac{9}{36} & 0 & 0 \\ 0 & 0 & 0 & \frac{8}{36} & \frac{20}{36} & \frac{8}{36} & 0 \\ 0 & 0 & 0 & 0 & \frac{5}{36} & \frac{26}{36} & \frac{5}{36} \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array} \right\} \tag{22}$$

To study the convergence of such a system, the eigenvalues λ_i and corresponding left and right eigenvectors $\mathbf{x_i}$ and $\mathbf{y_i}$ of $\mathbf{P}$ must be found.

$$\mathbf{x_i^T P} = \lambda \cdot \mathbf{x_i^T} \tag{23}$$

and

$$\mathbf{P y_i} = \lambda \cdot \mathbf{y_i} \tag{24}$$

The objective is to identify a number of *closed* states, any subset C of states so that there is no arc from any of the states in C to any of the states not in C. The first and last states in our case are clear examples of states where no transition to any other state is allowed anymore. This implies that the multiplicity of the eigenvalue $\lambda = 1$ is 2. The two corresponding left eigenvectors $\mathbf{x_i}$ are straithforward to identify. For $\mathbf{y_i}$ a system of equations must be solved. An example of how to solve such a system is described in [20] for a similar case. This permits to prove the convergence of systems using production and update rules similar to the ones of model A. However this framework does not seem adapted to the study of model B and C. Other forms of modeling must therefore be considered.

3.2 Interpretation of model B in the framework of stochastic games

Shoham and Tennenholtz have convincingly argued that the framework of stochastic games, popular for economic simulations, is relevant for the study of the emergence of social conventions [30]. By studying more formally the coordination game introduced by Lewis [24], they show several important results about the dynamics of convention emergence. A typical coordination game involved two players and is characterized by a payoff matrix like the following

$$M = \left\{ \begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \right\} \tag{25}$$

This means that both players receive rewards only if they coordinate their action. The problem is therefore very similar to the one studied in section 2, if we consider a population of agents playing such a game and having to choose between two conventions c_1 or c_2. Such forms of coordination games are said to have two kinds of *Nash equilibria*: joint strategies that are stable in the sense that no single agent benefits from switching to another strategy if all others remain unchanged. In our case, each Nash equilibrium corresponds a situation in which a single convention c_1 or c_2 is used by the entire population.

Shoham and Tennenholtz demonstrate that a way to reach such a collective agreement is to use a reward system called the *highest cumulative reward rule*. According to this rule, an agent switches to a new action if and only if the total payoff obtained from that action in the latest m iterations is greater than the payoff obtained from the currently-chosen action in the same time period. This rule bears important similarity with the update and production rules of model B.

The authors prove not only that the highest cumulative reward rule guarantees eventual emergence of coordination but also study the number of iterations required to reach such a Nash equilibrium. They present a general lower bound on the efficiency of convention evolution. This lower bound is in $N \cdot log(N)$, where N is the population size.

These are important results giving qualitative support of our empirical finding of the previous section. However, the models we studied cannot be strictly assimilated with models based on reinforcement like the ones studied by this kind of stochastic game framework. In model B and C agents do not adapt after receiving a coordination reward. Adaptation takes place while agents are listeners, observing the convention produced by other agent. In the case of the competition between two conventions, this difference may not play an important role but results may differ greatly when considering agreement for a larger number of conventions. This difference invites us to consider another framework.

3.3 Interpretation of models B and C with Polya processes

Model B and C can be interpreted using the formalism of Polya's urn problem. Polya processes are simple to state, rigorously tractable and yet they lead to complex phenomena. They have been mainly applied to model path-dependent processes in economical clustering (e.g. [2, 3, 4]). They have also been used as models for formal learning [16] and neural modeling [21]. The relevance of this form of modeling for studying the emergence of shared conventions was initially argued by Ferrer and Sole [12].

Let us consider an infinite urn that can contain red balls and white balls. Polya processes correspond to situations where the probability of adding a red or white ball depends on the current proportion of these balls in the urn. The following formalism can be used to model such kind of path-dependent process in the general case of an urn that can contain K kinds of balls [2, 3, 4]. Suppose vector $X_t = (X_t^1, X_t^2, ..., X_t^K)$ describes the proportion of color type 1 to K after n iterations. For $n = 1$ the initial vector of the urn present in the urn is $b_1 = (b_1^1, b_1^2, ..., b_1^K)$. A new ball is added after each iteration. Let us define a sequence of continuous functions $\{q_n\}$ from the space of color proportion to the space of probabilities (to add at each iteration a ball of a particular kind). The probability at iteration t to add a ball of color i is $\{q_t^i(X_t)\}$. Let $w = \sum_{i=1}^{K} b_1^i$ be the initial number of balls in the urn. We can define at iteration t for $i = 1, ..., K$ the following random variable.

$$\beta_t^i(x) = \left\{ \begin{array}{c} 1 \text{ with a probability } q_t^i(x) \\ 0 \text{ with a probability } 1 - q_t^i(x) \end{array} \right\} \tag{26}$$

The number of balls of color i at the next iteration is described by:

$$b_{t+1}^i = b_t^i + \beta_t^i(X_t) \tag{27}$$

The total number of balls at time t is $(w + t - 1)$. As a consequence the proportion X_t^i is:

$$X_t^i = \frac{b_t^i}{w + t - 1} \tag{28}$$

Equation 27 can be rewritten:

$$X_{t+1}^i \cdot (w + t) = X_t^i \cdot (w + t - 1) + \beta_t^i(X_t) \tag{29}$$

$$X_{t+1}^i \cdot (w + t) = X_t^i \cdot (w + t) + \beta_t^i(X_t) - X_t^i \tag{30}$$

$$X_{t+1}^i = X_t^i + \frac{1}{w + t}[\beta_t^i(X_t) - X_t^i] \tag{31}$$

This last equation can be rewritten in the following way:

$$X_{t+1}^i = X_t^i + \underbrace{\frac{1}{w + t}[q_t^i(X_t) - X_t^i]}_{\text{governing part}} + \underbrace{\frac{1}{w + t}[\beta_t^i(X_t) - q_t^i(X_t)]}_{\text{perturbation}} \tag{32}$$

This equation captures the basic dynamics of such kind of systems. The governing part is responsible for the overall evolution of the system and it can be shown that:

$$E[\beta_t^i(X_t) - q_t^i(X_t)|X_t] = 0 \tag{33}$$

As a consequence:

$$E[X_{t+1}^i|X_t] = X_t^i + \frac{1}{w + t}[q_t^i(X_t) - X_t^i] \tag{34}$$

The two particular cases that we have studied in section 2 correspond to two urn functions $q_n^i(X_t)$ that are indepedent from n: *max* and *id* [12].

- Function *max* consists in systematically choosing one kind of ball if the corresponding proportion in the population is higher than the others ($max(X_t^i) = 1$) when X_t^i is the maximal value and 0 otherwise) . In case of more that one maximum values, one of them is chosen at random. This is similar with the greedy production rule $\mathcal{P}_B$.

- Function *id* corresponds to a probabilistic choice proportional to the current proportion of balls in the urn ($id(X_t^i) = X_t^i$). This is similar to the production rule $\mathcal{P}_C$.

The convergence of such a system towards a fixed distribution is formally demonstrated by Arthur [2, 3] in the case of the *id* function. Ferrer and Sole introduced the idea of using the max function to model situations involving positive reinforcement and showed that an extreme consensus is reached in such a situation [12]. With the max function, dynamics lead to the rapid domination of a single ball color over the other ones. With the id function, dynamics corresponds to a stabilization of the relative proportion of the different balls in the urn.

Another general formulation can be obtained if we consider $q_t^i(X_t) = (X_t)^\gamma$ Chung, Handjani and Jungreis demonstrate that the system converges towards the use of a single ball when $\gamma > 1$ (positive reinforcement), maintains existing proportion when $\gamma = 1$ and tends to equalize the different proportions when $\gamma < 1$ (negative reinforcement) [8].

Can we directly extend results obtained in the framework of Polya processes to the models B and C studied in the previous section? Polya processes are models of a system interacting with itself. In that sense, distributed systems like the ones studied in section 2 are not strictly speaking Polya processes. A heuristic argument for the equivalence with such systems was presented by Ferrer and Sole [12].

In their model of distributed Polya process, each urn corresponds to one agent in a population of N agents . At time t, the interaction between the agents is modeled using a boolean connectivity matrix Ψ_t^{ij}, with $\Psi_t^{ij} = 1$ if the i-th agent is connected to the j-th agent at time t and 0 otherwise. Ψ_t is symmetric and, to avoid self-reinforcement, $\Psi_t^{ii} = 0$. An additional constraint is that all agents are always connected to the same number of agents C. As a consequence, $\sum_{j=1}^{N} \Psi_t^{ij} = C$. For instance the following matrix is compatible with this constraints with $N = 4$ and $C = 1$.

$$\Psi_t = \left\{ \begin{array}{c} 0\,1\,0\,0 \\ 1\,0\,0\,0 \\ 0\,0\,0\,1 \\ 0\,0\,1\,0 \end{array} \right\} \tag{35}$$

A random matrix of this kind is generated at each step.

An additional index i is now needed for vectors X_t and b_t, as every agent has its own urn. At time t, the proportion and the number of balls of type $1...K$ are now respectively $X_t^i = (X_t^{i1}, X_t^{i2}, ..., X_t^{iK})$ and $b_t = (b_t^1, b_t^2, ..., b_t^K)$. In the same manner, the probability at time t that agent i adds a ball of color j is defined by a sequence of continuous function $\{q_t^{ij}\}$.

Ferrer and Sole define the aggregation function Ω_t^{ij} (for agent i and ball color j), which combines the probabilistic choices of all agents connected to the i-th agent, in the following way

$$\Omega_t^{ij}(X_t) = \sum_{k=1}^{N} \Psi_t^{ik} \beta_t^{kj}(X_t^k) \tag{36}$$

where

$$\beta_t^{ij}(x) = \left\{ \begin{array}{c} 1 \text{ with a probability } q_t^{ij}(x) \\ 0 \text{ with a probability } 1 - q_t^{ij}(x) \end{array} \right\} \tag{37}$$

At time t, if agent i is chosen, the dynamics of the number of balls of type j, b_t^{ij}, and of the number of time T_t^i the agent i has been selected till time t are the following:

$$b_{t+1}^{ij} = b_t^{ij} + \Omega_t^{ij}(X_t) \tag{38}$$

$$T_{t+1}^i = T_t^i + 1 \tag{39}$$

And if the agent i has not been chosen

$$b_{t+1}^{ij} = b_t^{ij} \tag{40}$$

$$T_{t+1}^i = T_t^i \tag{41}$$

At time t, the number of balls contained in the urn of agent i is $w + T_t^i \cdot C$ and the proportion of balls of color j for agent i is the following

$$X_t^{ij} = \frac{b_t^{ij}}{w + T_t^i \cdot C} \tag{42}$$

In order rewrite equation 42 like equation 28, let us define T_t^{*i} as:

$$T_t^{*i} = T_t^i + 1 \tag{43}$$

$$X_t^{ij} = \frac{b_t^{ij}}{w + (T_t^{*i} - 1) \cdot C} \tag{44}$$

If the agent i has not been selected

$$X_{t+1}^{ij} = X_t^{ij} \tag{45}$$

If the agent i has been selected at time t, equation 38 can be rewritten as:

$$X_{t+1}^{ij} \cdot (w + (T_{t+1}^{*i} - 1) \cdot C) = X_t^{ij} \cdot (w + (T_t^{*i} - 1) \cdot C) + \Omega_t^{ij}(X_t) \tag{46}$$

$$X_{t+1}^{ij} \cdot (w + T_t^{*i} \cdot C) = X_t^{ij} \cdot (w + T_t^{*i} \cdot C - C) + \Omega_t^{ij}(X_t) \tag{47}$$

$$X_{t+1}^{ij} \cdot (w + T_t^{*i} \cdot C) = X_t^{ij} \cdot (w + T_t^{*i} \cdot C) + \Omega_t^{ij}(X_t) - C \cdot X_t^{ij} \tag{48}$$

$$X_{t+1}^{ij} = X_t^{ij} + \frac{\Omega_t^{ij}(X_t) - C \cdot X_t^{ij}}{w + T_t^{*i} \cdot C} \tag{49}$$

This equation can be rewritten in a form similar to the fundamental equation 32, by defining

$$\Phi_t^{ij}(X_t) = \sum_{k=1}^{N} \Psi_t^{ik} q_t^{kj}(X_t^k) \tag{50}$$

$$X_{t+1}^i = X_t^i + \underbrace{\frac{1}{w + T_t^{*i} \cdot C}[\Phi_t^{ij}(X_t) - C \cdot X_t^i]}_{\text{first part}} + \underbrace{\frac{1}{w + T_t^{*i} \cdot C}[\Omega_t^{ij}(X_t) - \Phi_t^{ij}(X_t)]}_{\text{second part}} \tag{51}$$

as

$$E[\Omega_t^{ij}(X_t) - \Phi_t^{ij}(X_t)|X_t] = 0 \tag{52}$$

only the first part of the equation directs the dynamics.

The formulation of equation 51 is not strictly equivalent to equation 32 as the denominator of the first part now depends not only of t but also of i, with the term T_t^{*i}. Based on this formulation, Ferrer and Sole study the conditions for spontaneous consensus in the case of the max and id urn function. Their conclusion support the experimental findings of section 2 [12].

4 A more complex model

Most of the distributed coordination systems studied so far in the context of emergent language processes are self-organizing lexicons [1, 5, 10, 11, 15, 17, 18, 19, 25, 27, 28, 31, 33, 34, 35]. In this section we will discuss how the properties characterized for simple models of distributed coordination scale to a classic model of self-organizing lexicon.

Model D. Each agent is now equipped with an associative memory where associations between a convention set $\mathcal{C} = \{c_1, c_2 ... c_{|\mathcal{C}|}\}$ and a set of states $\mathcal{S} = \{s_1, s_2 ... s_{|\mathcal{S}|}\}$ are stored. In classic models of self-organizing lexicons, states are often referred as *meanings*, *objects* or *referents* and conventions as *words* or *signals*. We prefer to use the terms *states* and *conventions* as they are more neutral and account for more diverse interpretations of the dynamics studied. In the matrix M_a, $m_{a,i,j}$ is the score of the association between the state s_i and the convention c_j.

$$M_a = \left\{ \begin{array}{c} m_{a,1,1} \ldots m_{a,1,|S|} \\ m_{a,2,1} \ldots m_{a,2,|S|} \\ \ldots \ldots \ldots \\ m_{a,|\mathcal{C}|,1} \ldots m_{a,|\mathcal{C}|,|S|} \end{array} \right\} \tag{53}$$

Like in the other models, two agents are picked at random in the population at each iteration. A state s_h is also chosen a random. Agent a_1 produces a convention c_k by choosing the convention associated with the biggest score in the column h.

$$\mathcal{P}_D(M_{a_1}, s_h) = c_{argmax_i(m_{a_1,i,h})} = c_k \tag{54}$$

Agent a_2 uses an interpretation rule $\mathcal{I}_D$ to decode c_k into a possible state using its own matrix. It chooses the state s_l corresponding to the strongest association with the convention c_k (highest score of line k).

$$\mathcal{I}_D(M_{a_2}, c_k) = s_{argmax_j(m_{a_2,k,j})} = s_l \tag{55}$$

If $l = h$ the communication is a success, otherwise it is a failure. In this model, different rules of adaptation are used depending on the cases. If communication is a success, agent a_2 increases the winning association (k, l) and decreases competition associations (this rule is called lateral inhibition by [37, 27]). If the communication is a failure association (k, l) is decreased and association (k, h) is increased (this supposes the existence of another type of signaling permitting agent a_2 to have access to the intended state s_l). Most models use adaptation rules similar to these ones. Some do not use different adaptation rules for success and failure and assume that (state, convention) pairs can be systematically observed by agent a_2 (e.g. [31]). The choice of the particular rules used in model D is motivated by empirical investigations conducted in [19].

$$\mathcal{U}_{D,l=h} : \left\{ \begin{array}{c} m_{a_2,i,j} \leftarrow m_{a_2,i,j} + \delta \text{ if } i = k \text{ and } j = l \\ m_{a_2,i,j} \leftarrow m_{a_2,i,j} - \delta \text{ if } i = k \text{ and } j \neq l \\ m_{a_2,i,j} \leftarrow m_{a_2,i,j} - \delta \text{ if } i \neq k \text{ and } j = l \end{array} \right\} \tag{56}$$

$$\mathcal{U}_{D,l\neq h} : \left\{ \begin{array}{c} m_{a_2,k,l} \leftarrow m_{a_2,k,l} + \delta \\ m_{a_2,k,h} \leftarrow m_{a_2,k,h} - \delta \end{array} \right\} \tag{57}$$

Initially, each agent has no preferences (all $m_{a,i,j} = 0$). We will assume that the number of possible conventions is much bigger than the number of states: $|\mathcal{C}| \gg |\mathcal{S}|$. This is equivalent to systems in which words are created on the fly (e.g. [33]). This permits to ensure that the population converges towards a shared coding [19].

We can describe the overall behavior of the population by defining a probabilistic function $p(c_i|s_j)$, giving the probability of using convention c_i for state s_j. In the same manner, the probabilistic function $i(s_i|c_j)$ can be used for the interpretation of convention c_j as state s_i. Both functions can be obtained by averaging the production and interpretation behavior resulting of the set of matrix $\{M_a\}$ at a given point in the evolution. We can thus define formally the *communication accuracy* ca of the population in the following way (see also [10, 26, 27, 31] for similar definitions):

$$ca = \frac{1}{|\mathcal{S}|} \sum_{i=1}^{|\mathcal{S}|} \sum_{j=1}^{|\mathcal{C}|} p(c_j|s_i).i(s_i|c_j) \tag{58}$$

By similarity with our previous definition of the coherence level, coherence level in production for state s_j can be defined as:

$$CL_P(s_j) = max_{i=1..|\mathcal{C}|}(p(c_i|s_j)) \tag{59}$$

By averaging over the different possible states, we can define the global coherence level in production:

$$CL_P = \frac{1}{|\mathcal{S}|}\sum_{j=1}^{|\mathcal{S}|} CL_P(s_j) \tag{60}$$

Similarly we can define the coherence level in interpretation for convention c_i and global coherence level in interpretation.

$$CL_I(c_j) = max_{j=1..|\mathcal{S}|}(i(s_j|c_i)) \tag{61}$$

$$CL_I = \frac{1}{|\mathcal{C}|}\sum_{i=1}^{|\mathcal{C}|} CL_I(c_i) \tag{62}$$

When $ca = 1$, all communication interactions between agents are successful. This neither implies $CL_I = 1$ nor $CL_P = 1$. A partial coherence in interpretation is possible as long as the coordination is complete for the convention actually produced. It does not matter for instance that agents give different interpretations of convention c_1, if this convention is never produced by any of the agents. In the same manner, a partial coherence in production ($CL_P < 1$) is possible if the different conventions used for the same state are systematically interpreted in the same manner. In an inverse manner, $CL_I = 1$ and $CL_P = 1$ does not impose $ca = 1$. For instance, s_1 and s_2 can be associated with the same convention c, c being systematically decoded into s_3 different from s_1 and s_2. In such a case, coordination of the system is complete but communication is impossible. This is why $ca = 1$ is usually chosen as the end criteria for simulations about the self-organization of conventional communication systems (see [10] for a related discussion about *perfect* communication systems).

Exp D.a ($N = 10$, $|\mathcal{S}| = 10$, ,$|\mathcal{C}| = 100$, End criteria: $ca = 1$, 1 run) Figure 10 shows a sample evolution of ca, CL_I and CL_P for 10 agents and 10 states. An efficient conventional communication system is established around iteration 1600. In the course of the evolution, coordinated interpretation arises before coordinated production.

Exp D.b (different N and$|\mathcal{S}|$, End criteria : $ca = 1$) We can now study the dependency of the convergence time T_c (time to reach $ca = 1$) to the population size N and the number of states$|\mathcal{S}|$. Figure 11 plots T_c divided by $N \cdot |\mathcal{S}|$ for different values of N and $|\mathcal{S}|$. In the various experiment $||\mathcal{C}|| = N \cdot |\mathcal{S}|$. Data suggests a linear dependency of $\frac{T_c}{N.|\mathcal{S}|}$ in $log(N)$ and $|\mathcal{S}|$ of the following type:

$$\frac{T_c}{N \cdot |\mathcal{S}|} \approx k_0 + k_1 \cdot LogN + k_2 \cdot |\mathcal{S}| \tag{63}$$

Values obtained by linear regression are $k_0 = -1.34$, $k_1 = 16.0$ and $k_2 = 1.17$. The corresponding plane is represented on figure 11. As k_2 is ten times smaller than k_1, T_c is approximatively proportional to $|\mathcal{S}|.N.LogN$:

$$T_c \propto |\mathcal{S}| \cdot N \cdot LogN \tag{64}$$

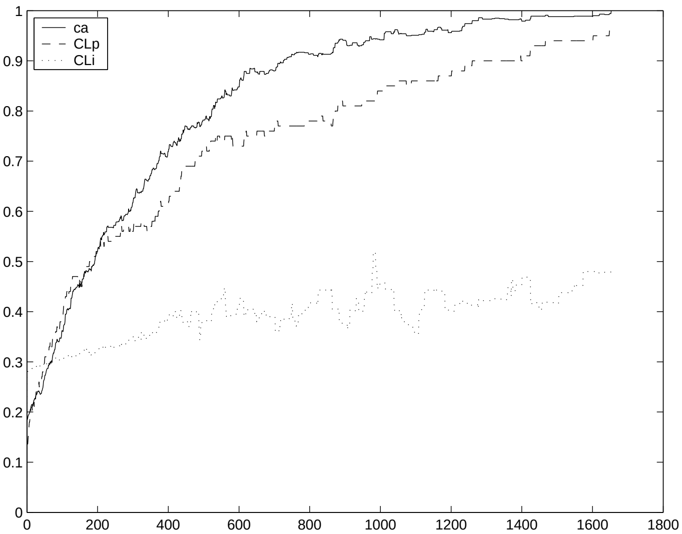

Figure 10: Lexicon self-organization. Evolution of the communicative accuracy ca, coherence level in production and interpretation CL_P and CL_I. 10 agents have to agree on shared mapping for 10 states, using a set of 100 conventions. An efficient conventional communication system is established around iteration 1600 (Exp D.a)

We can understand this finding intuitively. Because $|\mathcal{C}| \gg |\mathcal{S}|$, cases of competition of the same convention c_k for several different states are rare. The dynamics can be understood as $|\mathcal{S}|$ parallel competitions with only few interactions between them. This is similar to a situation in which these competitions would be conducted one after another. Therefore, it is natural to find again the $N \cdot LogN$ dependency multiplied by the number of states $|\mathcal{S}|$. However, for situations in which the different competitions would have complex interferences, the linear dependency in $|\mathcal{S}|$ may not be good approximation anymore.

Can model D be interpreted in one of the theoretical framework we considered in section 3? Model D, like models B and C do not respect the Markov property because of the historical character of the update rules used. The complexity of the model makes it also difficult to formulate in a stochastic game framework. Interpretation in terms of Polya processes is more promising. As suggested by Ferrer and Sole, extension of the model of equation 51 to allow more than one urn per agent can be realized with just a syntactic improvement, adding an additional index to distinguish the agent the urn belongs to [12]. They established a series of preliminary results in that direction. Working out the formal properties than can drawn from an interpretation of model D in such a framework will be the subject of future studies.

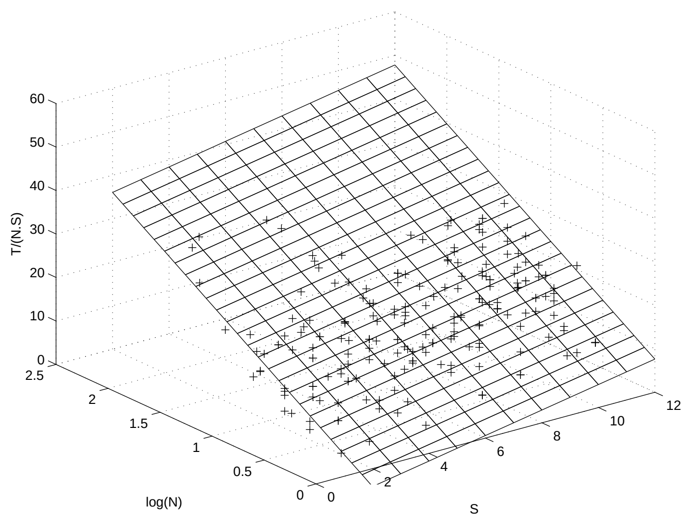

Figure 11: Convergence time T_c compared to population size N and to the size of the state space $|\mathcal{S}|$. Results suggest that in first approximation T_c increases in $|\mathcal{S}| \cdot N \cdot log(N)$ (Exp D.b)

5 General summary and conclusions

Simple models for distributed coordination have been studied in this paper from empirical, formal and qualitative perspectives. The models were deliberately simplified compared with architectures usually studied in research about self-organizing communication systems. Results and conjectures that were drawn from these models are the following.

- Two kinds of dynamics can lead to consensus. The slowest one has similar dynamics than a random walk, the faster one (self-reinforcing dynamics) has dynamics similar to several other systems with positive feedback loops.

- These models of distributed coordination can be interpreted using various formalism including Markov chains, stochastic games and Polya processes. The advantages and limitations of formal interpretations within these different framework were discussed. This discussion suggests that Polya processes are the most promising models to address formally distributed coordination in emergent communication systems.

- Both empirical results and qualitative interpretations suggest that convergence time of models with self-reinforcing dynamics is proportional to $N.Log(N)$, where N is the population size. This conjecture is experimentally verified with more complex models of lexicon self-organization.

The following questions arise naturally from this preliminary study. How much of the dynamics of more complex existing models described in the literature can be accounted with results described in this article? Are empirically observed convergences in these systems due to self-reinforcing dynamics (as it is in most of the cases assumed) or to dynamics similar to random walks? Are there intermediate cases between self-reinforcing dynamics produced by greedy production rules (like $\mathcal{P}_B$) and dynamics resulting of probabilistic rules (like $\mathcal{P}_C$)? And finally: How general is the $N \cdot LogN$ convergence?

6 Acknowledgments

Research funded by Sony CSL Paris with additional support from the ECAGENTS project founded by the Future and Emerging Technologies program (IST-FET) of the European Community under EU R&D contract IST-2003-1940. The authors would like to thank three anonymous reviewers who have greatly contributed to increase the quality of this article.

References

[1] T. Arita and Y. Koyama. Evolution of linguistic diversity in a simple communication system. *Artificial Life*, 4(1):109–124, 1998.

[2] B. Arthur, Y. Ermoliev, and Y. Kaniovski. A generalized urn problem and its applications. *Cybernetics*, 19:61–71, 1983.

[3] B. Arthur, Y. Ermoliev, and Y. Kaniovski. Strong laws for a class of path-dependent stochastic processes, with applications. In Arkin, Shiryayev, and Wets, editors, *Proceedings of the conference on stochastic optimization*, Lectures notes in control and information sciences, Berlin, 1984. Springer-Verlag.

[4] B. Arthur, Y. Ermoliev, and Y. Kaniovski. Path dependent processes and the emergence of macrostructure. In B. Arthur, editor, *Increasing returns and path dependence in the economy*, chapter 3, pages 33–48. The University of Michigan Press, Ann Arbor, MI, 1994.

[5] A. Cangelosi and D. Parisi. The emergence of a 'language' in an evolving population of neural networks. *Connection Science*, 10(2):83–97, 1998.

[6] A. Cangelosi and D. Parisi. *Simulating the Evolution of Language*. Springer, 2002.

[7] A. Cangelosi and D. Parisi. The processing of verbs and nouns in neural networks: Insights from synthetic brain imaging. *Brain and Language*, 89(2):401–408, 2004.

[8] F. Chung, S. Handjani, and D. Jungreis. Generalizations of polya's urn problem. *Annals of combinatorics*, pages 141–154, 2003.

[9] B. De Boer. *The origins of vowel systems*. Oxford University Press, 2001.

[10] E.D. De Jong and L. Steels. A distributed learning algorithm for communication development. *Complex Systems*, 14(4-5):315–334, 2003.

[11] C. Dircks and S. Stoness. Effective lexicon change in the absence of population flux. In D. Floreano, J-D Nicoud, and F. Mondada, editors, *Advances in Artificial Life (ECAL 99)*, Lecture Notes in Artificial Intelligence 1674, pages 720–724, Berlin, 1999. Springer-Verlag.

[12] R. Ferrer and R. Sole. Naming games through distributed reinforcement. unpublished report, 1998.

[13] R. Ferrer and R. Sole. Least effort and the origins of scaling in human language. *Proceeding of National Academy of Science USA*, 100:788–791, 2003.

[14] B. Huberman and T. Hogg. The behavior of computational ecologies. In B. Hberman, editor, *The ecology of computation*. Elsevier Science, 1988.

[15] E. Hutchins and B. Hazlehurst. How to invent a lexicon: the development of shared symbols in interaction. In N. Gilbert and R. Conte, editors, *Artificial Societies: The Computer Simulation of Social Life*, pages 157–189. UCL Press, London, 1995.

[16] M. Iosifescu and R. Theodorescu. *Random processes and learning*. Springer-Verlag, 1969.

[17] F. Kaplan. A new approach to class formation in multi-agent simulations of language evolution. In Y. Demazeau, editor, *Proceedings of the third international conference on multi-agent systems (ICMAS 98)*, pages 158–165, Los Alamitos, CA, 1998. IEEE Computer Society.

[18] F. Kaplan. Semiotic schemata: Selection units for linguistic cultural evolution. In M Bedau, J. McCaskill, N. Packard, and S. Rasmussen, editors, *Proceedings of Artificial Life VII*, pages 372–381, Cambridge, MA, 2000. The MIT Press.

[19] F. Kaplan. *La naissance d'une langue chez les robots*. Hermes Science, 2001.

[20] J. Ke, J. Minett, C-P. Au, and W. Wang. Self-organization and selection in the emergence of vocabulary. *Complexity*, 7(3):41–54, 2002.

[21] K. Khanin and R. Khanin. A probabilistic model for establishment of neuron polarity. Technical report, HPL-BRIMS-2000-16, 2000.

[22] R. Kinderman and S.L. Snell. *Markov random fields and their applications*. American mathematical society, 1980.

[23] S. Kirby. Natural language and artificial life. *Artificial life*, 8:185–215, 2002.

[24] D. Lewis. *Convention: A philosophical study*. Harvard University Press, 1969.

[25] D. Livingstone and C. Fyfe. Modelling the evolution of linguistic diversity. In D. Floreano, J-D Nicoud, and F. Mondada, editors, *Advances in Artificial Life (ECAL 99)*, Lecture Notes in Artificial Intelligence 1674, pages 704–708, Berlin, 1999. Springer-Verlag.

[26] M. Nowak and D. Krakauer. The evolution of language. *Proceedings of the National Academia of Science USA*, 96:8028–8033, 1999.

[27] M. Oliphant. *Formal approaches to innate and learned communicaton: laying the foundation for language*. PhD thesis, University of California, San Diego, 1997.

[28] P-Y. Oudeyer. Self-organisation of a lexicon in a structured society of agents. In D. Floreano, J-D Nicoud, and F. Mondada, editors, *Advances in Artificial Life (ECAL 99)*, Lecture Notes in Artificial Intelligence 1674, pages 726–729, Berlin, 1999. Springer-Verlag.

[29] P-Y. Oudeyer. The self-organization of speech sounds. *Journal of Theoretical Biology*, 233(3):435–449, 2005.

[30] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, 94(1–2):139–166, 1997.

[31] K. Smith. The evolution of vocabulary. *Journal of Theoretical Biology*, 228(1):127–142, 2004.

[32] K. Smith, S. Kirby, and H. Brighton. Iterated learning: a framework for the emergence of language. *Artificial Life*, 9(4):371–386, 2003.

[33] L. Steels. Self-organizing vocabularies. In C. Langton and T. Shimohara, editors, *Proceeding of Alife V*, Cambridge, MA, 1996. The MIT Press.

[34] L. Steels and F. Kaplan. Spontaneous lexicon change. In *Proceedings of COLING-ACL 1998*, pages 1243–1249, Montreal, August 1998. ACL.

[35] L. Steels and F. Kaplan. Stochasticity as a source of innovation in language games. In C. Adami, R. Belew, H. Kitano, and C. Taylor, editors, *Proceedings of Artificial Life VI*, pages 368–376, Cambridge, MA, June 1998. The MIT Press.

[36] L. Steels and F. Kaplan. Situated grounded word semantics. In T. Dean, editor, *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence IJCAI'99*, pages 862–867, San Francisco, CA., 1999. Morgan Kaufmann Publishers.

[37] L. Steels and F. Kaplan. Bootstrapping grounded word semantics. In T. Briscoe, editor, *Linguistic evolution through language acquisition: formal and computational models*, chapter 3, pages 53–73. Cambridge University Press, Cambridge, 2002.

[38] P. Vogt. Bootstrapping grounded symbols by minimal autonomous robots. *Evolution of Communication*, 4(1):89–118, 2000.

[39] P. Vogt. Minimum cost and the emergence of the zipf-mandelbrot law. In *Proceedings of the 9th Artificial Life conference*. MIT Press, 2004.

[40] W. Weidlich and G. Haag. *Concepts and models of a quantitative sociology; the dynamics of interacting populations*. Springer-Verlag, 1983.