

Agent Based Modelling of Communication Costs: Why Information Can Be Free

Ivana Čače¹ and Joanna J. Bryson²

¹ Universiteit Utrecht ivana@phil.uu.nl

² University of Bath jjb@cs.bath.ac.uk

Abstract. In this chapter we examine the costs and benefits of sharing information, using a simple spatial agent model. The information concerns the agents' food. The costs are for the signaling agent in terms of more competition over food, and the benefits are for the receiver in terms of access to more food. We will show that under some conditions agents that communicate do better — that is, have more offspring — than agents that do not.

1 Why We Have Language

What purposes, other than facilitating the sharing of information, can language have served? First, it may not have evolved to serve any purpose at all. It is possible that language is just a side effect of the large human brain — a spandrel or exaptation — that only became useful later. If language is adaptive, this does not necessarily mean that it is adaptive for the purpose of communication. For example Dennet (1996) and Chomsky (1980) have stressed the utility of language in thinking. Also, there are different ways to view communication. The purpose of language according to Dunbar (1993), is to replace grooming as a social bonding process and in this way to ensure the stability of large social groups.

Why would anyone think that communication might be a bad thing? Because communication is a form of altruism, it is a giving up of information to one's competitors. Remember, even friends are competitors in a biological sense, as they share the same food resources and mating opportunities. Consequently, many people find it unlikely that such altruism could evolve (Dessalles 2000; Marshall and Rowe 2003).

In the next section we will show how communication can be viewed in terms of costs and benefits. We will discuss some theories that fit in this framework and propose a slightly different approach. In the next section we will discuss communication in relation to parenting (kin-selection) and foraging abilities. We will then describe the experiment we used to show that communication is indeed a good thing.

2 The Signaling Framework

2.1 Introduction

The study of communication in terms of costs and benefits has originated in biology but has also been applied to the emergence of language. A scientifically effective way of looking at animal signals is in terms of mind-reading and manipulation (Krebs and Dawkins 1984). The signaling agent takes the role of manipulator while the receiving agent is the mind-reader. The manipulator sends some signal — intentionally or not — that will incite the mind-reader to some response. The mind-readers response might be costly to the mind-reader itself, to the manipulator or it could be advantageous to either or both.

When signaling is viewed in this way it can be formalized and then studied using for example game-theory or a simulation. Hopefully this will help explain some of the properties of signals found in nature. Animals will have evolved to maximize their gain and minimize their costs in their roles as signalers and receivers.

The honesty of the signal often determines the cost of/for the signal. Animals signal about themselves, most importantly their fitness or their intentions, or they signal about the environment. In a cooperative setting both signaler and receiver stand to gain from an accurate response to a honest signal. But in case of competition the signaler may gain more by manipulating the receiver to a response that may not be the most accurate for the given conditions. The signaler may try to pretend to be a fitter mate or a stronger adversary — or food instead of predator. So how does the receiver protect itself from being cheated? One solution, proposed first by Zahavi (1975), is to ensure signals are so costly, the weaker animals cannot afford to produce them. This is the so called “costly signaling” hypothesis and it has since been elaborated on by many others (Maynard-Smith and Harper 1995). This mechanism will obviously only work if the animal is signaling about itself. Other ‘solutions’ involve a penalty for lying or signals that are necessarily honest.

Formally the costs-and-benefits framework can be described as follows. Assume that not signaling has some effect on the signalers and receiver’s fitness, the null hypothesis receiver n_r and the null hypothesis signaler n_s . Then signaling will be adaptive if, on average:

$$f_r > n_r \text{ and } c + f_s > n_s \quad (1)$$

Where f_r , f_s and c are defined as in Fig. 1.

Altruism The costs-and-benefits framework can also be applied to the emergence of language. If language is simply perceived as signaling, the same question arises of how to keep it honest. Another problem occurs when language is taken to be primarily a tool for accurately sharing information. ³ If information is a

³ Simple signaling too transfers information in the sense that it reduces the receiver’s uncertainty about the signaler or the environment.

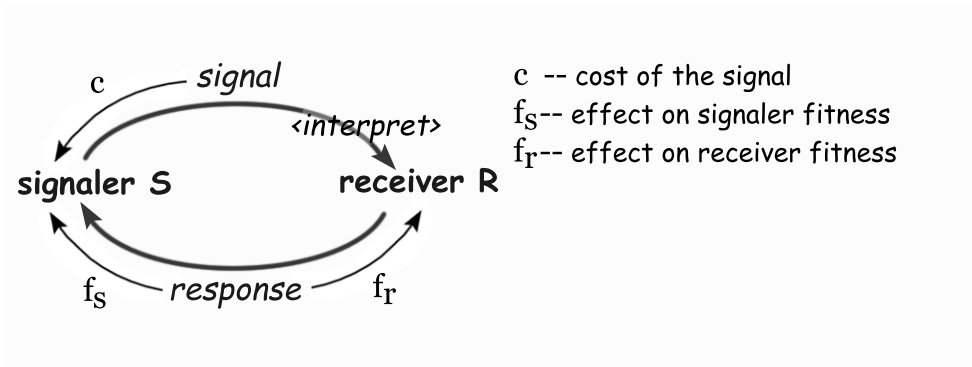


Fig. 1. The signaling game has fitness effects for both the signaller and receiver. Signaling will be adaptive if S and R both benefit from it.

resource then sharing it is altruistic behavior. The advantages of altruistic behavior to the group are evident, so long as it helps others more than it hurts the altruist. However, unless the behavior somehow increases the fitness of the altruistic individual it is not in itself adaptive, so it is unlikely to be selected. The two main mechanisms that are thought to underly apparent altruistic behavior are: reciprocal altruism and altruism in favor of kin.

Reciprocal altruism is possible if there are repeated interactions in which agents take turns in their roles as signallers and receivers. Over a longer period of time everyone involved could profit from these interactions. Reciprocal altruism presupposes a method of distinguishing between agents that will cooperate, or have cooperated in the past, and the ‘free-riders’ that will not. Cox et al. (1999) have shown that altruism will persist in an agent-based model if the agents are endowed with a memory that allows them to remember which other agents are likely to cooperate.

Cooperation among related individuals can exist even if it is not mutually beneficial. The success of an individual’s *genes* depends on the individual’s own reproductive success and that of others with the same genetic material. This notion has been formalized by Hamilton (1964). Hamiltons rule relates the inclusive fitness (i.e. reproductive success) f of an altruistic act to c , its cost to the actor, and to the relatedness r of the recipients to the actor and the benefit b they have from it:

$$f = c + \sum (r * b)$$

Note that kin selection implies that either the agent has control over its behavior and can distinguish between related and not related agents, or the altruistic behavior is involuntary but will mostly be triggered when related agents are around to profit. The close presence of kin is an important feature of so called viscous populations where individuals have limited mobility. Such popu-

lations are thought to facilitate cooperation and altruism⁴ (Marshall and Rowe 2003; Hamilton 1964). When agents remain close to their place of origin they are more likely to be related to nearby individuals. Any altruism towards neighboring agents is likely to benefit kin. But at the same time competition between nearby individuals will be counteracting cooperation. Kin selection is also possible if the population structure is such that individuals mingle with others they are not related to. Only now the altruistic behavior will be voluntary and directed towards kin, and there will be some level of cognitive load. The agent will have to know somehow if it is related to another agent (and to what degree). This is a situation similar to reciprocal altruism, when agents have to keep track of which other agents are likely to cooperate.

2.2 Models

A number of theories have been brought forward to explain why we have language despite the opportunities it offers for cheating and despite it being the foundation of information sharing, a basically altruistic behavior. Because language is an example of conventional signaling, the signals themselves are necessarily cheap. If there is a price it will come from the content of the signal. The information being transferred might have cost the signaler time and energy to acquire. Sharing this information may have fitness effects, positive or negative, for the signaler and the receiver. We will now discuss some of the theories that have applied the framework described above to language and that do not rely on either kin-selection or reciprocal altruism.

Table 1. Different ways to ensure that communication benefits both signaler and receiver. The second column shows how this is achieved (based on variables in Fig. 1).

variable	explanation
$\uparrow c$	S signals fitness to prospective mating partner. R cannot assess S's fitness, the signal will be honest if S cannot lie. One way to ensure this is if c of a signal indicating high quality is too high for a low quality individual.
$\uparrow f_r$	Over a longer period, the effects of acquiring useful information negates the effects of being cheated or manipulated (Simon 1990).
$\uparrow -f_s$	S signals its ability to fight to a competitor. R can choose to test S, if S lied it will run the risk of getting hurt more than a stronger individual (Lachmann and Bergstrom 2001).
$c + f_s > 0$	S tells R some fact K, this will cost it c . If K proves useful, S gains status: $\uparrow f_s$ (Dessalles 2000).

⁴ Viscosity is also thought to facilitate reciprocal altruism because of repeated interactions among individuals. This is the first requirement for cooperation among agents that are not related.

From the receivers perspective, trust will pay off if the received information is worth more (on average) than the price of being cheated. Simon (1990) made a convincing case for cooperation and altruism as social learning. He has argued that the complexity of the environment people live in makes docility, that is receptivity to social influence, an adaptive strategy. "... the docile individual will often be unable to distinguish socially prescribed behavior that contributes to fitness from altruistic behavior." Human life history provides the perfect setting for social learning. Children are born with a undeveloped brain and require many years of nurture by adults during a time in their development in which they are very sensitive to learning. As adults too, we continue to live in social groups. This facilitates communication. In today's society information is important but the survival skills of early humans too included information about plant and animal food sources and how to make tools.

On the signaling side there may be a cost to lying or a gain to producing accurate, useful information. We will give examples of both. We have already mentioned the model used by Lachmann and Bergstrom (2001) to show that a penalty for lying will incite honesty and make communication adaptive. They use a mathematical (game theoretic) model and describe the effect of various cost functions on the signaling-equilibrium between signaler and receiver. According to Lachmann and Bergstrom the results of their formal model are applicable to human language if there are socially imposed sanctions against liars. This requires that other group-members know and care about the accuracy of the information. If this is the case the costs of lying will not only serve to favor honesty but will also select for the more careful and redundant signals that become possible through linguistic communication. Language may also facilitate social enforcement by making it easier to share information about others' reliability.

The possible gain from honesty can be explained as classic costly signaling if one assumes a cost to gathering and/or dispensing information. Dessalles (2000) uses a simulation of his model to show that "the evolution of language is at least conceivable in a context of political competition between coalitions". First he proposes a model in which communication has a cost. He then shows that there is no selective pressure for communication as reciprocal altruism in this model. Next he proposes a model in which there is also a *gain* to communicating successfully. Receivers rate the usability of the information. Agents that are known to supply others with high quality information gain status through coalitions. Coalitions are defined as "groups of individuals showing solidarity in action". In Dessalles' model there is competition between coalitions and individual fitness is tied in with the success of the coalition. Agents that supply useful information are readily joined by others. In a sense the individuals signal their fitness only now it is not determined by physical strength or speed, but by the ability to enter a successful coalition.

Does there have to be a Price? In the models discussed above there was always a presupposed cost to (honest) communication. Without abandoning the costs-benefits framework altogether, it is possible to find support for the idea

that sharing knowledge may be useful if we take a more comprehensive view. Communication could offer advantages in the context of foraging and parenting, as we show in the next section. We have built a model to see *if* there is a cost to sharing knowledge, and found that this is not always the case. We challenge the hypothesis that compulsively sharing valuable knowledge is in itself not adaptive.

2.3 Food and Kin

Buckley and Steele (2002) have made an extensive survey and categorization of research in the area of the emergence of language and linked it to archaeological and palaeontological parameters that can aid in determining the plausibility of different hypotheses. This leads them to the conclusion of “life history strategy as the prime mover, and co-operative foraging and provisioning as the selective context for spoken language abilities.”

They organize social models for the evolution of language into 3 groups: parenting strategy, status and rank, and co-operation. Each of these models, or rather models that emphasize on each of these three social contexts, entails certain conditions that can be compared to archaeological evidence. They find no reliable archaeological evidence for increase in group size, and therefore reject the third type of model. Furthermore, hominin fossil anatomy shows that the differences between male and female (in terms of body mass and canine size) became smaller. According to Buckley and Steele this “does not suggest that male-male competition for mates was a driver for the evolution of social communication strategies.” Archaeological evidence does favor the first kind of model, which emphasizes the benefits of communication in the provisioning effort for child rearing. Increase in provisioning effort is tied in to the increase in hominin brain size and the subsequent prolonged immaturity in children. We will return to this shortly when we consider the role of communication in foraging. First we discuss the connection between communication and child-rearing.

Communication among Kin We have discussed the altruism in favor of kin above. An example of communication that can be viewed in this way is teaching. Communication among kin can also *facilitate* the development of communication abilities. The selective advantage that communication provides may not be enough to ensure that the inclination to communicate will spread through the population. The selective process will be frequency dependent, that is the benefits of communication will depend on the proportion of communicating individuals. When the trait first appears the proportion of the communicating individuals in the group is very small and therefore the advantages of communicating will be very small or rare. As a result the speed with which the trait spreads through the population could become so low, that it can easily be disrupted by random events.

Cavalli-Sforza and Feldman (1983) have shown that this bottleneck is removed if the communication occurs among close relatives or “if aggregation of communication occurs because of assortative mating or meeting”. Cavalli-Sforza

and Feldman assume that the ability to communicate — that is the combined ability to emit the signal and to understand what is being communicated — is hereditary. Communication is used to transmit some fitness enhancing skill, that all adult communicators have already acquired. If these conditions are met the communicators' offspring will have a sufficiently high fitness increase, due to contact with (at least one of) their parents, for the evolution of communication to take off. Communication with other family members or a preference for the company of others to communicate with, may speed up this process even more. This way the payoff of communication might be higher (if there is enough useful information to be communicated) for both adults and children. In this analytical model the adaptive pressure towards communication will persist, provided the relation between costs, benefits and the chances of meeting another communicator are favorable. For the process described above to work, the property in question need not be hereditary. Communication (cooperation, altruism) could very well be a learned trait, that is, the whole population has been born with the ability but only those who have been raised in a 'communicating environment' will display the actual behavior.

The Expensive Tissue Hypothesis Language is one of the things that sets us apart from the other primates, having a much larger brain is another one. The brain uses about 20-25% of the body's energy. Humans have approximately the same calorie-intake as chimpanzees, so how can we afford to feed our brain? The answer is that our intestine needs a lot less to uptake as many calories. The Expensive Tissue Hypothesis (Aiello and Wheeler 1995) states that the large human brain has been made possible by our small gut which was in turn made possible by a rich diet. Food of low digestibility requires relatively large guts while food of high digestibility (such as sugary fruits, protein and oil rich seeds and animal material) requires relatively smaller guts. This is true both in the general case (compare for example a cow with any given carnivore) as in the specific case among primates. The idea that language could have influenced diet was derived from one due to Steele (2004). Steele's idea was that language and a larger brain may have co-evolved. A large brain does not only offer many exiting possibilities in terms of cognitive functions, it also has a price. Because of the large brain children take a longer time to mature and place a high burden on those who provide for them. This would make it necessary to find new and more efficient ways of finding food. What if language is one of the factors that enabled us as a species to obtain this richer — but harder to get — food? Then having a larger brain could in turn have facilitated language.

3 Experiment

3.1 Agent Based Modeling

In this section we briefly introduce Agent Based Modeling (ABM) and describe our model. We will than move on to describing the experiment and the results.

As a research tool, ABM is used in a variety of areas including: social science and economics (Epstein and Axtell 1996; Cederman and Rao 2001), animal behavior (Hemelrijk 2000) and complex systems in general (Baray 1998; Esteva et al. 2001). ABM has also been applied before to the question of communication (Noble et al. 2002; Dessalles 2000). Because communication is a form of social interaction, examining this behavior is best done by simulating social interactions. ABM is the most suitable method for doing this.

Simulations and theoretical models are only tools to test the soundness of a hypothesis. We found that freely sharing information about food does not need to be costly. Our results signify something about the preconditions for the emergence of language. However, the use of models can only help clarify the mechanisms behind emergence of language, real-world data is needed to gain insight in the actual development undergone by early man. In other words we only describe a possibility. We do not make any claims about the evolutionary path from low level signaling to human language.

When we describe our model in the following section, it will become clear that we have not simulated any existing population. For example, the agents are presented with opportunities to feed at random, if there is more competition there will be less food. At this level of abstraction there is no distinction between collecting food or hunting for food.

3.2 Building the Model

The purpose of the model is to determine *if* communication about food is beneficial and to understand the mechanism that causes the costs and benefits in the context of honest communication about food. The model is intended to clarify rather than predict, so a simpler model is favored over a more complicated one. Other requirements are that it simulates communication and provides some way of measuring the effects it has on both the population as a whole and on individuals. That is, it needs to be suitable for comparing different behavioral *strategies*, communicating and not communicating. A number of simplifying assumptions have to be made in order to construct an abstract model with only a few factors, the influence of which can then be examined in turn. Bear in mind that the model should produce some kind of result within a reasonable amount of time when run on a computer, this has consequences for the choice of model. A model with a smaller number of agents runs faster but if the number drops below a certain threshold the agents may die out due to random population fluctuations. So the first crude constraint could be expressed as follows: the number of agents should be as small as possible but large enough to prevent arbitrary extinctions or other impairments. This entails a preference for a combination of model parameters that results in relatively small fluctuations in population size.

Our models were built in NetLogo (Wilensky 1999). NetLogo is a freely available multi-agent modeling environment, specifically designed for the ABM of natural and social phenomena. The NetLogo world consists of two kinds of programmable agents: an environment divided into patches and mobile agents called

turtles. Patches can represent environmental change, such as growing food, while turtles are typically used to represent animals, including humans.

In the model we used in this experiment the environment consists of 101x101 patch square on a torus space, which is presented on screen as a square. This means that agents that walk of the edge of the square will reappear on the opposite edge.

The agents move round, eat and reproduce. At every time-step they lose a small amount of their energy and if their energy level drops under 0 they die. The environment ‘grows food’, that is at every time-step every patch is filled with a unit of food with probability P . These food-growth settings for the model we used were inspired by Wilensky (1998). For example, with a probability of 0.016 and a field size of 10201 ($101 * 101$), around 163 food-units will be added to the environment. A patch can only contain one unit of food. To keep the model as simple as possible the agents reproduce a-sexually. When an agent reproduces its energy is divided between it and its offspring to keep the total energy constant, the new agent gets 20% of the parents’ energy. Note that this means that if agents that have low energy levels are allowed to reproduce, their offspring (and maybe even they themselves) have a small survival chance. We have made the agents’ probability of reproduction dependent on their energy, so as to give the weaker agents a better chance. With the progression of the agents with the better strategy slowed down, there is more time to study the system. New agents are instantiated on the same patch as their parent.

The agents have no cognitive architecture. After all we were testing the effects of communication and adding extra mechanisms to the model (e.g. perception, action selection) could cloud the main question. This does mean that intelligent walking patterns have to be simulated too. We accomplish this by generating a walking pattern that is found in foraging animals as well as in evolutionary optimized foraging agents (van Dartel et al. 2002). Prior to taking a step the agents make a random turn, then they take a step of length l (one unit is equivalent to the length of a patch) where:

$$l = (0.7/n)^{0.3}$$

n is a random value between 0 and 1, 0.7 is a normalizing factor. The value of 0.3 ensures that the agents’ most frequent step length is 1.⁵ This way viscosity is high and when communication is added to the model the impact of communication, in terms of more competition, is large. The density of the population has no impact on their mobility: any number of agents can be standing in the same patch.

These settings yield a stable model, that is, regardless of the *initial* values for the number of agents or the amount of food, the model will gravitate towards equilibrium values. Also, all else remaining equal, the number of agents depends

⁵ For larger values, up to 0.5, walking is even better optimized for foraging. The resulting pattern is one of random walk with a occasional large jump. Large jumps would significantly lower viscosity and competition and are therefore not desirable in this case.

on the amount of energy that is added to the environment. This means that adding more food will not serve to make the agents less ‘needy’ and decrease competition, it will only increase the number of agents the environment can sustain (the carrying capacity).

On top of this model the knowledge distribution and communication is modeled as follows. A second type of food is added to the environment. This ‘special food’ comes in 6 varieties and has twice the calorie value of the regular food. In order to collect any of the varieties, the agent needs specialized knowhow.

There are two ways to obtain the knowhow to eat one of the special food-types. Every agent has a small (5%) chance of knowing how to eat any one of the 6 items at birth. And, at every time-step, all the communicating agents pick one thing they know and tell all their neighbors. Agents are killed off when they have been around for more than 50 time-steps, to keep knowledge sparse. The knowhow itself is implemented as a 6-bit string. For example, an agent with string [0 1 0 0 1 0] knows two things. The pseudo-code for agents becomes:

```
take-credit
if (random energy) > 30 [offspring ]
set energy (energy - 1)
move
set age (age + 1)
if (breed = talker) [ communicate ]
```

At every time-step all the agents go through this routine, then the food in the environment is replenished (grown) as described earlier.

To prevent the special food from clogging up the environment when there is not enough knowhow among the agents to eat it, the regular food can overgrow the special food. The pseudo-code for the environment is:

```
for all patches:
fill patch with special food, with probability  $P_{special}$ 
fill patch with regular food, with probability  $P_{regular}$ 
```

A Priori Some things can be said about this model and models in general on forehand. Note that, as always in the agent based approach, only the agents and their environment are modeled. When the model is run, the complex behavior of the system emerges from the local interactions between the agents and between the agents and their environment. In this model the cost of communication is not given but emerges from the spatial interactions between the agents and their environment. If an agent shares food-information with its neighbors then it will have to deal with increased food competition. Both communication and competition take place locally. Small step size results in a viscous population as we discussed earlier. The question is which will be of greater influence: competition between nearby (possibly related) individuals or the positive effect of altruism towards nearby kin. In our model the costs of the communication are caused solely by competition. If the agents were to be more mobile communication

would remain without consequences *for the individual communicating agent*. It would quickly leave the vicinity of the agents it had shared its knowhow with.

It would be tempting to make an analogy between evolution proper and the way in which our model evolves. This would also be wrong. Obviously the scale is wrong, both in terms of population size as well as in terms of time. But on a more fundamental level, any model that only models behavior and not genes is not a valid model of evolution. It only checks phenotypical properties for their evolutionary usefulness. It does not say anything about the genetic possibilities and dependencies (like the one between language, brain size and general cognitive abilities). The fake evolution in the model should be regarded as a *tool* for determining the usefulness of some agent trait in some context. If in the model, certain behavioral traits are made hereditary and after some length of time agents with particular traits have outnumbered agents with other behaviors, then the former behavior is *better within the model context*.

3.3 Getting Results

There are only two types of agents, a distinction based on their strategies. The agents' fitness depends on the amount of food they are able to collect, so the ratio between the two types of agents is a direct measure of the success of their strategies. The most straightforward way to decide which is the better strategy, communicating or not communicating, is to run the model several times and see which agents reproduce better.

Figure 2 shows the proportion of communicating agents in the population for running the model under nine different settings. When either type of agent had a clear advantage over the other, the model was run 3 times with the fitter agents initially making up around⁶ 10% of the population. This way the fraction of the agent-type goes through values from 0.1 to 1. The results show no evidence that there can be a stable mixture of communicators and non-communicators. In the model with the smallest number of agents, 10% corresponds to around 30 agents. Should the number of agents fall below 20, the population will almost certainly die out. In the models where neither strategy provides a clear advantage we have run the model 4 times, twice for a ratio of 1 : 1 and twice for a ratio of 1 : 10. The models were run for a maximum of 5000 time-steps.

We have varied the stable number of agents by adding twice as much regular food, and the profitability of knowhow, by adding four times more special food. Note that with a higher agent density the agents have less probability of monopolizing knowhow about a particular type of food. These alternatives were run with and without communication, the two top rows in Fig. 2. In a model without communication there is no difference between the two types of agents. The proportion of communicating agents will still vary, indicating the influence of chance in this kind of agent model.

⁶ When the model is set up, the new agents are randomly assigned a type, with a probability of 0.1.

The communicating agents have a clear advantage, and this advantage is larger if more food that requires knowhow is available. Querying the running model about the agents' knowhow per age category shows that both the agent types learn during their lifetime, but that the talker agents get a head start. The offspring of communicating agents on average know how to eat a little more than one type of food. The average sum of the knowhow of the offspring of silent agents is 0.3 ($6 * 0.05$). This is a result of spatial interactions, the new talker agents are 'born', instantiated on the same spot as their talker parent. In a sense one could say the agents profit from parental investment. The parent may lose out on feeding opportunities but it does so in favor of its offspring. How much learning goes on in the rest of the agent's life depends on the number of communicators in the environment.

Now the next question is, without the parental investment, is communication still adaptive? The answer is no, as the bottom row in Fig. 2 shows. If the new agent is born at a random location in the field, the benefit from having a talker parent disappears. The agent shares its knowhow with new cohabitants that are not related. However the speed with which the talkers die out is lower than the speed with which the non-communicating agents died out in the runs described above. This shows that the cost of honest signaling is not very high. Another way to remove parental investment is by increasing agent mobility. We also ran the model with the three different food settings for larger steps. The agents took steps of random length between 0 and 20 (units in patches), as well as between 0 and 10. The results from these runs are similar to those of the runs without communication.




Fig. 2. These graphs show the proportion of communicating agents, under different settings. Multiple lines in a single graph reflect multiple runs (see text). The top row shows what happens when the model is initialized with 10% communicating agents. In the middle row there are again two types of agents only now there is no difference between the two types. The bottom row shows the situation without parental investment.

Another metric is the number of agents the environment can support. The more the agents know the more food becomes available to them. It is no surprise then, that as the communicating agents take over the population, knowhow spreads and the total agent population increases (a little). Table 2 shows the average number of agents and average knowhow per agent type (i.e. the number of food stuffs agents know about) for every one of the nine settings. With both communication and parental investment (middle row), there is a clear difference in knowhow between silent agents and agents that communicate. The spread of knowledge depends on the proportion of communicating agents, therefore the average number of agents and their knowledge will vary over time. The average knowhow was measured when half the population consisted of communicating agents.

Table 2. The carrying capacity (N) and spread of knowledge for communicating agents and ‘silent’ agents, for each of the nine model set-ups.

<i>conditions/food</i>	initial values		special food ↑		regular food ↑	
	N	talk silent	N	talk silent	N	talk silent
no communication	316	0.3 0.3	352	0.3 0.3	700	0.3 0.3
with communication	370	1.5 0.7	531	2.0 0.6	745	2.1 1.3
no parental care	360	0.6 0.6	430	0.6 0.6	738	1.2 1.2

The benefit of the parental investment in the form of communication, outweighs any costs that may result from communication towards ‘strangers’. One possible explanation for the low cost would be that the agents form pockets of knowledge and that within these pockets communication is without consequences because all agents know the same things. It is true that only a small percentage (between 5% and 15%) of communications lead to agents acquiring new knowledge. But the effects even these seemingly ineffective communications have, both positive and negative, are apparent. When agents teach their surrounding kin new things the fitness effect is positive. If they share their knowhow with competing strangers, the effect is negative.

We have not fully quantified our results. That is we have not measured exactly how often agents meet their parent (or other related agents) compared to how often they meet other agents. Or how the ratio between interactions with strangers and interactions with kin changes during an agents lifetime. We have not examined which effect of communication dissolves sooner when the agents’ mobility is increased: offspring benefit or increased competition. This is a valid question that has been under debate before (Marshall and Rowe 2003). The answer depends on very specific population properties and cannot be answered using a general model. In order to validate a model of these sorts of details we would need much more complete data on real-world communication. However,

our model as it stands shows the important general characteristic we set out to test, whether communication could in fact be adaptive.

3.4 Conclusions

We found a model that supports honest, cost-free communication about food. The communication is cost-free in the sense that the cumulative effect of sharing knowhow is positive, even in the presence of free-riders who only benefit from the information without sharing their own know-how. Communication benefits the offspring of those that communicate. The agents' low mobility ensures they remain in the vicinity of their offspring for at least some time. Large steps disrupt both the positive (kin-selection) and the negative (increased competition) effects of communication.

One benefit of this freedom to communicate is that the cognitive load of recognizing individual agents and deliberating the possible effects of sharing knowledge with them is taken away. The positive effect of communication among kin will persist in any (agent) society where frequent interactions between parent and offspring are common during 'childhood'.

Our model shows that the positive effect of communication as parental investment is large enough to undo any kind of negative effect from random communication. These findings support the theory of language emergence in the parenting context, as brought forward by Buckley and Steele.

As we stated in our introduction, our results in no way challenge whether there are other selective forces that have affected the evolution of language, particularly language as we know it. But we have conclusively shown that arguments in favor of such mechanisms as selection for prestige cannot rely for evidence on free communication being non-adaptive. Communication may still have been the first selective advantage of language.

References

- Aiello, L. C., Wheeler, P.: The expensive-tissue hypothesis: The brain and the digestive-system in human and primate evolution. *Current Anthropology* **36**(2) (1995) pp.199–221
- Baray, C.: Effects of Population Size upon Emergent Group Behavior. *Complexity International* **6** (1998)
- Buckley, C., Steele, J.: Evolutionary ecology of spoken language: co-evolutionary hypotheses are testable. *World Archaeology* **34**(1) (2002) pp.26–46
- Cavalli-Sforza, L. L., Feldman M. W.: Paradox of the evolution of communication and of social interactivity. *Proceedings of the National Academy of Sciences of the USA* **80**(7) (1983) pp.2017–21
- Cederman L.-E., Rao, M. P.: Exploring the dynamics of the democratic peace. *Journal of Conflict Resolution* **45**(6) (2001) pp.818–833
- Chomsky, N.: Rules and Representations. *Brain and Behavioral Sciences* **3** (1980) pp.1–61
- Cox, S. J., Sluckin, T. J., Steele, J.: Group Size, Memory, and interaction Rate in the Evolution of cooperation. *Current Anthropology* **40**(3) (1999) pp.369–377

- van Dartel, M.F., Postma, E. O., van den Herik, H.J.: Universal properties of adaptive behaviour In Blockeel, H., Denecker, M. eds.: Proceedings of the 14th Belgium-Netherlands Conference on Artificial Intelligence (BNAIC 2002) Leuven, Belgium (2002) pp.59–66
- Dennett, D. C.: Kinds of minds: Towards an understanding of consciousness. Weidenfeld and Nicolson 1996
- Dessalles, J.-L.: Language and hominid politics. In Knight, C., Studdert-Kennedy, M., Hurford, J., eds.: The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form. Cambridge University Press, Cambridge, UK (2000) pp.62–79
- Dunbar, R. I. M.: Coevolution of neocortical size, group size and language in human. Behavioral and Brain Sciences. **16**(4) (1993) pp.681–735
- Epstein, J. M., Axtell, R.: Growing Artificial Societies: Social Science from the Bottom Up. Brookings Institution Press, MIT Press Cambridge, MA (1996)
- Esteve, M., Padget, J., Sierra, C.: Formalizing a language for institutions and norms. In Meyer, J.J., Tambe, M., eds.: Intelligent Agents VIII. Volume 2333 of Lecture Notes in Artificial Intelligence., Springer Verlag (2001) pp.348–366
- Gibson, K. R.: Social transmission of facts and skills in the human species: neural mechanisms. In Box, H. O., Gibson, K. R., eds.: Mammalian Social Learning, Cooperative and Ecological Perspectives Cambridge University Press, Cambridge, UK, (1999) pp.351–366
- Hamilton, W. D.: The Genetical Evolution of Social Behaviour I Journal of Theoretical Biology **7**(1) (1964) pp.1–16
- Hemelrijk, C. K.: Towards the integration of social dominance and spatial structure. Animal Behaviour **59**(5) (2000) pp.1035–48
- Krebs, J. R., Dawkins, R.: Animal Signals: Mind-Reading and Manipulation. In Krebs, J. R., Davies, N. B., eds.: Behavioural ecology: An evolutionary approach (second ed.) Blackwell Scientific Publications, Oxford (1984) pp.380–402
- Lachmann, M., Számadó, S., Bergstrom C. T.: Cost and conflict in animal signals and human language. Proceedings of the National Academy of Sciences of the United States of America **98**(23) (2001) pp.13189–94
- Marshall, J. A. R., Rowe, J. E.: Viscous Populations and Their Support for Reciprocal Cooperation. Artificial Life **9**(3) (2003) pp.327–334
- Maynard-Smith, J., Harper, D.: Animal Signals: Models and Terminology. Journal of Theoretical Biology **177**(1) (1995) pp.305–311
- Noble, J., Paolo, E. A. D., Bullock, S.: Adaptive Factors in the Evolution of Signaling Systems. In Cangelosi, A., Parisi, D., eds.: Simulating the Evolution of Language. London: Springer Verlag, London, UK, (2002) pp.53–78
- Simon, H. A.: A Mechanism for Social Selection and Successful Altruism. Science **250**(4988) (1990) pp. 1665–68
- James Steele, J.: What can archaeology contribute to solving the puzzle of language evolution? plenary talk at The Evolution of Language (April 2004)
- Wilensky, U.: NetLogo Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL (1999)
- Wilensky, U.: The NetLogo Rabbits-Grass-Weeds Model Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL (1998) <http://ccl.northwestern.edu/netlogo/models/RabbitsGrassWeeds>
- Zahavi, A.: Mate selection – A selection for a handicap. Journal of Theoretical Biology **53**(1) (1975) pp.205–214