

Submitted to *Cognitive Science*.

Learning and the Emergence of Coordinated Communication

Michael Oliphant & John Batali

Department of Cognitive Science
University of California at San Diego
9500 Gilman Drive
La Jolla, CA 92093-0515
oliphant@cogsci.ucsd.edu
batali@cogsci.ucsd.edu

DRAFT

November 18, 1996

DRAFT

For the members of a population of animals to enjoy the benefit that might accrue from the exchange of information, their communicative behavior must be coordinated — most of the time that an animal sends a signal in some type of situation, others respond to the signal in a manner appropriate to the situation that inspired it. We investigate how coordinated communication could emerge among animals capable of producing and responding to simple signals, and how such coordination could be maintained, when new members of a population learn to communicate by observing the other members. We describe a learning procedure that enables an individual to achieve the maximum possible accuracy in communicating with a given population. If all new members of the population use this procedure, or one of the approximations to it we describe, the coordination of the population's communication will steadily increase, ultimately yielding a highly coordinated system. Our results are derived mathematically from a formal model of simple communication systems, and are illustrated with computational simulations. We discuss their biological plausibility and their relevance to more complex communication systems, including human language.

1 Introduction

Communication enables animals to influence each others' behavior, often in beneficial ways. This benefit requires that the communicative behavior of the animals be "coordinated," in the sense that most of the time that one animal sends a signal in some type of situation, other animals respond appropriately. Populations of a number of species of primates, other mammals, and birds, have highly coordinated alarm call systems in which the presence of a specific type of predator is indicated with a specific type of vocalization (Struhsaker, 1967; Cheney and Seyfarth, 1990; Slobodchikoff, et al., 1991; Evans, et al., 1993; Hoogland, 1983, 1995; Dasilva, et al., 1994; Blumstein, 1995; Evans and Marler, 1995; Smith, 1996; Hauser, 1996).

Though these alarm call systems appear to be mostly innate, our focus in this paper is on procedures whereby new (e.g., juvenile) members of a population could learn to communicate with the other members by observing their communicative behavior. Two apparently distinct issues are relevant to the evaluation of such learning procedures. First, the procedure must enable the new members to accurately acquire the communication system of the population, even though their observations may be limited, noisy, or otherwise misleading. Second, the learning procedure used by its new members will affect the population's communication system over time. The use of a particular procedure might result in the population's communication increasing in coordination, ultimately yielding a nearly optimally coordinated system. If a learning procedure were to satisfy both criteria, it could explain how learned communication systems are maintained over time, as well as how they are established in the first place.

In section 2 we describe a formal model of simple communication systems that we use to explore these issues. We treat an animal's communicative behavior as the manifestation of a set of behavioral and perceptual dispositions. Upon encountering certain types of situations, the animal may tend to perform specific signaling actions. We refer to this as the animal's "transmission behavior." For each of a set of signaling actions that it can recognize others performing, it may evince some sort of awareness that a specific type of situation is occurrent, perhaps by performing actions appropriate to that type of situation. We refer to this as the animal's "reception behavior." We call such systems "simple" because the signals are produced and responded

to as individual, discrete tokens.

In section 3 we consider procedures for learning to communicate with the members of a population. We show that while procedures based on imitating the transmission and reception behavior observed in the population can accurately acquire an existing highly coordinated communication system, and can maintain it against degradation, such procedures cannot be guaranteed to improve the coordination of a system in which communication is only rarely accurate.

We then derive a learning procedure that yields, for any population, a communication system that has the highest possible communicative accuracy with that population. Unless the population's communication is already optimally coordinated, individuals using the systems produced by our procedure will be able to communicate with the members of the population more accurately than the members can among themselves. The addition of each new member will slightly increase the coordination of the population's communication. Therefore the coordination of the population's communication will steadily increase, ultimately resulting in a highly coordinated system.

We describe computational simulations that illustrate our results in section 4. In section 5 we discuss a number of issues raised by our model and our results. We argue that to satisfy our model, animals must be capable of learning to predict and influence each other's behavior, and they must share an environment in which information exchange is often mutually beneficial. Finally, we explore the relevance of our model to the emergence and learning of human language.

2 Simple Communication Systems

In this section we present our formal model of simple communication systems. A similar model is introduced by Lewis (1969) in his philosophical analysis of convention. Hurford (1989) uses a generalized version of Lewis' model that is essentially identical to the one described here.

2.1 Meanings and Signals

We assume a population of animals that can recognize some set of situation types such that there is a distinct, appropriate response to each. We call

these pairs of situation types and their appropriate responses “meanings.” We also assume that the animals have available to them a set of distinct types of actions that can be performed with little or no cost by an animal, and can be recognized by others. We call the elements of this set “signals.”

Our analysis focuses on communicative episodes of the following sort. One member of a population, upon noticing that a situation of a particular type is happening, produces a signal. We call this “encoding” the meaning. Other animals, upon recognizing the signal, respond to it. We call this “interpreting” the signal as a meaning. The communicative episode is considered successful if the responding animals act appropriately to the situation noticed by the first.

For example, when vervet monkeys (*Cercopithecus aethiops*) see predators, they make vocalizations that alert the rest of the group. The type of alarm call that is given depends on the specific kind of predator in the vicinity. A loud barking call is given for leopards, a short, double syllable cough for eagles, and a “chutter” sound is made for snakes. Though the production of an alarm call is not part of the activity of seeking shelter, and the calls are not similar to sounds the predators make, the response of other monkeys to a given type of call is appropriate to evade the predator that inspired it. When the leopard call is heard, the monkeys run to the trees; the eagle call provokes them to look up into the air and seek shelter; hearing the snake call makes the monkeys stand up on two legs and look in the grass (Struhsaker, 1967; Seyfarth, et al., 1980; Cheney and Seyfarth, 1990).

In describing animals as encoding meanings as signals, and as interpreting signals as meanings, we do *not* thereby assume that the animals are necessarily aware of the communicative nature of their actions, nor that they recognize any relation among situation types, signals, and response types. Both the production of, and the response to, signals will be treated as behavioral dispositions, as described in the next section. The notion of meanings is entirely a formal device, simplifying the model by allowing us to refer to a situation type and its appropriate response type as a single entity.

We assume that there are at least as many available signals as there are meanings. Given that virtually any behavior that an animal can perform could be used as a signal, this seems plausible. Indeed our results hold even if there is an infinite number of signals.¹

¹For the most part our results also hold if the set of meanings is larger than the set of

We assume that each signal can be accurately performed, and its performance accurately recognized, in more than one of the situation types. However we do not assume that there is no relationship at all among some situations and signals, thus we do not assume that signals must be “arbitrary” in the sense of de Saussure (1916) or Hockett (1960a, 1960b).

It may be, for example, that some signaling actions are more likely to be performed in some of the situation types than others, or that a signal is more likely to be interpreted as indicating some particular meanings than others, perhaps due to similarities between the signal and aspects of the situation or its appropriate response. Many animal signals apparently arise from the exaggeration and “ritualization” of intention movements, protective behaviors, autonomic reactions, and other responses to situations (Tinbergen, 1952; Moynihan, 1970; Smith, 1977).

While our model does not require any such relations between signals and meanings, it can apply whether or not they exist. Such clues may speed up the establishment of a coordinated communication system, and may simplify learning the system, but so long as each signal can, in principle, be paired with more than one of the meanings, the population must somehow settle on one of them, and each learner must determine which one it is.

2.2 Send and Receive Functions

We characterize an individual’s communicative behavioral dispositions with two probability functions: a “send function,” and a “receive function.” For each signal and meaning, the send function gives the probability that an animal will send the signal when it notices that the meaning’s situation type is occurrent, and the receive function gives the probability that an animal, upon recognizing the signal, will perform the meaning’s response. Example send and receive functions are shown in figure 1.

An idealization of the communicative behavior of adult vervet monkeys is given by the send function s_1 and the receive function r_1 . The send function s_2 illustrated in figure 1 might describe the communicative behavior of an individual who hasn’t quite mastered the vervets’ system. It is reasonably accurate for leopards, in that it sends the bark signal with probability 0.7, but it also sends a cough with probability 0.1 and a chatter with probability

signals, with differences that will be noted when appropriate.

s_1	<i>bark</i>	<i>cough</i>	<i>chutter</i>
<i>leopard</i>	1.0	0.0	0.0
<i>eagle</i>	0.0	1.0	0.0
<i>snake</i>	0.0	0.0	1.0

<i>bark</i>	<i>cough</i>	<i>chutter</i>	r_1
1.0	0.0	0.0	<i>leopard</i>
0.0	1.0	0.0	<i>eagle</i>
0.0	0.0	1.0	<i>snake</i>

s_2	<i>bark</i>	<i>cough</i>	<i>chutter</i>
<i>leopard</i>	0.7	0.1	0.2
<i>eagle</i>	0.4	0.5	0.1
<i>snake</i>	0.0	0.4	0.6

<i>bark</i>	<i>cough</i>	<i>chutter</i>	r_2
0.6	0.4	0.2	<i>leopard</i>
0.2	0.5	0.2	<i>eagle</i>
0.2	0.1	0.6	<i>snake</i>

Figure 1: Example send and receive functions. In the send functions s_1 and s_2 , each entry is the probability that the signal at the top of the column is sent for the presence of the predator to the left of the row. In the receive functions r_1 and r_2 , each entry is the probability that signal at the top of the column is responded to in a way appropriate to evade the predator to the right of the row.

0.2. For eagles, s_2 is even less like the vervet's system, sending the cough signal only slightly more often than it sends the bark signal. The receive function r_2 is also similar to, but not exactly, that of the vervets.

In the definitions and calculations below, we represent the set of meanings relevant to a population of animals as M , and the set of signals available to them as S . The probability that a signal σ is sent for a meaning μ by a send function s is represented as $s(\mu, \sigma)$. We represent the probability that a signal σ is interpreted as a meaning μ by receive function r as $r(\sigma, \mu)$.

2.3 Communicative Accuracy

Suppose that one animal possesses the function s_1 and another possesses the receive function r_1 shown in figure 1. If the first animal notices a leopard, it will, with probability 1.0, produce the cough signal. If the second animal hears the cough, it will, with probability 1.0, react in a manner appropriate to the presence of leopard. Similar outcomes would obtain if the first notices an eagle or a snake. All signals sent by the first animal will be correctly interpreted by the second.

Now suppose that an animal whose send function is s_1 sends a signal to an animal whose receive function is r_2 when a leopard is present. The first animal will produce a bark with probability 1.0. The second animal, upon hearing the bark, will respond as if a leopard were present with probability 0.6. But it also may respond as if an eagle or a snake were present, each with probability 0.2.

These considerations motivate the following definition: Given a send function s , and a receive function r , we define the “communicative accuracy” from s to r , which we write as $ca(s, r)$, as the probability that signals sent by an individual using the send function s will be correctly interpreted by an individual using the receive function r .

To compute this value, first consider a specific meaning μ and a specific signal σ . The probability that an animal using send function s will send σ for μ is $s(\mu, \sigma)$, and the probability that an animal using receive function r will interpret σ as μ is $r(\sigma, \mu)$. These two events are independent, so the probability that both occur is given by the product:

$$s(\mu, \sigma) r(\sigma, \mu) \tag{1}$$

For a given meaning μ , the probability that a signal sent by an individual according to send function s will be correctly interpreted by an individual possessing receive function r can be computed by taking the sum, over all of the possible signals, of the probability that, if a given signal is sent for μ , it will be correctly interpreted:

$$\sum_{\sigma} s(\mu, \sigma) r(\sigma, \mu) \tag{2}$$

We assume that each of the situation types corresponding to the meanings occurs with equal frequency. This allows us to compute the communicative

accuracy from s to r by taking the average probability that signals for each meaning are interpreted correctly. This is the average of the quantity in formula 2, taken over all of the meanings:

$$ca(s, r) = \frac{1}{|M|} \sum_{\mu} \sum_{\sigma} s(\mu, \sigma) r(\sigma, \mu) \quad (3)$$

where $|M|$ is the number of meanings. This probability has a maximum value of 1.0 if the set of signals is at least as large as the set of meanings. If the number of signals is smaller than the number of meanings, the maximum value of $ca(s, r)$ is $|S|/|M|$, where $|S|$ is the number of signals.

Here are the communicative accuracy values for the send and receive functions in figure 1:

$$\begin{array}{ll} ca(s_1, r_1) = 1.00 & ca(s_2, r_2) = 0.42 \\ ca(s_1, r_2) = 0.57 & ca(s_2, r_1) = 0.60 \end{array}$$

The above measure of communicative accuracy only takes into account one direction of information transfer. We define the “two-way communicative accuracy” (ca_2) between two individuals as the probability that a signal sent by either member of the pair is interpreted correctly by the other. If the first individual possesses send and receive functions s_1 and r_1 , and the other individual has send and receive functions s_2 and r_2 , and each is as likely to send as to receive, the two-way communicative accuracy between them can be computed as follows:

$$ca_2(s_1, r_1, s_2, r_2) = \frac{1}{2} (ca(s_1, r_2) + ca(s_2, r_1)) \quad (4)$$

2.4 Communicating with a Population

We now consider an animal that possesses a specific send and receive function interacting with the members of a population, each of which possess their own send and receive functions. Given the send and receive functions of the individual, and those possessed by the members of the population, it is possible to determine the probability that communicative events involving the individual and members of the population will be successful.

If we assume that the individual interacts with members of the population with equal likelihood, the probability that a signal sent by the individual

to members of the population can be determined by taking the average, over each member of the population, of the communicative accuracy between the individual's send function, and the receive functions of the members of the population. A similar calculation could be performed to compute the probability that the individual will correctly interpret signals sent by members of the population.

A mathematically equivalent way to compute these probabilities is to first determine, for each meaning and signal, the average probability that the signal is sent for the meaning, and the average probability that the signal is interpreted as the meaning, by members of the population. These probabilities can be found by taking the average of the entries in the send and receive functions of the members in the population, as follows:

$$S(\mu, \sigma) = \frac{1}{N} \sum_i s_i(\mu, \sigma) \quad R(\sigma, \mu) = \frac{1}{N} \sum_i r_i(\sigma, \mu) \quad (5)$$

where N is the number of individuals in the population, and s_i and r_i refer to the send and receive functions, respectively, of the i th member of the population.

If we assume that each pair of animals interacts with equal likelihood, the functions S and R can themselves be treated as send and receive functions, describing the average communicative behavior of the members of the population. In a sense, we are taking the whole population as if it were an individual that can communicate according to the send function S and the receive function R .

Thus the average probability that an individual with send function s will be correctly interpreted by a member of a population whose average receive function is R is $ca(s, R)$. The average probability that an individual with receive function r will correctly interpret signals sent by a member of the population whose average send function is S is $ca(S, r)$. The two-way communicative accuracy between the individual and the members of the population is $ca_2(s, r, S, R)$.

The accuracy of intercommunication among the members of a population can now be computed. Since the probability that a given signal σ is sent for meaning μ by members the population is $S(\mu, \sigma)$, and the probability that signal σ is interpreted as μ is $R(\sigma, \mu)$, the value of $ca(S, R)$ is the probability that communicative events involving members of the population

will be successful.²

The maximum possible communicative accuracy among the members of a population is equal to 1.0 if the set of signals is at least as large as the set of meanings, and equals $|S|/|M|$ otherwise. If $ca(S, R) = 1.0$ for some population, we say that its communication system is “optimally coordinated.”

2.5 Game-Theoretic and Evolutionary Accounts of the Emergence of Coordinated Communication

Lewis (1969), explored whether a coordinated communication system could emerge and remain stable in a game among players whose benefit depends on the accuracy of their communication with the other players. Further development and analysis of Lewis’ model has been done using the mathematical theory of games (Spence, 1973; Crawford and Sobel, 1982; Cho and Kreps, 1987; Farrell, 1988, 1993). In the terminology of game theory, our model is an instance of what are called “signaling games,” specifically those involving “cheap talk,” in which the benefit (or cost) to the players depends only on whether their communication is successful.

An important result from this approach is that while highly coordinated communication systems can occur (i.e., are Nash equilibria in such games), they are not inevitable. Other stable outcomes are “babbling” or “pooling” equilibria in which the participants send the same signal for several meanings, or interpret different signals as the same meaning, and thereby fail to transmit much information. Furthermore, in these models the players choose their communicative behavior based on explicit analyses of the alternatives, and with knowledge of the communication systems used by the other players. It would seem that a model in which coordinated systems are more likely to occur, and which requires less cognitive sophistication, is needed to account for the emergence of coordinated communications among animals.

One possibility is that the communicative dispositions of a species are innate and heritable, and the benefits of successful communication are sufficient for coordinated communication systems to evolve by natural selection.

²This is strictly correct only if we allow for an animal interpreting its own signals to count as a “communicative event.” Whatever the plausibility of this assumption, its influence is negligible if the population size is large, and the results below would be identical, though their derivations would be slightly more tedious.

Maynard-Smith (1965) shows that the behavior of making an alarm call can benefit an individual's inclusive fitness, and is therefore possible to acquire through natural selection. It is also possible that the sender of a signal obtains a more immediate benefit from its being correctly interpreted. For example, if by signaling the presence of a predator, the other animals respond by hiding, the predator might choose to leave the area. Another example is the proposal and acceptance of an offer to form an alliance and thereby achieve a mutually beneficial outcome that neither participant could produce alone (Hinde, 1981; Hoogland, 1983).

When the game-theoretic models described above are used to analyze populations of agents whose communicative behaviors are innate and fixed during their lifetimes, but whose reproductive fitness depends on the accuracy of their communication, it can be shown that an optimally coordinated communication system is the only evolutionarily stable outcome (Warneryrd, 1993; Blume, et al., 1993; Kim and Sobel, 1995; Skyrms, 1996). Computational simulations of the evolution of innate communication systems illustrate these results (Werner and Dyer, 1991; MacLennan and Burghardt, 1993; Ackley and Littman, 1994; Oliphant, 1996; Parisi, forthcoming; Di Paolo, 1996).

3 Learning Simple Communication Systems

Our concern in this paper is with the question of how individuals could *learn* to communicate, and the effect that learning has on the coordination of a population's communication system.

3.1 Learning Procedures

We assume that a new member learns to communicate by observing communicative events among members of that population for a while, and then, based on those observations, constructs its own send and receive functions, according to some "learning procedure."

The learner must be able, at least in some communicative episodes, to tell which meaning is being encoded by a signal, and how the receiver interprets it. In the extreme, of course, this assumption invalidates the need for communication — such a learner could just continue reading the minds of others rather than figuring out what their signals mean. But it seems

plausible that for highly social animals, for whom accurate communicative ability is an essential skill to acquire, some deference to learners would be made. This could take the form of an explicit indication of the correct response to a signal when it is produced, or of the correct signal to perform in a given type of situation. Also, in many situations where a signal is produced or interpreted, its significance is immediately apparent to observers in the immediate vicinity. By recording such observations, the learner might ultimately be able to accurately send and interpret signals when it is more distant from the situation or from other individuals, or they are otherwise occluded.

The introduction of new members to a population, whose communicative behaviors are based on observations of the existing members, can have profound effects on how the population's communication system changes over time. Suppose that a new individual, whose learned send and receive functions are s and r , is added to a population whose average send and receive functions are initially S and R . If the population size is finite, and if:

$$ca_2(s, r, S, R) \geq ca(S, R) \quad (6)$$

the communicative accuracy may increase slightly as a result of the addition of the new member. If all new members use a learning procedure that strictly satisfies this inequality, the communicative accuracy will increase as a result of adding the new members, and may ultimately approach a state of optimal coordination.

The assumption that new members of the population do all their learning before being added to the population, and use a fixed communication system thereafter, is made to simplify the mathematical analysis. If juveniles were added to the population before they begin learning, their attempts at communication would introduce noise into the observations made by other learners. If the fraction of such untrained members were low enough, this noise would not affect the emergence and maintenance of highly coordinated communication systems, though it would limit the maximum achievable communicative accuracy in such populations to a value less than 1.0.

3.2 Related Work

Hurford (1989) explores the evolutionary advantages of alternative learning procedures for simple communication systems. The learning procedures dis-

S_1	<i>bark</i>	<i>cough</i>	<i>chutter</i>	
<i>leopard</i>	0.500	0.100	0.400	
<i>eagle</i>	0.425	0.450	0.125	
<i>snake</i>	0.100	0.475	0.425	
	<i>bark</i>	<i>cough</i>	<i>chutter</i>	R_1
	0.400	0.525	0.325	<i>leopard</i>
	0.375	0.425	0.250	<i>eagle</i>
	0.225	0.050	0.425	<i>snake</i>

Figure 2: Send and receive functions describing the average communicative behavior of a hypothetical population. In the send function S_1 each entry gives the average probability that a member will send the indicated signal for the meaning. The receive function R_1 gives the average probability that members of the population will interpret a given signal as a given meaning.

cussed below are based on considerations of Hurford’s results. Hutchins and Hazlehurst (1991) present a somewhat different model, in which groups of neural networks develop a communication system that involves shared patterns of activation values on the networks’ hidden layers as they are trained. Canning (1992), applying game-theoretic models, shows that coordinated communication can occur as a result of learning, provided that each agent utilizes a set of assumptions about the other agents’ systems. Yanko and Stein (1993) show how a simple communication system can develop among a group of mobile robots trained with reinforcement learning as they perform a cooperative task. Steels (1996) applies symbolic learning methods to explore the development and subsequent modification of communication systems. While the results from these projects are consistent with ours (Oliphant, forthcoming), the model of communication and learning we employ is meant to more abstractly characterize the cognitive and other biological requirements for the emergence of coordinated communication.

3.3 Learning by Imitation

Consider a population whose average send and receive functions are as shown in figure 2. What send and receive functions would enable a learner to most accurately communicate with this population?

The simplest way to use the average send and receive functions would

For each meaning μ :

s.1: Find the signal κ_μ for which $S(\mu, \kappa_\mu)$ is maximum.

s.2: Set $s_{ic}(\mu, \kappa_\mu) = 1.0$, and set $s_{ic}(\mu, \sigma) = 0$ for all $\sigma \neq \kappa_\mu$.

For each signal σ :

r.1: Find the meaning η_σ for which $R(\sigma, \eta_\sigma)$ is maximum.

r.2: Set $r_{ic}(\sigma, \eta_\sigma) = 1.0$, and set $r_{ic}(\sigma, \mu) = 0$ for all $\mu \neq \eta_\sigma$.

Figure 3: The Imitate-Choose learning procedure.

s_{ic}	<i>bark</i>	<i>cough</i>	<i>chutter</i>
<i>leopard</i>	1.0	0.0	0.0
<i>eagle</i>	0.0	1.0	0.0
<i>snake</i>	0.0	1.0	0.0

<i>bark</i>	<i>cough</i>	<i>chutter</i>	r_{ic}
1.0	1.0	0.0	<i>leopard</i>
0.0	0.0	0.0	<i>eagle</i>
0.0	0.0	1.0	<i>snake</i>

Figure 4: Send and receive functions derived using the Imitate-Choose learning procedure, for a population whose average communicative behavior is described by the send and receive functions in figure 2.

be to imitate them — to use S_1 as the learner’s send function, and to use R_1 as the learner’s receive function. Such a learner’s two-way communicative accuracy with the population will be exactly the same as the average communicative accuracy among its members. This learning procedure could acquire an optimal system if it were in place, but could not improve a sub-optimal system, not even one that was only slightly degraded from optimal.

Instead of simply imitating observed communicative behavior, a learner could use a population’s average send and receive functions to determine which send and receive behaviors are most popular. For each signal and meaning pair, the learner would send the signal most often sent for that meaning, and would interpret each signal the way most of the population does. This learning procedure, which we call “Imitate-Choose,” is described in figure 3.³ If a population’s communication is optimally coordinated, this procedure will accurately acquire the system. When applied to the average send and receive functions shown in figure 2, the Imitate-Choose learning procedure will yield the new send and receive functions shown in figure 4.

An individual using the new receive function r_{ic} will correctly interpret signals sent by a population whose average send and receive functions are shown in figure 2 with a probability equal to $ca(S_1, r_{ic}) = 0.342$. The average communicative accuracy for the population is 0.33, so the Imitate-Choose procedure can sometimes yield individuals that are slightly better than average at interpreting signals sent by members of a population whose communication is less than optimally coordinated.

The send function obtained by the Imitate-Choose procedure does not do as well. Signals sent according to s_{ic} will be interpreted correctly by the population with a probability of $ca(s_{ic}, R_1) = 0.292$. This is less than the population average. These two values can be used to compute the two-way communicative accuracy between an individual using the new system and the population: $ca_2(s_{ic}, r_{ic}, S_1, R_1) = 0.317$. So adding an individual using this system to the population will decrease the average communicative accuracy.

In general, Imitate-Choose exaggerates the communicative dispositions in the population. If the system is highly coordinated, the use of Imitate-Choose

³The descriptions of the learning procedures here and in the next section leave out a few details having to do with cases where values in the population average send and receive functions are zero, or where there is no unique maximum. These cases don’t affect the results we derive. In appendix A.1 we show how they can be handled for the learning procedure described in the next section.

can increase coordination, thus maintaining a near optimal system against degradation. If the population’s communication is not very well coordinated, on the other hand, Imitate-Choose can degrade it even more.

3.4 Learning by Obverting

Consider again the average send and receive functions shown in figure 2. While the bark signal is sent for leopards most often, it is only *interpreted* as indicating the presence of leopards with a probability of 0.400. If the cough signal were sent to indicate leopards, it would be correctly interpreted with a probability of 0.525. Similar considerations hold for interpreting signals. As just noted, most of the population interprets the cough signal as indicating a leopard, while this signal is most often sent when a snake is actually present.

For one’s signals to be interpreted correctly, one should send for each meaning the signal that is most likely to be interpreted as that meaning. To maximize the probability that one will correctly interpret signals sent by others, one should interpret a signal as the meaning it most often encodes.

A learning procedure that works this way, which we call “Obverter,” is described in figure 5. The send and receive functions that result from applying Obverter to a population with the average send and receive functions in figure 2 are shown in figure 6. They have the following probabilities of accurate communication with a population whose average send and receive functions are shown in figure 2:

$$ca(s_{ob}, R_1) = 0.46 \quad ca(S_1, r_{ob}) = 0.47 \quad ca_2(s_{ob}, r_{ob}, S_1, R_1) = 0.46$$

All of these values are greater than the population average of 0.33.

It is proved in appendix A.1 that the Obverter learning procedure, when applied to any population’s average send and receive functions, will produce send and receive functions that have the highest possible communicative accuracy with that population. An individual using the Obverter learning procedure will thus acquire a send and a receive function whose two-way communicative accuracy with the population will be strictly greater than the average communicative accuracy, unless the population already possesses an optimally coordinated system. Therefore the average communicative accuracy of a population whose new members use the Obverter learning procedure will steadily increase, and an optimally coordinated system will ultimately emerge.

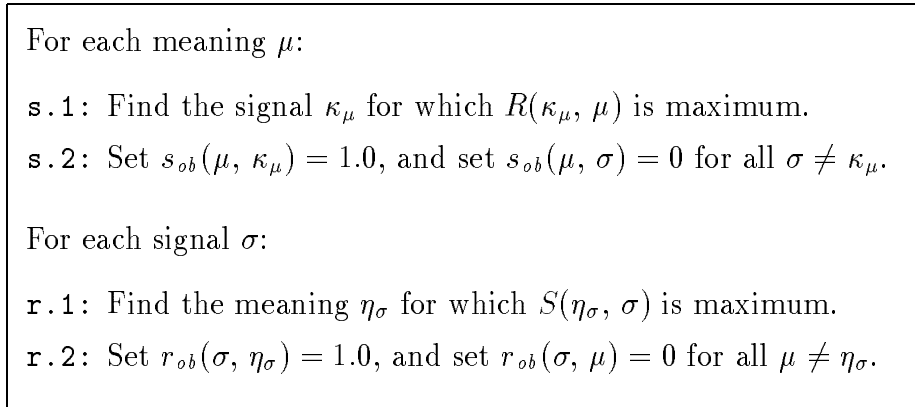


Figure 5: The Obverter learning procedure.

s_{ob}	<i>bark</i>	<i>cough</i>	<i>chutter</i>
<i>leopard</i>	0.0	1.0	0.0
<i>eagle</i>	0.0	1.0	0.0
<i>snake</i>	0.0	0.0	1.0

<i>bark</i>	<i>cough</i>	<i>chutter</i>	r_{ob}
1.0	0.0	0.0	<i>leopard</i>
0.0	0.0	0.0	<i>eagle</i>
0.0	1.0	1.0	<i>snake</i>

Figure 6: Send and receive functions derived using the Obverter learning procedure, for a population whose average communicative behavior is described by the send and receive functions in figure 2.

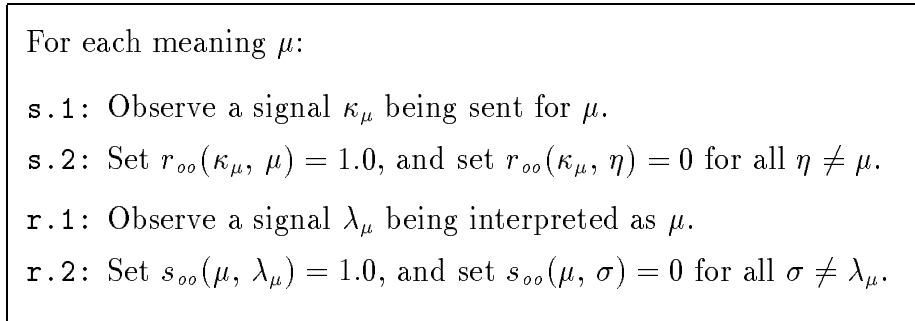


Figure 7: The Unit-Statistic learning procedure.

3.5 Approximating Obverter

To use the Obverter procedure, a learner must have access to the population average send and receive functions. It is more plausible to assume that the learner has access only to approximations of these functions, based on a finite number of observations.

Let us suppose that a learner performs a set of observations, and collects two arrays of statistics, S_{obs} and R_{obs} , based on those observations. These arrays can be used in place of S and R in the Obverter procedure. Depending on the number of observations, S_{obs} and R_{obs} might accurately approximate S and R , and the resultant send and receive functions will be close to optimal for that population.

With unbiased sampling, errors introduced by the finite sample size (and other random noise) will cancel out, and so the approximations to the Obverter procedure that use a finite number of observations can yield, and maintain, a highly coordinated communication system. It is shown in appendix A.3, however, that a learning procedure that uses a finite number of observations cannot be guaranteed to observe each of the meanings being conveyed. This means that a population using such a learning procedure cannot quite reach a state of optimal coordination, though it can come arbitrarily close as the number of observations made by each learner increases.

3.6 Unit-Statistic Learning

The specific set of observations required to guarantee that the average communicative accuracy will increase, such that the population may ultimately attain an optimally coordinated system, turns out to be very small. The “Unit-Statistic” learning procedure, presented in figure 7, requires only a single observation, for each meaning, of a signal being sent for that meaning, and of some signal being interpreted as that meaning. These single observations could be the most recent, the first observed, the ones the learner is most confident about, or could be chosen at random.

The Unit-Statistic procedure is essentially the same as Obverter, except that it uses the single observations of each meaning being sent, instead of the population average send function, to create the new receive function, and it uses the single observations of each meaning being interpreted, instead of the population average receive function, to create the new send function. The Unit-Statistic procedure thus makes extremely small demands on the learner’s memory and other cognitive faculties, since it only needs to record one observation of each meaning being sent and received, and does not need to compute averages or determine maximal values.

It is shown in appendix A.2 that, in a population whose new members use this learning procedure, the expected value of the communicative accuracy will increase over time. The general intuition behind the proof is that the more often a given signal is sent for a given meaning, the more likely it is to be observed, and therefore to be interpreted as that meaning by a receive function created by the Unit-Statistic procedure. The learner’s ability to correctly interpret signals will tend to be greater than average. A similar argument holds for learner’s send function.

Thus, if its new members use the Unit-Statistic learning procedure the population’s average communicative accuracy will increase (though more slowly than with the Obverter procedure), and can ultimately achieve and maintain an optimally coordinated system.

4 Computational Simulations

Computational simulations were performed to illustrate these results. In each simulation, a population of individuals is created, each of which possesses

<p>Obverter: The Obverter learning procedure.</p> <p>Obs-25: An approximation to the Obverter learning procedure using observations of 25 episodes of communication.</p> <p>Obs-10: As above, but with 10 observations.</p> <p>Unit-Stat: The Unit-Statistic learning procedure.</p>
--

Table 1: Learning procedures used in the computational simulations.

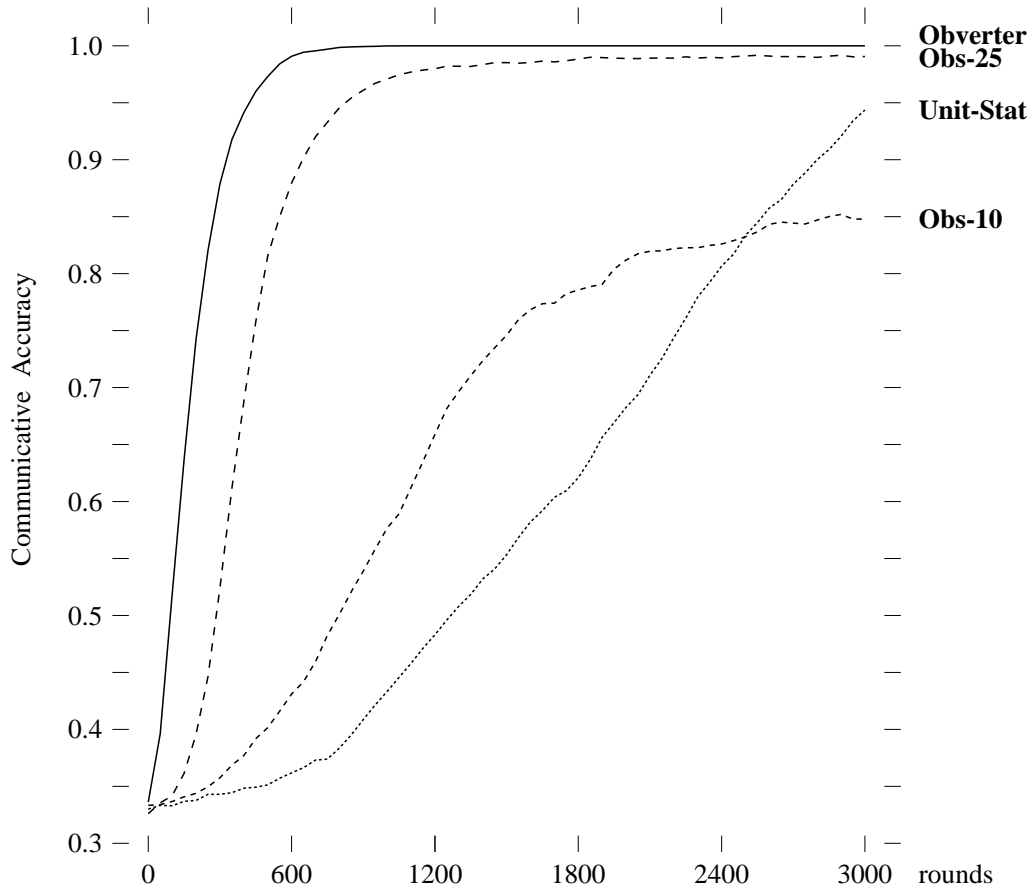
send and receive functions implemented as arrays of values representing the relevant probabilities. At the start of a simulation, the contents of each individual's send and receive arrays are initialized to random values.

In each round of a simulation run, one individual is chosen at random and removed from the population. A new individual is created, and uses one of the learning procedures described in table 1 to determine values for its send and receive arrays. In a given simulation, every individual uses the same learning procedure. After the new individual is trained, it is added to the population, whose average communicative accuracy is then recorded.

Results of such simulations are shown in figure 8. Each simulation involves 100 individuals capable of sending five signals for three meanings. The plots in figure 8 show, for each of the learning procedures, the communicative accuracy of the populations, averaged over ten simulation runs.

Populations using the Obverter learning procedure do quite well. Their average communicative accuracy reaches 0.99 after only 600 rounds of the simulation (when each individual in the population has been replaced an average of six times). The Obs-25 procedure does almost as well, reaching an accuracy of 0.98 after 1200 rounds. The Obs-10 procedure is less effective, but still reaches a high value of communicative accuracy. The Unit-Statistic learning procedure also performs well. Details of how the learning procedures affect the temporal dynamics of a population's average communicative accuracy are discussed in appendix A.3.

Figure 9 shows send and receive functions produced during a simulation of the Obs-25 learning procedure. In round 0, the send and receive func-



<i>learning procedure</i>		<i>Round</i>				
		600	1200	1800	2400	3000
Obverter	<i>mean</i>	0.991	1.000	1.000	1.000	1.000
	<i>sd</i>	0.005	0.000	0.000	0.000	0.000
Obs-25	<i>mean</i>	0.879	0.980	0.988	0.989	0.991
	<i>sd</i>	0.035	0.019	0.010	0.012	0.012
Obs-10	<i>mean</i>	0.431	0.659	0.786	0.826	0.848
	<i>sd</i>	0.031	0.076	0.098	0.032	0.030
Unit-Stat	<i>mean</i>	0.362	0.483	0.621	0.807	0.944
	<i>sd</i>	0.036	0.087	0.125	0.156	0.078

Figure 8: Simulation runs of a populations learning to communicate using different learning procedures. Each plot shows the average communicative accuracy of the population during the run. Results shown are the average of ten runs; statistical properties of the runs are given in the table.

S	a	b	c	d	e	Round 0	s	a	b	c	d	e	
1	.21	.18	.20	.20	.21	$ca = 0.33$	1	.00	.00	.00	.00	1.0	
2	.19	.21	.22	.19	.19		2	.00	.00	1.0	.00	.00	
3	.20	.20	.21	.19	.21		3	.00	1.0	.00	.00	.00	
	a	b	c	d	e	R		a	b	c	d	e	r
	.32	.32	.35	.32	.37	1		.00	1.0	1.0	.00	.00	1
	.35	.32	.34	.32	.31	2		.00	.00	.00	1.0	1.0	2
	.33	.35	.31	.35	.32	3		1.0	.00	.00	.00	.00	3

S	a	b	c	d	e	Round 200	s	a	b	c	d	e	
1	.39	.15	.09	.12	.25	$ca = 0.40$	1	1.0	.00	.00	.00	.00	
2	.09	.13	.59	.13	.06		2	.00	.00	1.0	.00	.00	
3	.14	.24	.13	.27	.23		3	.00	.00	.00	.00	1.0	
	a	b	c	d	e	R		a	b	c	d	e	r
	.50	.30	.10	.28	.35	1		1.0	1.0	1.0	.00	.00	1
	.23	.28	.78	.28	.17	2		.00	.00	.00	.00	.00	2
	.27	.43	.12	.44	.48	3		.00	.00	.00	1.0	1.0	3

S	a	b	c	d	e	Round 400	s	a	b	c	d	e	
1	.73	.13	.01	.03	.09	$ca = 0.69$	1	1.0	.00	.00	.00	.00	
2	.00	.03	.93	.02	.01		2	.00	.00	1.0	.00	.00	
3	.04	.20	.02	.46	.27		3	.00	.00	.00	.00	1.0	
	a	b	c	d	e	R		a	b	c	d	e	r
	.88	.37	.01	.18	.36	1		1.0	.00	.00	.00	.00	1
	.05	.11	.95	.06	.04	2		.00	1.0	1.0	.00	.00	2
	.06	.51	.03	.76	.60	3		.00	.00	.00	1.0	1.0	3

S	a	b	c	d	e	Round 600	s	a	b	c	d	e	
1	.91	.05	.00	.00	.04	$ca = 0.88$	1	1.0	.00	.00	.00	.00	
2	.00	.00	1.0	.00	.00		2	.00	.00	1.0	.00	.00	
3	.00	.27	.00	.51	.22		3	.00	.00	.00	.00	1.0	
	a	b	c	d	e	R		a	b	c	d	e	r
	.97	.14	.00	.01	.13	1		1.0	.00	.00	.00	.00	1
	.01	.05	1.0	.00	.03	2		.00	.00	1.0	.00	.00	2
	.02	.81	.00	.99	.84	3		.00	1.0	.00	1.0	1.0	3

Figure 9: Send and receive functions, and the average communicative accuracy (ca), during one of the simulation runs shown in figure 8. Three meanings (labeled ‘1’, ‘2’, and ‘3’), and five signals (labeled ‘a’, ‘b’, ‘c’, ‘d’, and ‘e’) are used. The left column presents the population average send and receive functions at each of the indicated rounds. The right column shows outputs of the learning procedure Obs-25 when applied during that round.

tions of each member of the population are random, as reflected in the fact that most of the entries in the average send and receive functions for that round are approximately equal, and the average communicative accuracy is 0.33. However the learned send and receive functions still incorporate some appropriate mappings. For example the signal ‘e’ is interpreted most often as meaning ‘1’, and the learner’s send function incorporates this mapping. On the other hand, signal ‘a’ is interpreted most often as meaning ‘2’ but the learner sends signal ‘c’ for this meaning. But the difference between the two probabilities is small (0.35 versus 0.34) and so is not always picked up in only 25 samples.

As the simulation progresses, the population begins to settle on a more accurate system, and the learner is able to acquire it accurately, and often improve on it. By round 600, when the average communicative accuracy is near 0.9, the signal ‘a’ is almost always sent for meaning ‘1’, and is almost always interpreted correctly. For meaning ‘2’ the signal ‘c’ is always sent and received accurately. The population uses the signals ‘b’, ‘d’, and ‘e’ to encode meaning ‘3’. Even so, the average communication accuracy is high, because all three of these signals are interpreted correctly as meaning ‘3’.

We have run many simulations of this model, varying the population sizes, numbers of signals and meanings, learning procedures, and other aspects of the model. In all of our simulations, approximations to the Obverter procedure yield near optimal communication systems.

5 Discussion

5.1 Why Obverter Works

If the point is for one’s signals to be understood, the transmission behavior of others is irrelevant. What is important is their reception behavior. The transmission behavior of others should be consulted, on the other hand, when one is determining how to interpret their signals. Procedures based on imitation, or on other non-communicative considerations, for example the ability of the learner to accurately interpret its own signals, can not be guaranteed to increase communicative accuracy.

Another reason why Obverter performs optimally is because it explicitly avoids ambiguities. When choosing a signal to transmit for a meaning, or

a meaning to interpret in response to a signal, the procedure will select the single most popular. If each new member does this, the ambiguities in the population's system will be shaken out over time. Such ambiguity-avoidance alone is not sufficient for a procedure to improve communicative accuracy however, as is shown by the inferiority of the performance of Imitate-Choose in section 3.3, even though that procedure also involves choosing the most popular alternative.

The Obverter procedure combines the properties of determining communicative behavior according to its communicative function, and of avoiding ambiguity in the resultant send and receive functions. These two properties not only guarantee that its use will improve the communicative accuracy of a population, but that each learner will acquire the optimal system for communicating with the population.

5.2 Cognitive Requirements

To instantiate our model, animals must be capable of classifying different types of situations, and must be capable of recognizing that a situation of a given type is occurrent. They must be capable of performing the correct kinds of responses to the situations they recognize. They must also be capable of producing several different types of signaling actions in the situations, and of recognizing when other animals are performing them. All of these abilities must be at least partly learned. These general requirements are not very stringent, and they are satisfied by many mammal species, in particular primates and rodents (Barnett, 1975; Cheney and Seyfarth, 1990; Hoogland, 1995), probably by many species of birds (Burton, 1985), and possibly by species in other vertebrate taxa (Pough, et al., 1996).

Our model imposes two further requirements that are satisfied in fewer species. The first such requirement is not cognitive but *ecological* — the animals must inhabit an environment in which the accurate exchange of information is mutually beneficial, at least in many situations. The second additional requirement our model imposes is that the animals be capable of understanding and predicting each other's behavior well enough to use some behaviors as signals, and to be capable of reliably affecting the behavior of others by sending signals. Both of these requirements are most likely to be met by animals that live socially (Humphrey, 1976; Rubenstein and Wrangham, 1986; Heyes and Galef, 1996).

For animals that satisfy all of these requirements, the Obverter learning procedure is an instance of Bayesian inference. The construction of an individual's receive map is based on determining the most likely cause of each signal (the situation type in which it is most often performed). An individual's send map is based on determining, for each potential desired effect, the most likely way to cause it (by sending the signal which will most likely be responded to in a way that will benefit the sender). Thus, in environments in which accurate information exchange is often beneficial, the Obverter procedure is an instance of a generally useful cognitive ability applied to the domain of social behavior.

The requirements imposed by our model suggest why highly coordinated learned communication systems seem to be rare. While many animal species are capable of learning to exploit regularities in their environments, our model requires in addition that the animals learn regularities in each other's behavior. It is possible that very few animal species are capable of doing this well enough to instantiate our model. For many animals species whose behavior is largely innate anyway, a more direct evolutionary pathway to obtain the benefit of accurate communication would have been to acquire innate transmission and reception behaviors.

On the other hand, the cognitive and ecological requirements our model imposes, and the coordinated communication that results, need not emerge independently. As the social interactions among the members of a population become more complex, it is more and more likely the animals will encounter situations in which coordinated activity is of mutual benefit. To exploit such situations, the animals might develop better abilities to predict and influence each others' behavior, thus enabling even more complex social interactions. Further development of the abilities to coordinate social activity might require learning, especially if the animals must rapidly adapt their behavior to new environments. Thus sociality, learning, and communication might emerge synergistically, each ability drawing on, and contributing to, the development of the others.

In section 2.1, we pointed out that our model does not require that the animals be aware of the communicative nature of their actions. In particular, we do not require that an animal entertain a "communicative intent" when sending signals, nor that it recognize any such intent on the part of those whose signals it observes, as suggested by Grice (1957, 1989), whose ac-

count has had a profound influence on the debate about the nature of animal communication (Bennett, 1976; Dennett, 1983; Premack, 1983; Cheney and Seyfarth, 1990, chapter 5).

While often taken as a general theory of communicative action, Grice's account is primarily intended to explain how a speaker could use an utterance to convey something *other* than its conventional meaning. Our results demonstrate that communicative conventions can arise among animals with less cognitive sophistication than Grice requires. Our model does not, however, explain how animals might convey anything other than the conventional meanings of their signals. Given that animals apparently do not often resort to sarcasm, irony, or theatrical productions — cases central in the analyses of Grice, Austin (1962) and Searle (1969) — this does not seem like a fatal shortcoming.

Animals are sometimes observed to use their communicative abilities to *deceive* one another (Andersson, 1980; Krebs and Dawkins, 1984; Munn, 1986; Byrne and Whiten, 1988). While the ability to do this is also taken as evidence for the possession of, and the recognition of, higher-order intentionality by animals (Whiten, 1991), our model is consistent with the occasional occurrence of deception among animals without such cognitive abilities. It is possible that situations could occur in which one of the participants in a communicative event would benefit from producing a signal that is ordinarily sent in a different situation type. To take advantage of such situations, an animal need only be aware of how other individuals respond to signals that it sends. It is also possible for animals to *resist* deception if they can somehow determine that despite the signal performed by another animal, the situation type ordinarily conveyed by that signal could not be occurrent. These abilities require that the animals be capable of subtle reasoning about each other's behaviors, but not the internal cognitive states that may be causing them (Maynard-Smith and Price, 1973; Johnson, 1993). Of course for a communication system to become established in the first place, and coordinated enough to allow for the possibility of deception, situations in which successful deception is possible must be relatively rare.

5.3 Relevance to Human Language

As the story of Adam making up names for animals (Genesis 2: 19–20) illustrates, it has long been understood that the specific sounds or gestures as-

sociated with a given concept are matters of more or less arbitrary convention (de Saussure, 1916), but little has been said about how the conventions that constitute a language's lexicon could arise. Lewis' (1970) account of convention involves participants explicitly reasoning about the potential payoffs of their communicative actions and of those of the others. Aristotle (see Sorabji, 1996) argues that the establishment of conventions requires language (*logos*) in the first place. Strawson (1974), Lewis (1975), and Davidson (1984) argue that the establishment of conventions requires the ability to form complex intentions and to recognize them in others. As mentioned above, our results demonstrate that the conventions required for a lexicon could be established among animals whose cognitive skills are relatively modest.

Though communication systems capable of conveying quite large numbers of meanings can emerge in populations using the Obverter procedure, our model does not address the emergence of communications systems that use the sequential structure of signals to convey complex meanings. However the fact that highly coordinated simple communication systems can emerge suggests that populations of animals who use them might become more and more dependent on communication to organize their activity, thus providing adaptive benefit for improving their communicative abilities still further.

It is possible, for example, that as a simple communication system increases in coordination, and thus more useful to a population of animals, their cognitive and social abilities relevant to communication would improve. New meanings could be added to the system, along with new signals to convey them. Eventually the difficulty of accurately performing and recognizing the different signals could lead to the development of complex signals with sequential regularities. Such regularities are observed in bird songs (Catchpole and Slater, 1995) and the long calls of some primate species (Robinson, 1984; Mitani and Marler, 1989). These sequential regularities could then be associated with the semantic structure of complex meanings. Batali (forthcoming) describes computational simulations of agents capable of sending and interpreting signals composed of sequence of tokens, and conveying structured meanings, who develop systems of coordinated communication as a result of using a learning procedure similar to Obverter. Such populations develop communication systems with sequential regularities reminiscent of grammatical phenomena in human languages.

In considering the problem of learning the syntactic regularities in a language, a number of theories have been proposed that assume that the child

must learn values for a fairly small set of parameters that determine the syntactic properties of the language, for example its standard word order, how thematic roles are assigned, constraints on the movement of constituents, and so forth. In some models the set of parameters, their possible values, and the learning mechanisms needed to set the values appropriately, are assumed to be language-specific and innate (Chomsky, 1987; Atkinson, 1992; Lightfoot, 1992). In other models the finite set of choices is based on functional considerations (Bates and MacWhinney, 1982). It has also been suggested that most of the information that speakers possess about the syntax of their language involves properties of its lexical items (Bresnan, 1982; Langacker, 1986; Chomsky, 1995). In all of these approaches, syntactic regularities are argued to be the result of interactions among a set of learned conventions, and our results show how such conventions can become established in a population, and can be acquired by an individual.

Research into lexical acquisition by children usually focuses on the problem of learning the relations between the categories and the sounds or gestures of an extant, highly coordinated communication system — a human language (Macnamara, 1982; Clark, 1993). An influential hypothesis is that children rely on assumptions about the relations among words and the concepts they denote, for example that two words that sound different must contrast in meaning, or that words correspond to categories that are organized taxonomically (Clark, 1983; Markman, 1987).

The Obverter learning procedure also incorporates a set of assumptions that constrain the relations among the meanings and signals that can be learned. However those assumptions derive entirely from considerations of communicative function, and not from any assumed organization of the set of meanings. This is consistent with suggestions by Baldwin (1993) and Tomasello (1995) that children attend to clues about communicative function, for example, pragmatic and contextual information, to acquire the meanings of words, and that this information reduces the need for the child to rely on additional constraints about how word meanings are related.

A subtle assumption in much of the research into lexical acquisition is perhaps inspired by de Saussure's (1916) notion of the bidirectional nature of the linguistic sign. No distinction is made between the task of learning to produce the correct word for a concept, and the task of learning to comprehending a word as referring to a specific concept. While evidence that

a child understands a word may come from experiments that involve comprehension or production or both, understanding the meaning of a word is usually treated as unitary phenomenon. Little justification is made for this assumption except that it is more or less true for adult language users.

We make no such assumption. Indeed, we separate the tasks of learning transmission and reception behavior completely. In a population whose new members use Obverter, the bidirectionality of signals is a *consequence* of how the learning procedure affects the population's communication system, not an assumed constraint on how learning is done.

Our model makes empirical predictions that could be used to test whether or not children use something like Obverter when acquiring a lexicon. We predict that the major influence on a child's production of words will be the child's observation of those words being understood by others, and that the major influence on a child's comprehension of words will be the child's observation of those words being produced by others. This would require experiments in which the two sets of observations are not as coordinated as they would be if words from human languages were used. A paradigm involving made-up words and concepts, similar to those described by Nelson and Bonvillian (1973), Woodward, et al. (1994), and Schafer and Plunkett (1996), could be employed.

Research involving children exposed to impoverished linguistic input suggests that they use learning strategies that would tend to improve the communicative accuracy of a population. In studies of deaf children exposed to the relatively unsystematic gestures of their parents, Goldin-Meadow and Mylander (1990) find that the children emerge with systems more complex, and more consistent, than the ones they were exposed to. Bickerton (1990) argues that properties of Creole languages show that children possess learning procedures that systematize and regularize the pidgin languages they observe. Although mostly concerned with the acquisition of syntax, Pinker's (1994) summary of such evidence as indicating that "children actually reinvent [language], generation after generation" [p. 32], is consistent both with the way that Obverter works (as opposed to imitation-based procedures), and with the fact that the use of Obverter by the new members of a population can literally result in the invention of a coordinated communication system.

6 Conclusion

The questions of how a coordinated system of communication could originate, and how individuals could learn such a system, might seem unrelated. To the contrary, we have shown that learning procedures exist that will result in the communication system of a population increasing in accuracy, ultimately yielding a highly coordinated system. Since these procedures will improve sub-optimal systems, they will also maintain the accuracy of an established system against degradation due to the noisy and limited input available to its learners.

The main reason that these procedures work so well is that they explicitly take the *communicative* aspects of communicative behavior into account. This is done by first separating the problem of learning to communicate into two subproblems, that of acquiring appropriate transmission behaviors, and of acquiring appropriate reception behaviors. Thus separated, the two problems can be solved in related ways. The Obverter procedure, and the approximations to it we describe, work by using the reception behavior of a population to determine how to send signals, and by using the transmission behavior of a population to determine how to receive signals. The resultant systems can be guaranteed to enable accurate communication with the population.

Our learning procedures involve generally useful cognitive abilities applied to the domain of social interaction. Our model is therefore most relevant to species of social animals capable of adjusting their behavior to that of others, and for whom the accurate exchange of information is often mutually beneficial. Though such species may be rare, humans almost certainly descended from one of them.

Appendix: Mathematical Details

A.1 Obverter Yields the Best Possible System

The Obverter learning procedure, described in figure 5, yields send and receive functions that have the highest possible communicative accuracy with a given population. This will be proved as follows: We first derive the maximum values that can be achieved for the population, and then show that

the Obverter procedure will produce send and receive functions that achieve those maximum values.

For each meaning μ , the sum $\sum_{\sigma} s(\mu, \sigma)$, where σ ranges over the set of signals, represents the probability that *some* signal is sent for μ . Likewise, for each signal σ , the sum $\sum_{\mu} r(\sigma, \mu)$, where μ ranges over the set of meanings, represents the probability that σ is interpreted as some meaning. So:

$$\sum_{\sigma} s(\mu, \sigma) \leq 1 \quad \text{and} \quad \sum_{\mu} r(\sigma, \mu) \leq 1 \quad (7)$$

The inequalities allow for the possibility that not all meanings are signaled, and not all signals are interpreted. In the case of the send functions in figures 1 and 2, for example, some signal is always sent for each meaning, and so all of the rows of the send functions sum to one. For the receive functions in figures 1 and 2, each signal is always interpreted as some meaning, and so all of the columns of the receive functions sum to one.

Suppose that a given population's average send and receive functions are S and R . The probability that a send function s is correctly interpreted by members of the population when it sends signals for some meaning μ is

$$\sum_{\sigma} s(\mu, \sigma) R(\sigma, \mu) \quad (8)$$

(See equation 2.) Let κ be the signal most often interpreted as μ by members of the population. Thus:

$$\sum_{\sigma} s(\mu, \sigma) R(\sigma, \mu) \leq \sum_{\sigma} s(\mu, \sigma) R(\kappa, \mu) \quad (9)$$

since $R(\kappa, \mu)$ is greater than or equal to each $R(\sigma, \mu)$ in the summation on the left. Since $R(\kappa, \mu)$ is a constant in equation 9, we can move it out of the summation to obtain:

$$\sum_{\sigma} s(\mu, \sigma) R(\sigma, \mu) \leq R(\kappa, \mu) \sum_{\sigma} s(\mu, \sigma) \quad (10)$$

Finally, since $\sum_{\sigma} s(\mu, \sigma) \leq 1$, we have:

$$\sum_{\sigma} s(\mu, \sigma) R(\sigma, \mu) \leq R(\kappa, \mu) \quad (11)$$

So $R(\kappa, \mu)$ is the maximum possible probability that signals about meaning μ are correctly interpreted by members of the population.

The Obverter learning procedure will create a send function s_{ob} with $s_{ob}(\mu, \kappa) = 1.0$. It will communicate about meaning μ with members of the population successfully with probability $R(\kappa, \mu)$. This, as we have just shown, is the maximum value that any send function could achieve.

Given a send function s_{ob} derived according to the Obverter procedure, we have, for any other send function s

$$\frac{1}{|M|} \sum_{\mu} \sum_{\sigma} s(\mu, \sigma) R(\sigma, \mu) \leq \frac{1}{|M|} \sum_{\mu} \sum_{\sigma} s_{ob}(\mu, \sigma) R(\sigma, \mu) \quad (12)$$

since each term in the sum over μ on the right is greater than or equal to the corresponding term on the left. Thus:

$$ca(s, R) \leq ca(s_{ob}, R) \quad (13)$$

and so the send function s_{ob} has the highest possible probability of having its signals correctly interpreted by members of the population.

Two considerations are in order: (1) Suppose that for some meaning μ , all $R(\sigma, \mu) = 0$. In this case, no member of the population interprets signals about μ at all, and so no send function could communicate about μ correctly; (2) Suppose that for some μ , several κ have equal, maximal values of $R(\kappa, \mu)$. In both cases there is more than one signal whose probability of being interpreted as a given meaning is equal to the maximum value. In both cases, the inequality in formula 9 still holds for any one of the signals with the maximum value of being interpreted as μ (whether it is zero or not). Accordingly, the description of the Obverter procedure in figure 5 should be modified to handle this case, as follows:

If, during step **s.1** of the Obverter procedure, while trying to find a signal to send for meaning μ , it is found that more than one κ have equal, maximal values for $R(\kappa, \mu)$, choose one at random, call it κ_{μ} . Then continue with step **s.2**.

Similar results obtain when considering the best possible receive function for communicating with a population. First we note that for a given signal σ , a receive function r will correctly interpret σ with a probability given by

$$\sum_{\mu} S(\mu, \sigma) r(\sigma, \mu) \quad (14)$$

Let η be the meaning most often encoded as σ . So:

$$\sum_{\mu} S(\mu, \sigma) r(\sigma, \mu) \leq \sum_{\mu} S(\eta, \sigma) r(\sigma, \mu) = S(\eta, \sigma) \sum_{\mu} r(\sigma, \mu) \leq S(\eta, \sigma) \quad (15)$$

Thus $S(\eta, \sigma)$ is the maximum probability that any receive function will correctly interpret signal σ when sent by members of the population. The receive function r_{ob} that Obverter will create will have $r_{ob}(\sigma, \eta) = 1.0$.

For any receive function r :

$$\frac{1}{|M|} \sum_{\mu} \sum_{\sigma} S(\mu, \sigma) r(\sigma, \mu) \leq \frac{1}{|M|} \sum_{\mu} \sum_{\sigma} S(\mu, \sigma) r_{ob}(\sigma, \mu) \quad (16)$$

So:

$$ca(S, r) \leq ca(S, r_{ob}) \quad (17)$$

The receive function r_{ob} has the highest possible probability of correctly interpreting signals sent by members of the population.

As in the discussion above, the possibility arises that $S(\mu, \sigma)$ is zero for all of the meanings, or that several meanings have equal, maximal values for a given σ . As in the previous derivation, the result still holds. In the case where the maximal value is nonzero for a given signal, we will, as before, choose a meaning at random. However in the case where a signal is never sent, there is no point in defining *any* interpretation for it at all — indeed if the set of available signals is infinite, this would be impossible. Hence the Obverter procedure should be modified as follows:

If during step **r.1** of the Obverter procedure, more than one η have equal, maximal, nonzero values for $S(\eta, \sigma)$, choose one at random, call it η_{σ} . Continue with step **r.2**. If, during step **r.1**, no $S(\eta, \sigma) > 0$, $r_{ob}(\sigma, \mu)$ is undefined for all μ .

In a population of finite size, whose older members die off and are replaced by learners using the Obverter procedure, there will come a point where, for some meaning μ , a certain signal σ , is interpreted most often as that meaning, and that signal is rarely, if ever, interpreted as any other meaning.

After that point, all send functions produced by Obverter will send σ for μ with probability 1.0. The fraction of the population that sends σ for μ will then increase, as new members replace older ones, and ultimately σ will be the signal sent most often for μ . After this point, all receive functions produced by obverter will interpret σ as μ with probability 1.0. Eventually, all members will both send σ for μ with probability 1.0, and interpret σ as μ with probability 1.0, thus achieving the maximum possible communicative accuracy for that meaning. As the same process will happen for all meanings (assuming that the number of signals is at least as large as the number of meanings), the population will ultimately develop an optimally coordinated communication system.

A.2 Unit-Statistic Learning

The Unit-Statistic learning procedure, described in figure 7, requires that the learner record exactly one observation, for each meaning, of some signal being sent for that meaning, and of some signal being interpreted as that meaning. These observations are used to create send and receive functions by using a procedure similar to Obverter. In this appendix we show that the resultant send and receive functions will have an expected value of two-way communicative accuracy with the members of the population that is higher than the population average communicative accuracy.

Consider a meaning μ . Signal σ is sent for this meaning with average probability $S(\mu, \sigma)$ by members of the population. If the learner records exactly one instance of a signal being sent for μ , it will record that σ is sent for μ with probability $S(\mu, \sigma)$. If the learner now always interprets σ as μ (in accordance with the Unit-Statistic procedure), it will be correct with probability $S(\mu, \sigma)$. The probability that a set of learners using the Unit-Statistic learning procedure will correctly interpret signals about meaning μ sent by a population whose average send function is S is therefore:

$$\sum_{\sigma} S(\mu, \sigma)^2 \tag{18}$$

The expected value of the communicative accuracy of the receive functions for all meanings is:

$$\frac{1}{|M|} \sum_{\mu} \sum_{\sigma} S(\mu, \sigma)^2 \tag{19}$$

Similar considerations hold for learners constructing their send functions based on one observation in which, for example, signal σ is observed being interpreted as meaning μ . According to the Unit-Statistic procedure, they will always send σ for μ . Their signals will be correctly interpreted by the population with probability:

$$\sum_{\sigma} R(\sigma, \mu)^2 \quad (20)$$

The expected value of the communicative accuracy of the learned send functions for all meanings is:

$$\frac{1}{|M|} \sum_{\mu} \sum_{\sigma} R(\mu, \sigma)^2 \quad (21)$$

So the expected value of the two-way communicative accuracy between Unit-Statistic learners and the population is:

$$\frac{1}{2} \left(\frac{1}{|M|} \sum_{\mu} \sum_{\sigma} R(\sigma, \mu)^2 + \frac{1}{|M|} \sum_{\mu} \sum_{\sigma} S(\mu, \sigma)^2 \right) \quad (22)$$

This can be simplified to:

$$\frac{1}{2|M|} \left(\sum_{\mu} \sum_{\sigma} R(\sigma, \mu)^2 + S(\mu, \sigma)^2 \right) \quad (23)$$

We now subtract the average communicative success of the members of the population with each other from the quantity in 23:

$$\begin{aligned} & \frac{1}{2|M|} \left(\sum_{\mu} \sum_{\sigma} R(\sigma, \mu)^2 + S(\mu, \sigma)^2 \right) \\ & - \frac{1}{2|M|} \left(\sum_{\mu} \sum_{\sigma} S(\mu, \sigma) R(\sigma, \mu) + S(\mu, \sigma) R(\sigma, \mu) \right) \end{aligned} \quad (24)$$

This expression can be converted to:

$$\frac{1}{2|M|} \left(\sum_{\mu} \sum_{\sigma} R(\sigma, \mu) [R(\sigma, \mu) - S(\mu, \sigma)] + S(\mu, \sigma) [S(\mu, \sigma) - R(\sigma, \mu)] \right) \quad (25)$$

Which equals:

$$\frac{1}{2|M|} \left(\sum_{\mu} \sum_{\sigma} R(\sigma, \mu) [R(\sigma, \mu) - S(\mu, \sigma)] - S(\mu, \sigma) [R(\sigma, \mu) - S(\mu, \sigma)] \right) \quad (26)$$

Yielding:

$$\frac{1}{2|M|} \left(\sum_{\mu} \sum_{\sigma} [R(\sigma, \mu) - S(\mu, \sigma)]^2 \right) \quad (27)$$

This quantity is always greater than or equal to zero, and therefore the expected value of the communicative accuracy of learners using the Unit-Statistic procedure be higher than that of the population average.

The value of formula 27 equals zero only when all corresponding values of $S(\mu, \sigma)$ and $R(\sigma, \mu)$ are exactly equal. This could occur if population has reached a state of optimal communication, in which, for each meaning μ there is a signal σ that is sent by all with probability 1.0, and which is interpreted by all with probability 1.0. In this case, of course, the learner could do no better (and will, in fact, do as well as) the rest of the population.

It is also possible that formula 27 could equal zero in some population for which an optimal communication system has not been achieved. Such cases are going to be highly unstable, as random fluctuations in the S and R entries will cause the quantity in formula 27 to attain a positive value, and the average communicative accuracy will increase.

Though in principle a population of finite size whose new members use the Unit-Statistic procedure could achieve an optimally coordinated system, this outcome cannot be guaranteed. For it to occur, each new learner's single recorded observation of each of the meanings would have to be identical. While the probability of this happening increases as the coordination of the system increases, and the coordination of the system *will* increase, as has just been shown, the probability equals 1.0 only if the system is optimally coordinated. Thus the expected value of the population's communicative accuracy will asymptotically approach 1.0, but may or may not actually achieve that value in any specific population in a finite amount of time.

A.3 Temporal Dynamics of the Learning Procedures

The different shapes of the plots of the computational simulations in figure 8 illustrate the qualitative differences in the temporal dynamics of communica-

tive accuracy among populations using different learning procedures.

To analyze the differences, let us consider a meaning μ and a signal σ such that, at some point in the simulation, σ is most often sent for μ and is also the signal most often interpreted as μ . Let p be the probability that the signal σ is sent for μ , and is then interpreted correctly, that is:

$$p = S(\mu, \sigma)R(\sigma, \mu) \quad (28)$$

Given these assumptions, the Obverter learning procedure will result in each learner sending σ for μ with probability 1.0, and interpreting σ as μ with probability 1.0. Thus p will increase by an amount proportional to $(1 - p)/N$ in each round of the simulation, where N is the number of individuals in the population. If the population size is large, this expression will be proportional to the derivative of p as a function of time, and therefore:

$$p(t) = 1 - ke^{-rt/N} \quad (29)$$

Where $k = 1 - p(0)$ and r depends on details of the simulation, for example the number of meanings and signals, and the rate at which members of the population are replaced.

In the case of Unit-Statistic learning, the value of p will not necessarily increase each round, as the chance that the learner will observe μ being encoded as σ , or σ being interpreted as μ , depends on the specific entries in the population's send and receive functions. The probability p describes the chance that both events are observed in a single communicative event, and thus the probability that a Unit-Statistic learner will encode μ as σ and interpret σ as μ . So the value of p will increase by $(1 - p)/N$ with a probability of p . If the population size is large enough, we can take the time derivative of p to be $p(1 - p)/N$, and therefore:

$$p(t) = \frac{1}{1 + ke^{-rt/N}} \quad (30)$$

Where $k = p(0)/(1 - p(0))$, and r depends on the simulation details.

The shapes of the curves labeled ‘‘Obverter’’ and ‘‘Unit-Stat’’ in figure 8 are consistent with the functions in equations 29 and 30.

For the learning procedures based on a fixed number of observations, it might be expected that their performance would approximate that of Obverter if

the number of observations is large enough, and this is apparently the case for Obs-25. We might expect that fewer observations would result in performance similar to the Unit-Statistic procedure, but this is not entirely correct. The reason that the Unit-Statistic procedure is guaranteed to increase the communicative accuracy of the population is that it must observe, for each meaning, one episode of that meaning being encoded as a signal, and one episode of some signal being interpreted as that meaning. However there is no guarantee, for any finite number of observations, that one of each such episode will be seen. If a learner with a finite number of observations fails to observe one or more such events, it will send or receive randomly for the meanings whose transmission or reception was not observed. The presence of such a learner will degrade the average communicative accuracy of a population whose system is optimal, and the communicative accuracy of a population whose members use this learning procedure will be limited to some value below 1.0.

To estimate the limiting value, we first compute the probability that each meaning will be observed in some number of observations. With three meanings and ten observations, this probability is 0.948. This value represents the probability that an Obs-10 learner's send map entry will be correct, and the probability that its receive map entry will be correct. So the maximum communicative accuracy such a population could achieve is $0.948^2 = 0.899$. In the case of Obs-25 the probability that all three meanings will be observed is approximately 0.9998, and so the threshold is greater than 0.999. Though these estimates ignore other influences on the limiting values, they are roughly consistent with the performances of the Obs-10 and Obs-25 learners shown in figure 8.

References

- [1] David H. Ackley and Michael L. Littman. Altruism in the evolution of communication. In R. Brooks and P. Maes, editors, *Proceedings of the Fourth Artificial Life Workshop*, Cambridge, MA, 1994. MIT Press.
- [2] M. Andersson. Why are there so many threat displays? *Theoretical Biology*, 86:773–781, 1980.

- [3] Martin Atkinson. *Children's Syntax: An Introduction to Principles and Parameters Theory*. Blackwell, Oxford, 1992.
- [4] J. L. Austin. *How to do Things with Words*. Clarendon Press, Oxford, 1962.
- [5] D. A. Baldwin. Infants' ability to consult the speaker for cues to word reference. *Journal of Child Language*, 20:395–418, 1983.
- [6] S. A. Barnett. *The Rat: A Study in Behavior*. The University of Chicago Press, 1975.
- [7] John Batali. Computational simulations of the emergence of grammar. forthcoming.
- [8] E. Bates and B. MacWhinney. Functionalist approaches to grammar. In E. Wanner and L. R. Gleitman, editors, *Language Acquisition: The State of the Art*. Cambridge University Press, 1982.
- [9] J. Bennett. *Linguistic Behavior*. Cambridge University Press, 1976.
- [10] Derek Bickerton. *Language and Species*. University of Chicago Press, 1990.
- [11] Andreas Blume, Yong-Gwan Kim, and Joel Sobel. Evolutionary stability in games of communication. *Games and Economic Behavior*, 5:547–575, 1993.
- [12] D. T. Blumstein. Golden-Marmot alarm calls: The production of situationally specific vocalizations. *Ethology*, 100(2):113–125, June 1995.
- [13] Joan Bresnan, editor. *The Mental Representation of Grammatical Relations*. MIT Press, Cambridge, MA, 1982.
- [14] Robert Burton. *Bird Behavior*. Knopf/Random House, 1985.
- [15] R. Byrne and A. Whiten. Towards the next generation of data quality: A new survey of primate tactical deception. *Behavioral and Brain Sciences*, 11:267–273, 1988.

- [16] David Canning. Learning language conventions in common interest signaling games. Unpublished manuscript, Columbia University, 1992.
- [17] C. K. Catchpole and P. J. B. Slater. *Bird Song: Biological Themes and Variations*. Cambridge University Press, 1995.
- [18] Dorothy L. Cheney and Robert M. Seyfarth. *How Monkeys See the World: Inside the Mind of Another Species*. University of Chicago Press, 1990.
- [19] I.-K. Cho and D. Kreps. Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102:179–222, 1987.
- [20] Noam Chomsky. *Knowledge of Language: Its Nature, Origin, and Use*. Praeger, New York, 1987.
- [21] Noam Chomsky. *The Minimalist Program*. MIT Press, 1995.
- [22] Eve Clark. Convention and contrast in acquiring the lexicon. In T. B. Seiler and W. Wannenmacher, editors, *Cognitive Development and the Development of Word Meaning*, pages 67–89. Lawrence Erlbaum, 1983.
- [23] Eve Clark. *The Lexicon in Acquisition*. Cambridge University Press, 1993.
- [24] Vincent P. Crawford and John Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, November 1982.
- [25] K. B. Dasilva, D. L. Kramer, and D. M. Weary. Context-specific alarm calls of the eastern chipmunk, *tamias striatus*. *Canadian Journal of Zoology — Revue Canadienne de Zoologie*, 72(6):1087–1092, June 1994.
- [26] Donald Davidson. Communication and convention. In *Inquiries into Truth and Interpretation*, pages 265–280. Clarendon/Oxford University Press, 1984.
- [27] Ferdinand de Saussure. *Course in General Linguistics*. Duckworth, London, 1983. Translated by R. Harris; Originally published in 1916.

- [28] Daniel C. Dennett. Intentional systems in cognitive ethology: The “Panglossian paradigm” defended. *Behavioral and Brain Sciences*, 6:343–355, 1983.
- [29] C. S. Evans, L. Evans, and P. Marler. On the meaning of alarm calls — functional reference in an avian vocal system. *Animal Behavior*, 46(1):23–38, July 1993.
- [30] Christopher Evans and Peter Marler. Language and animal communication: Parallels and contrasts. In H. L. Roitblat and J.-A. Meyer, editors, *Comparative Approaches to Cognitive Science*. MIT Press/Bradford Books, 1995.
- [31] Joseph Farrell. Communication, coordination, and Nash equilibrium. *Economics Letters*, 27:209–214, 1988.
- [32] Joseph Farrell. Meaning and credibility in cheap-talk games. *Games and Economic Behavior*, 5:514–531, 1993.
- [33] S. Goldin-Meadow and C. Mylander. Beyond the input given: The child’s role in the acquisition of language. *Language*, 66(2):323–355, 1990.
- [34] H. Paul Grice. Meaning. *Philosophical Review*, pages 377–388, 1957.
- [35] H. Paul Grice. *Studies in the Way of Words*. Harvard University Press, 1989.
- [36] Marc D. Hauser. *The Evolution of Communication*. MIT Press/Bradford Books, 1996.
- [37] Cecilia M. Heyes and Jr. Bennett G. Galef, editors. *Social Learning in Animals: The Roots of Culture*. Academic Press, 1996.
- [38] Robert A. Hinde. Animal signals: Ethological and games-theory approaches are not incompatible. *Animal Behavior*, 2:535–542, 1981.
- [39] Charles F. Hockett. Logical considerations in the study of animal communication. In W. E. Lanyon and W. N. Tavolga, editors, *Animal Sounds and Communication*. American Institute of Biological Sciences, Washington, D.C., 1960.

- [40] Charles F. Hockett. The origin of speech. *Scientific American*, 203, September 1960.
- [41] J. L. Hoogland. Nepotism and alarm calling in the black-tailed prairie dog, *cynomys ludovicianus*. *Animal Behavior*, 36:472–479, 1983.
- [42] J. L. Hoogland. *The Black-Tailed Prairie Dog: Social Life of a Burrowing Mammal*. University of Chicago Press, 1995.
- [43] Nicholas K. Humphrey. The social function of intellect. In P. P. G. Bateson and R. A. Hinde, editors, *Growing Points in Ethology*, pages 303–317. Cambridge University Press, 1976.
- [44] J. R. Hurford. Biological evolution of the Saussurean sign as a component of the language acquisition device. *Lingua*, 77:187–222, 1989.
- [45] Edwin Hutchins and Brian Hazlehurst. Learning in the cultural process. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II: SFI Studies in the Sciences of Complexity, Volume X*, pages 689–706, 1991.
- [46] Christine M. Johnson. Animal communication via coordinated cognitive systems. *Ethology*, 10:187–207, 1993.
- [47] Yong-Gwan Kim and Joel Sobel. An evolutionary approach to pre-play communication. *Econometrica*, 63(5):1181–1193, 1995.
- [48] John R. Krebs and Richard Dawkins. Animal signals: Mind-reading and manipulation. In J. R. Krebs and N. B. Davies, editors, *Behavioral Ecology: An Evolutionary Approach*, chapter 15, pages 380–402. Blackwell Scientific Publications, Oxford, second edition, 1984.
- [49] Ronald W. Langacker. An introduction to cognitive grammar. *Cognitive Science*, 10:1–40, 1986.
- [50] David K. Lewis. *Convention: A Philosophical Study*. Harvard University Press, 1969.
- [51] David K. Lewis. Languages and language. In Keith Gunderson, editor, *Minnesota Studies in the Philosophy of Science, Vol. VII: Language, Mind and Knowledge*. University of Minnesota Press, 1975.

- [52] David Lightfoot. *How to Set Parameters: Arguments from Language Change*. MIT Press/Bradford Books, 1992.
- [53] Bruce J. MacLennan and Gordon M. Burghardt. Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2(2):161–187, 1993.
- [54] John Macnamara. *Names for Things: A Study of Human Learning*. MIT Press/Bradford Books, 1982.
- [55] Ellen M. Markman. How children constrain the possible meanings of words. In U. Neisser, editor, *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, pages 255–287. Cambridge University Press, 1987.
- [56] J. Maynard-Smith. The evolution of alarm calls. *American Naturalist*, 99, 1965.
- [57] J. Maynard-Smith and G. R. Price. The logic of animal conflict. *Nature*, 246(15), 1973.
- [58] J. Mitani and P. Marler. A phonological analysis of male gibbon singing behavior. *Behavior*, 109:20–45, 1989.
- [59] M. Moynihan. Control, suppression, decay, disappearance and replacement of displays. *Journal of Theoretical Biology*, 29:85–112, 1970.
- [60] C. A. Munn. The deceptive use of alarm calls by sentinel species in mixed-species flocks of neotropical birds. In R. W. Mitchell and N. S. Thompson, editors, *Deception: Perspectives on Human and Nonhuman Deceit*, pages 169–176. S.U.N.Y. Press, Albany, 1986.
- [61] K. E. Nelson and J. D. Bonvillian. Concepts and words in the 18 month-old: Acquiring concept names under controlled conditions. *Cognition*, 2(4):435–450, 1973.
- [62] Michael Oliphant. The dilemma of Saussurean communication. *BioSystems*, pages 31–38, 1996.

- [63] Michael Oliphant. Communication as altruistic behavior. Ph.D. Dissertation, Department of Cognitive Science, U.C.S.D., in preparation, 1997.
- [64] Ezequiel A. Di Paolo. An investigation into the evolution of communicative behaviors. Cognitive Science Research Paper 445, School of Cognitive and Computing Sciences, University of Sussex., 1996.
- [65] Domenico Parisi. An artificial life approach to language. *Brain and Language*, in press.
- [66] Steven Pinker. *The Language Instinct: How the Mind Creates Language*. William Morrow and Company, New York, 1994.
- [67] F. Harvey Pough, John B. Heiser, and William M. McFarland. *Vertebrate Life*. Prentice-Hall, fourth edition, 1996.
- [68] D. Premack. The codes of man and beast. *Behavioral and Brain Sciences*, 6:125–167, 1983.
- [69] J. G. Robinson. Syntactic structure in the vocalizations of wedge-capped capuchin monkeys, *cebus nigrivittatus*. *Behavior*, 90:46–79, 1984.
- [70] Daniel I. Rubenstein and Richard W. Wrangham, editors. *Ecological Aspects of Social Evolution: Birds and Mammals*. Princeton University Press, 1985.
- [71] Graham Schafer and Kim Plunkett. Rapid word learning in 15-month-olds under tightly controlled conditions. *CRL Newsletter, U.C.S.D., La Jolla, CA*, March 1996. Submitted to *Child Development*.
- [72] John R. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, 1969.
- [73] R. M. Seyfarth, D. L. Cheney, and P. Marler. Vervet monkey alarm calls: semantic communication in a free-ranging primate. *Animal Behavior*, 28:1070–1094, 1980.
- [74] Brian Skyrms. *Evolution of the Social Contract*. Cambridge University Press, 1996.

- [75] C. N. Slobodchikov, J. Kiriazis, C. Fischer, and E. Creef. Semantic information distinguishing individual predators in the alarm calls of gunnison prairie dogs. *Animal Behavior*, 42:713–719, November 1991.
- [76] W. John Smith. *The Behavior of Communicating: An Ethological Approach*. Harvard University Press, 1977.
- [77] W. John Smith. Communication and expectations: A social process and the cognitive operations it depends upon and influences. In M. Bekoff and D. Jamieson, editors, *Readings in Animal Cognition*, pages 143–255. Bradford Books/MIT Press, 1996.
- [78] Richard Sorabji. Rationality. In Michael Frede and Gisela Striker, editors, *Rationality in Greek Thought*. Clarendon/Oxford University Press, 1996.
- [79] J. M. Spence. Job market signaling. *Quarterly Journal of Economics*, 87:296–332, 1973.
- [80] Luc Steels. Emergent adaptive lexicons. In P. Maes, editor, *Proceedings of the Simulation of Adaptive Behavior Conference*. MIT Press, 1996.
- [81] P. F. Strawson. Intention and convention in speech acts. *Philosophical Review*, 73, 1974.
- [82] T. T. Struhsaker. Auditory communication among vervet monkeys (*cercopithicus aethiops*). In S. A. Altmann, editor, *Social Communication Among Primates*. University of Chicago Press, 1967.
- [83] N. Tinbergen. Derived activities: Their causation, biological significance, origin and emancipation during evolution. *Quarterly Review of Biology*, 27:1–32, 1952.
- [84] Michael Tomasello. Pragmatic contexts for early verb learning. In M. Tomasello and W. E. Merriman, editors, *Beyond Names For Things: Young Children's Acquisition of Verbs*. Lawrence Erlbaum, 1995.
- [85] Karl Warneryd. Cheap talk, coordination, and evolutionary stability. *Games and Economic Behavior*, 5:532–546, 1993.

- [86] G. M. Werner and M. G. Dyer. Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II: SFI Studies in the Sciences of Complexity, Volume X*, pages 511–547, 1991.
- [87] Andrew Whiten, editor. *Natural Theories of Mind: Evolution, Development and Simulation of Everyday Mindreading*. Blackwell, Oxford, 1991.
- [88] A. L. Woodward, E. M. Markman, and C. M. Fitzsimmons. Rapid word learning in 13- and 18-month-olds. *Developmental Psychology*, 30(4):553–566, 1994.
- [89] H. Yanko and L. Stein. An adaptive communication protocol for cooperating mobile robots. In J-A Meyer, H. L. Roitblat, and S. Wilson, editors, *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. MIT Press, 1993.