CHAPTER 3

The cultural evolution of communication

The first step in an investigation into the evolution of the distinguishing design features of language — cultural transmission, symbolicism and compositionality — is to consider the cultural evolution of simple, unstructured communication systems. Such systems are the equivalent of a symbolic vocabulary. In particular, I will search for the circumstances under which optimal communication emerges. My line of reasoning, for the purpose of this chapter, is as follows. Humans have a culturally transmitted, symbolic vocabulary. Therefore, humans must have the necessary mental apparatus to support such a symbolic vocabulary. I will make the default adaptationist assumption that this capacity must have provided a fitness payoff at some point in evolutionary history, and this payoff must have been due to the communicative benefits of symbolic vocabulary. Therefore, humans have the necessary mental apparatus to support a *communicatively useful*, culturally transmitted symbolic vocabulary. The aim of this chapter is therefore to identify the learning bias required to support a communicatively useful vocabulary, in the strongest case an optimal communication system, and equate this with the mental capacity of humans.

The assumption that the human capacity for symbolic vocabulary evolved due to fitness payoff arising from communication will be reexamined in Chapter 4.

In Section 3.1 I review previous computational models which tackle this issue. In Section 3.2 a simple model of communication is developed. In Section 3.3 a new ILM is introduced. This model shows that the key determinant of the population's communicative behaviour is the direct bias on cultural transmission resulting from the learner's bias. In light of this, a more sophisticated model is developed in Section 3.4 to investigate a wider range of learning biases. In Section 3.5 the key bias for the cultural evolution of vocabulary is identified and defined. Finally, parallels are drawn between this learning bias and the learning bias applied by human language learners to the task of vocabulary acquisition. This comparison suggests that the human learning bias (and perhaps not the learning bias of other, closely related species) exhibits the appearance of design for the cultural evolution of communicatively-optimal symbolic vocabulary.

3.1 Models of the evolution of vocabulary

In the review carried out in Chapter 2 I covered the models of the cultural evolution of vocabulary systems described in Hutchins & Hazelhurst (1995) and Nowak *et al.* (1999). To summarise briefly, Hutchins & Hazelhurst (working within the NM framework) report that communicatively-optimal symbolic vocabulary evolves culturally. I attributed this to a direct bias acting on cultural transmission, a consequence of the autoassociator network architecture used in their model. Nowak *et al.* also demonstrate (working within the ILM framework) that communicatively-optimal symbolic vocabulary can evolve culturally. However, in their model this is a consequence of the natural selection of cultural variants — successful communicators are more likely to act as cultural parents, therefore successful communication systems are preferentially retained in the population.

Hurford (1989) describes possibly the first computational investigation into the evolution of communication. One of the central concerns of Hurford's paper is the biological evolution of learning strategies, and his work in this area will be returned to in Chapter 4. However, Hurford does cover purely cultural evolution as well.

The communicative behaviour of individuals in Hurford's model is represented with two probability matrices — a production matrix, which gives the probability of producing a particular signal given a certain meaning, and a reception matrix, which gives the probabilistic reception behaviour of the individual. A generational ILM is used. Learners form their production and reception matrices based on a sample of the observable behaviour produced by the previous generation. Each learner samples once from the population's production and reception behaviour, yielding a set of observed meaning-signal pairs (based on a stochastic sample of the population's production behaviour) and a set of observed signal-meaning pairs (based on a stochastic sample of the population's reception behaviour). Hurford considers three learning strategies. *Imitator* agents form their transmission matrix based on observed transmission behaviour, and their reception matrix based on observed reception behaviour. *Calculators* base their production behaviour on observed reception and their reception on observed production. *Saussureans* base their production behaviour on observed production, then derive their reception behaviour from this matrix.

Imitator learners form their production and reception matrixes on the basis of direct observation of production and reception behaviour — if, for example, an Imitator observes signal $s1$ being produced for meaning $m1$, it will set the probability of producing $s1$ for $m1$ to 1 in its own production matrix.

Saussurean learners derive their production matrix from production behaviour in a similar manner, then design their reception matrix so as to make it optimally coordinated with their own production behaviour — for example, if a Saussurean learner arrives at a production matrix where $s1$ is produced for meanings $m1$ and $m2$, then it will interpret $s1$ as meaning $m1$ with probability 0.5 and $m2$ with probability 0.5.

Calculators form their reception matrix on the basis of observed production behaviour, in the same way that a Saussurean learner forms its reception matrix on the basis of its own production behaviour. For example, if a Calculator learner observes a population where $s1$ is produced for meanings $m1$ and $m2$, then it will interpret $s1$ as meaning $m1$ with probability 0.5 and $m2$ with probability 0.5. By the same optimisation process, Calculators calculate their production matrix on the basis of observed reception behaviour.

Hurford reports two results with relation to this ILM. Firstly, populations of Calculator agents are unable to preserve an optimal communication system over time. Secondly, populations of Imitator and Saussurean learners are capable of creating communication systems which lead to intermediate levels of communicative accuracy through purely cultural processes.

There are two candidate pressures acting on cultural transmission in Hurford's model. Firstly, the behaviour of the populations could be explained by natural selection of cultural variants — more successful communicators are more likely to act as cultural parents in Hurford's model. Secondly, the different learning strategies could result in different direct biases on cultural transmission. I will demonstrate in Section 3.5.3 that this later pressure is probably the key one.

Oliphant & Batali (1997) introduce another Iterated Learning Model of vocabulary. Individuals are required to communicate about a small set of meanings using a small set of signals. As in Hurford's (1989) model, individuals are modelled using probabilistic functions, with each individual being characterised by a production function, which gives the probability of each signal being sent for a given meaning, and a reception function, which gives the probability of a given signal being interpreted as a particular meaning.

Oliphant & Batali use a gradual population turnover model. At each time-step a single individual is removed from the population and replaced by a new individual. This individual estimates the average production and reception functions in use in the population, by making a number of observations of the population's production and reception behaviour. Based on these estimated functions, the new individual then creates its own production and reception functions according to one of two learning procedures, termed *Imitate-Choose* (a slight variation on Hurford's Imitator) and *Obverter* (related to Hurford's Calculator).

The learner's estimation of the probability with which the population produces signal σ_j for meaning μ_i is given by $P(\mu_i, \sigma_j)$ and the learner's estimation of the probability with which the population interprets σ_j as meaning μ_i is given by $R(\sigma_j, \mu_i)$. The learner must choose their own production and reception probabilities, given by $p(\mu_i, \sigma_j)$ and $r(\sigma_j, \mu_i)$.

The Imitate-Choose learner proceeds as follows:

For each meaning μ_i:
- Find the signal σ_j for which $P(\mu_i, \sigma_j)$ is maximum.
- Set $p(\mu_i, \sigma_j) = 1$ and $p(\mu_i, \sigma_k) = 0$ for all $k \neq j$.

For each signal σ_j:
- Find the meaning μ_i for which $R(\sigma_j, \mu_i)$ is maximum.
- Set $r(\sigma_j, \mu_i) = 1$ and $r(\sigma_j, \mu_k) = 0$ for all $k \neq i$.

The Imitate-Choose learner therefore bases its production behaviour on the average production behaviour of the population, selecting the most frequently used signal for each meaning. Similarly, reception behaviour is based on the population's reception behaviour, with the most frequent interpretation of a given signal being learned as the *only* interpretation of that signal. Note that there is no coupling between production and reception behaviour.

The Obverter learning procedure proceeds as follows:

For each meaning μ_i:
- Find the signal σ_j for which $R(\sigma_j, \mu_i)$ is maximum.
- Set $p(\mu_i, \sigma_j) = 1$ and $p(\mu_i, \sigma_k) = 0$ for all $k \neq j$.

For each signal σ_j:
- Find the meaning μ_i for which $P(\mu_i, \sigma_j)$ is maximum.
- Set $r(\sigma_j, \mu_i) = 1$ and $r(\sigma_j, \mu_k) = 0$ for all $k \neq i$.

Obverter learners will produce the signal which is most commonly interpreted by the rest of the population as conveying the meaning they wish to convey. Similarly, Obverter learners will interpret a signal as meaning the meaning it is most frequently produced for. The Obverter learner therefore bases its production behaviour on the population's reception behaviour and its reception behaviour on the population's production behaviour. Note that, unlike in the Imitate-Choose strategy, this results in the coupling of production and reception behaviour — the population's production behaviour at time t will influence its reception behaviour at time $t + 1$.

Oliphant & Batali define a measure of communicative accuracy for individuals using these probabilistic send and receive functions, and measure how the communicative accuracy of a population changes over time as new individuals are introduced and learn according to one of the two strategies. Communicative accuracy within the population is measured according to a variant of the canonical formula given in Chapter 2, Section 2.2.2.1 — simply put, a communicative episode between two individuals is a success if the hearer interprets the form produced by the speaker as conveying the meaning that the speaker intended.

Oliphant & Batali report that the Imitate-Choose strategy can increase communicative accuracy among a population where communicative accuracy is already high. However, in poorly coordinated populations the use of Imitate-Choose can result in further degradation. In contrast, use of the Obverter strategy always results in a steady increase in communicative accuracy until optimal levels are reached.

Why does the Obverter learning strategy result in optimal communication, but the Imitate-Choose strategy does not? Oliphant & Batali attribute the success of the Obverter strategy to its implicitly communicative aims — during learning, signals are selected so as to maximise their probability of being understood. Imitator agents do not have this built-in understanding of the communicative task. They suggest that both strategies build in ambiguity-avoiding measures — in both cases the most popular meaning-signal combinations are selected to the exclusion of other possible combinations. As we will see in the remainder of this Chapter, these comments are somewhat wide of the mark. Firstly, it is not necessary to build in an implicit understanding of the communicative task — in Section 3.4 I will demonstrate that optimal communication can emerge in a population of learners who do not select signals so as to maximise the probability of being understood. Secondly, building in an understanding of the communicative task does not necessarily lead to optimal communication — Hurford's Calculators are similar to Oliphant & Batali's Obverters, but cannot even preserve an optimal system. Finally, it will be shown in Section 3.5.3 that the Obverter and Imitate-Choose strategies respond differently to different types of ambiguity, and that this difference is crucial in understanding the behaviour of populations of such learners.

What is clear, however, is that the emergence or non-emergence of optimal communication in these populations is driven by what B&R term directly biased transmission — the Imitate-Choose and Obverter strategies have different biases as to how they acquire communication systems, and over the course of repeated cultural transmission the cultural variants which most closely match these biases come to dominate the population. We can surmise that the Imitate-Choose and Obverter strategies have different biases, with only the Obverter strategy being biased in favour of communication systems which maximise communicative accuracy.

Livingstone & Fyfe (1999) investigate the evolution of vocabulary using a model which is a hybrid NM-ILM. Livingstone and Fyfe's main concern is the emergence of diversity of vocabulary, but they do make some observations of the overall structure of the vocabularies in their populations. Individuals are modelled using neural networks, mapping from input signals to output meanings[1]. The neural network model of an agent has N input units and M output units, where these units can take values of ± 1. Meanings are represented by patterns of activation over the M nodes where a single node has an activation of $+1$. This yields M distinct meanings. Signals are represented by arbitrary patterns of activation over the N signal nodes, yielding 2^N possible signals. The network's behaviour while producing signals for a given meaning and arriving at the interpretation of a particular received signal are determined by the single layer of connection weights in the network, connecting all nodes in N with all nodes in M.

At each generation each individual in the generation $g + 1$ receives t exposures to the communicative behaviour of generation g individuals, where each exposure consists of an observation of a single meaning-signal pair. Generation $g + 1$ individuals then receive a further $t/2$ exposures to the communicative behaviour of other generation $g + 1$ individuals. This model therefore exhibits a degree of hybridization between the Iterated Learning and Negotiation models. However, it is more appropriate to classify the model as of the Iterated Learning type, as the exposures to the previous generation's communicative behaviour occur first and will have the greatest impact. The model outlined in Section 3.3 also suggests that the behaviour of the model would be qualitatively similar if the negotiation portion of learning were omitted.

[1] Livingstone & Fyfe (1999) actually present the network as one which maps from input meanings to output signals. However, during learning signals are treated as input and meanings as output. This turns out to be the key factor in understanding the behaviour of the model.

At each training episode the learner is presented with a signal-meaning pair. The learner takes the signal as input and produces a pattern of activation over the output meaning nodes, x', representing that individual's interpretation of that signal. The teacher's meaning x is then used to perform weight adjustment according to:

$$\Delta w_{ij} = \eta \left(x_i - x'_i \right) y_i$$

where w_{ij} is the weight of the connection between input node i and output node j, y_i and x_i gives the activation levels of the ith input and output unit respectively and η is the learning rate. This type of network model is typically referred to as an Obverter network (as is the network described in Batali (1998), discussed in Section 2.3.3.4), by analogy with the Obverter learning strategy of Oliphant & Batali (1997) — observed production behaviour is used to acquire reception behaviour.

Livingstone and Fyfe report that, over time, populations of such agents converge on shared, stable mappings between meanings and signals which would be optimal in terms of the communicative accuracy measures used by Oliphant & Batali. What drives the emergence of this optimal vocabulary system? As with the models of Hutchins & Hazelhurst (1995), Oliphant & Batali (1997) and Batali (1998), the learning bias of these agents results in directly biased transmission, with the learners happening to favour communication systems which are optimal in terms of communicative accuracy. A discussion of the nature of this bias is postponed until later in this Chapter.

3.2 The communication model

A communication system C consists of a *production* function $p(m)$, mapping from unstructured meanings m to unstructured signals s, and a *reception* function $r(s)$, mapping from signals s to meanings m. m and s are selected such that $m \in \mathcal{M}$ and $s \in \mathcal{S}$ where $\mathcal{M} = \left\{ m_1, m_2 \ldots m_{|\mathcal{M}|} \right\}$ and $\mathcal{S} = \left\{ s_1, s_2 \ldots s_{|\mathcal{S}|} \right\}$. This simple model is suitable for studying the emergence of conventionalised symbolic vocabulary.

How can we evaluate the communicative accuracy of a population using such a communication system? The accuracy of a single communicative event involving a producer P with production function $p(m)$, a receiver R with reception function $r(s)$ and a meaning $m_i \in \mathcal{M}$, $ca(P, R, m_i)$, is defined as:

$$ca\left(P,R,m_i\right) = \left\{ \begin{array}{ll} 1 & \text{if } r\left(p\left(m_i\right)\right) = m_i \\ 0 & \text{otherwise} \end{array} \right\}$$

When $ca\left(P,R,m_i\right) = 1$ the communication is successful. A population's communicative accuracy can be estimated by taking the average $ca\left(P,R,m_i\right)$ for a random sample of P, R and m_i. In a population possessing an *optimal communication system* $ca\left(P,R,m_i\right) = 1$ for any choice of P, R and m_i. This method of measuring communicative accuracy is adopted in Section 3.3.

Equivalently, if the production function $p(m)$ is viewed as a probabilistic function $p(s_j|m_i)$, which gives the probability of producing signal s_j given meaning m_i, and the reception function $r(s)$ is similarly viewed as a probabilistic function $r(m_i|s_j)$ then the communicative accuracy between two individuals with respect to a single meaning, $ca\left(P,R,m_i\right)$, is given by:

$$ca(P,R,m_i) = \sum_{j=1}^{j=|\mathcal{S}|} p(s_j|m_i) \cdot r(m_i|s_j)$$

The communicative accuracy of P and R over all meanings, $ca(P,R)$ can then be defined as the average of their communicative accuracy over each meaning $m_i \in \mathcal{M}$ e.g.

$$ca\left(P,R\right) = \frac{\sum_{i=1}^{i=|\mathcal{M}|} \sum_{j=1}^{j=|\mathcal{S}|} p\left(s_j|m_i\right) \cdot r\left(m_i|s_j\right)}{|\mathcal{M}|}$$

In a population possessing an optimal communication system $ca(P,R) = 1$ for any choice of P and R. This method of evaluating communicative accuracy is more appropriate for the model outlined in Section 3.4.

3.3 Model 1: a feedforward network model

In this Section a simple ILM is described, which is designed to allow the investigation of the impact of learner bias and natural selection of cultural variants on emergent communication systems. This model is based on my undergraduate dissertation (Smith 1998), and has been published in more recent form in Smith (in press).

In this ILM, communicative agents are modelled using feedforward neural networks. Neural networks were chosen for several reasons. Firstly, there is some tradition of using neural networks in research on the evolution of communication — neural networks of some form are used by Batali (1994), Hutchins & Hazelhurst (1995), Batali (1998), Cangelosi & Parisi (1998), Cangelosi (1999), Livingstone & Fyfe (1999) and Kirby & Hurford (2002). Continuing this tradition provides several benefits. In particular, using a similar model allows the results of this research to be more easily related to previous research and the generality of the results of earlier simulations to be tested.

Secondly, well-established mechanisms exist for training neural networks to learn input-output mappings (i.e. backpropagation). Using an established learning mechanism reduces the amount of novel elements contained in the model, as well as allowing our understanding of that mechanism to be expanded.

Finally, using neural networks allows both genetically-transmitted and culturally-transmitted information to influence, in principle, the eventual behaviour of agents in the model. This will prove useful in Chapter 4, when I will consider dual-transmission models.

3.3.1 The communicative agent

The model of a communication system is as described above in Section 3.2. Communicative agents must be capable of representing, using and learning such systems.

3.3.1.1 Representation

Feedforward neural networks are used to model communicative agents. Each individual is modelled using a single network mapping between meanings and signals. There are two possible types of networks: one which takes a representation of a meaning as input and produces a representation of a signal as output, and one which takes a signal as input and produces a meaning as output. The structure of the two networks are shown in Figure 3.1. Feedforward networks mapping from input meanings to output signals will be termed *imitator* networks, whereas networks mapping from input signals to output meanings will be referred to as *obverter* networks. The precise nature of the meaning-signal mapping in these networks is determined by the network connection weights.

Given that the input and output layers in these networks have three nodes, communication systems are mappings between three-dimensional meaning vectors and three-dimensional signal vectors. Binary vectors are used, giving 2^3 possible meanings and 2^3 possible signals. A subset of the set of possible meaning vectors, $\mathcal{M}_{CRS}$, are considered

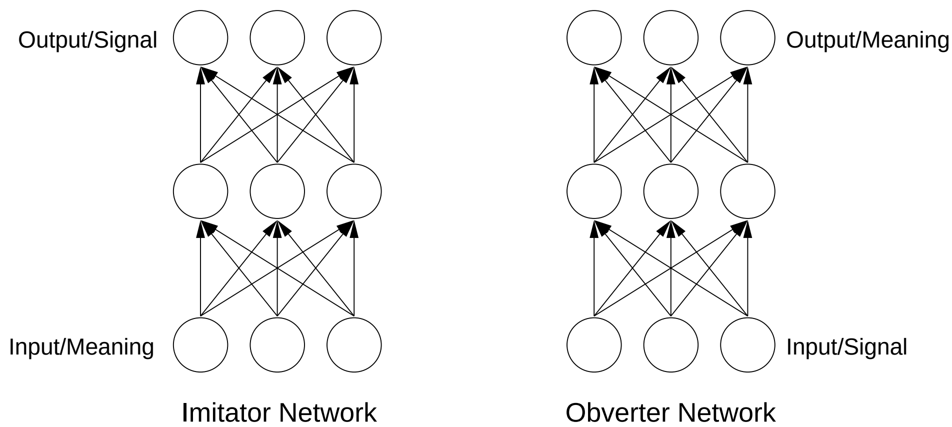

Figure 3.1: The two network architectures. Imitator networks map from input meanings to output signals. Obverter networks map from input signals to output meanings.

to be communicatively relevant situations, meaning the agents are required to communicate about them. For all simulations outlined in this section, $\mathcal{M}_{CRS}$ consists of the unit meanings represented by the vectors $(1\ 0\ 0)$, $(0\ 1\ 0)$ and $(0\ 0\ 1)$. The full set of possible signals are allowed. Therefore $|\mathcal{M}| = 8$, $|\mathcal{M}_{CRS}| = 3$ and $|\mathcal{S}| = 8$.

3.3.1.2 Production and reception

These neural network models of agents embody the production and reception functions $p(m)$ and $r(s)$. Deriving $p(m)$ from an imitator network or $r(s)$ from an obverter network is straightforward. For imitator networks $p(m)$ is derived by presenting the pattern of activation corresponding to each $m \in \mathcal{M}_{CRS}$ to the network, propagating activations forward through the network and thresholding the resultant real-valued output pattern of activation to give the signal $s \in \mathcal{S}$ associated with the given meaning. Similarly, for obverter networks $r(s)$ is derived by presenting the pattern of activation corresponding to each $s \in \mathcal{S}$ to the network, propagating activations forward through the network and thresholding the resultant real-valued output pattern of activation to give the meaning $m \in \mathcal{M}$ associated with the received signal.

Reception for imitator networks and production for obverter networks is slightly more complex, given that the networks are not bidirectional. To derive $r(s)$ for an imitator network each signal $s \in \mathcal{S}$ is considered in turn. All $m \in \mathcal{M}_{CRS}$ are propagated through a given agent's network to produce a real-numbered output pattern of activation for each meaning. Each output pattern is given a confidence rating, corresponding to how closely that pattern matches the signal currently under consideration, s. The meaning which

Process	Network Type	
	Imitator	Obverter
Production	propagate	confidence measure
Reception	confidence measure	propagate

Table 3.1: A summary of the production and reception procedures for the two types of networks. Imitators produce a signal for a given meaning by propagating activations forward through their network, and arrive at a meaning given a received signal using the confidence measure. Obverter networks produce using the confidence measure process, and receive by propagating activations.

yields the real-numbered output vector closest to s, according to the confidence measure, is chosen as the interpretation of s. This method is based on the method used by Batali (1998) and Kirby & Hurford (2002) for producing outputs for similar networks. Similarly, to derive an obverter network's $p(m)$ each meaning $m \in \mathcal{M}_{CRS}$ is considered in turn. All $s \in \mathcal{S}$ are propagated through a given agent's network to produce a real-numbered output pattern of activation for each signal. Each output pattern is given a confidence rating, corresponding to how closely that pattern matches the meaning currently under consideration, m. The signal which yields the real-numbered output closest to m, according to the confidence measure, is chosen as the network's production for m.

The confidence measure that a given real-numbered output vector, o, of length n matches a target binary vector t of length n is given by $C(t, o)$. $C(t, o)$ is simply the product of the confidence scores for each individual node $1...n$ in the output vector i.e.

$$C(t[1 \ldots n], o[1 \ldots n]) = \prod_{i=1}^{n} C(t[i], o[i])$$

where the confidence measure for node i is

$$C(t[i], o[i]) = \begin{cases} o[i] & \text{if } t[i] = 1, \\ (1 - o[i]) & \text{if } t[i] = 0. \end{cases}$$

(Equations adapted from Kirby & Hurford (2002))

The production and reception processes for both types of networks are summarised in Table 3.1.

The deterministic nature of these networks during production means that a definition of ambiguity for communication systems can be formally stated. Communication systems used by neural networks will be termed:

- *Unambiguous* if $p(m)$ is a one-to-one function.
- *Partially ambiguous* if $p(m)$ is a many-to-one function, but the range of $p(m)$ is not a singleton set.
- *Fully ambiguous* if the range of $p(m)$ is a singleton set.

3.3.1.3 Learning

In common with most implementations of the NM and ILM, I assume here that individuals learn from observed meaning-signal pairs. Well-established procedures exist for training feedforward networks to associate pairs of input-output pairs — here I use the backpropagation method (Rumelhart *et al.* 1986). For imitator agents, the training process involves attempting to associate an input meaning with an output signal. Imitators are therefore learning their production function on the basis of observed production behaviour. Obverter networks learn to associate input signals with output meanings — obverters learn their reception function on the basis of observed production behaviour, as in Batali (1998), Livingstone & Fyfe (1999) and others.

3.3.2 The Iterated Learning Model

As discussed in Chapter 2, the results of repeated cultural transmission can be investigated using an Iterated Learning Model. In an ILM agents acquire their competence through learning from observations of the behaviour of other agents. This competence is then used to generate behaviour which is observed in turn by other agents. In the case of this model, the culturally-transmitted behaviour of interest is a communication system.

The process of iterated learning requires a model of population turnover. In this model I use a generational population turnover model, illustrated in Figure 2.6 (a) in Chapter 2. At every time-step a new population of a certain size is created. The pre-existing population produces some observable behaviour and the members of the new population observe and learn from that behaviour. The pre-existing population is then removed and replaced by the newly-created population and the process repeats.

More formally, the generational ILM consists of an initialisation process and an iteration process:

Initialisation Create a population $population_{g=0}$ of N agents[2]. Each agent is either an imitator or obverter, as described above, with populations being homogeneous in this

[2] $N = 100$ for all ILMs outlined in this section.

respect. Each agent has random initial connection weights in the range $[-1, 1]$. Each agent's communication system is determined by these random initial connection weights.

Iteration

1. Evaluate the communicative accuracy of every member of $population_g$ by evaluating every individual's communicative accuracy as both producer and receiver with two randomly selected partners according to the measure $ca\,(P, R, m)$, for every $m \in \mathcal{M}_{CRS}$.
2. For every member of the population $population_g$, generate a set of meaning-signal pairs by applying the network production process to every $m \in \mathcal{M}_{CRS}$. Noise is added to each meaning-signal pair[3] with probability p_n.
3. Create a new population $population_{g+1}$ of N agents of the same type (imitator or obverter) as $population_g$, where each member of $population_{g+1}$ has random initial connection weights in the range $[-1, 1]$.
4. Each member of $population_{g+1}$ receives e exposures to the observable behaviour generated by $population_g$. During each of these e exposures the new agent observes the complete set of meaning-signal pairs generated by a member of $population_g$ selected randomly from among the t most successful communicators in $population_g$. For each exposure the learner updates their connection weights according to the observed meaning-signal pairs using the backpropagation learning algorithm[4].
5. $population_g$ is removed and replaced with $population_{g+1}$. Return to 1.

Each pass through the iteration process will be termed a *generation.* Note that the selection of individuals to observe depends on t and therefore allows the possibility of *natural selection* of cultural variants, as described by B&R. If $t = N$ then selection of individuals to act as cultural parents is independent of the communicative success of those individuals and there is no natural selection of cultural variants. When $t < N$ the probability of an individual being observed and learned from will depend on their evaluated communicative success, and there will be natural selection, acting on cultural transmission, in favour of communication systems which result in successful communication.

The fact that every individual in a population begins their life with a particular network type (imitator or obverter) and a particular set of connection weights (randomly

[3]In order to add noise to a meaning-signal pair $\langle m_i, s_j \rangle$, s_j is replaced with a randomly-selected $s_k \in \mathcal{S}$, where $k \neq j$.

[4]A learning rate of 0.5 is used

distributed within some range) suggests some kind of innate endowment of these components. It is our goal to investigate the impact of this innate endowment on the communicative behaviour of the population. However, every agent begins life with the *same* endowment – there is no possibility of genetic variation within the population. The emergent behaviour of the population will therefore be determined by the dynamics resulting from the iterated cultural transmission of communication systems among individuals with a common genetic endowment. In Chapter 4 I will investigate how the biological evolution of this innate endowment in a genetically heterogeneous population can impact on the evolution of communication systems.

3.3.3 Network architecture, learning bias and natural selection

The goal of this Chapter of the thesis is to identify the learning mechanisms necessary to create, through cultural processes, a communicatively useful vocabulary. The ILM described above can be used to investigate whether imitator and obverter agents construct an optimal, unambiguous communication system from random initial behaviour, and under what circumstances. To this end, runs of the iterated learning model were carried out. In these simulations, the communication system used by agents in the initial population is dependent on their random connection weights, and is therefore random. 10 runs were carried out for each set of experimental conditions, with runs proceeding for 1000 generations. We are primarily interested in the *end states* of these runs, rather than their progress through time. In order to evaluate the end state communication system in use in the populations, the average communicative accuracy of the population is recorded for the last 10 generations of each run. Each point in the plots that follow therefore represents the average communicative accuracy of 10 populations over a period of 10 generations.

3.3.3.1 Learning bias and no natural selection

Figure 3.2 shows the results for simulation runs for imitator and obverter populations where $t = N$, (every member of the population is a potential cultural parent) for various numbers of learning exposures (e), in the absence of noise on cultural transmission ($p_n = 0$).

The different network architectures clearly result in very different behaviour, when placed in the context of the ILM. For imitator networks, the populations converge on communication systems which result in communicative accuracy of 0.33. This level of communicative accuracy is a consequence of the population using a shared, fully ambiguous

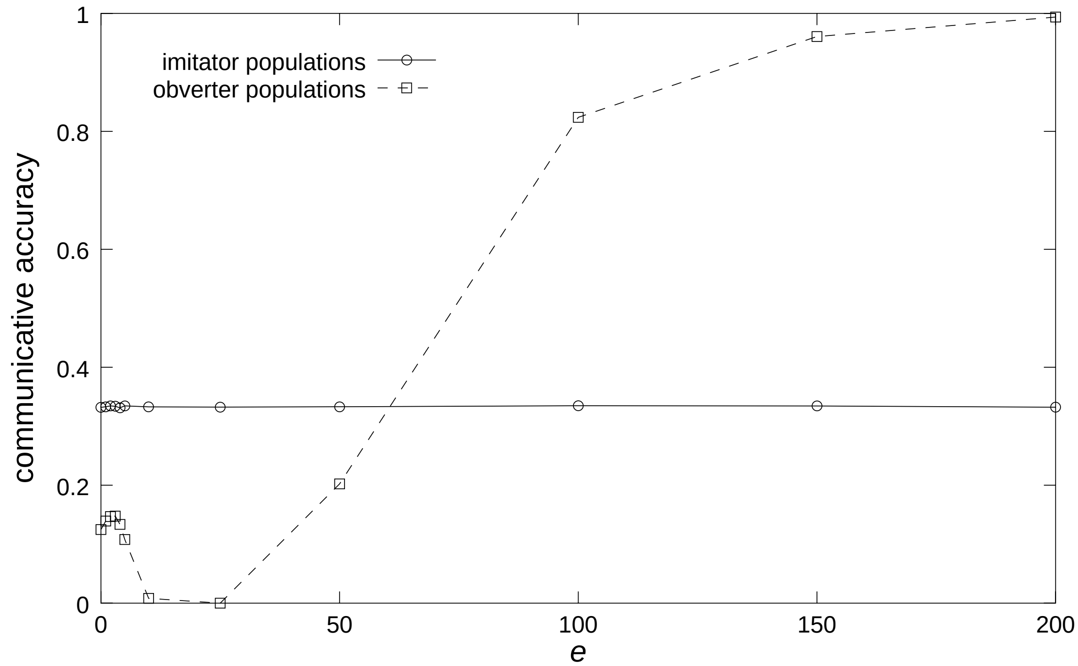

Figure 3.2: The average final communicative accuracy of imitator (solid line) and obverter (dashed line) populations as a function of the number of learning exposures (e) used during the simulation runs. These results are for the case where there is no natural selection of cultural variants ($t = N$). Imitator populations converge on systems which yield chance levels of communicative accuracy, regardless of e. In contrast, given high enough e, obverter populations converge on communication systems which give high levels of communicative accuracy.

communication system. In contrast, obverter populations converge on levels of communicative accuracy close to optimal when e is large — given large enough e, obverter populations converge on a shared, unambiguous vocabulary.

Why do the two different network architectures display this behaviour when placed in the context of the ILM? We can rule out natural selection of cultural variants (because $t = N$). This behaviour therefore must be due to direct bias pressure operating on cultural transmission.

In order to understand the source of this bias, it is necessary to assess the ability of individual agents, in isolation, to acquire systems of various levels of ambiguity, with varying levels of exposure to such systems. The possible range of meaning-signal mappings is actually rather small — $|\mathcal{M}_{CRS}| = 3$ and $|\mathcal{S}| = 8$ gives $|\mathcal{S}|^{|\mathcal{M}_{CRS}|} = 512$ possible meaning-signal mappings. It is therefore possible to examine the ability of agents to acquire every possible system. Of the 512 possible systems, 8 are fully ambiguous, 168 are partially ambiguous and 336 are unambiguous. For each system, 100 networks with random initial weights in the range $[-1, 1]$ were given e exposures to that system of meaning-signal

e	System Type		
	Fully Ambiguous	Partially Ambiguous	Unambiguous
1	25.4	0.3	0.0
2	47.1	0.8	0.0
3	74.0	1.4	0.0
4	91.8	1.4	0.1
5	98.0	1.7	0.0
10	100.0	0.5	0.0
25	100.0	1.6	0.1
50	100.0	34.2	13.1
100	100.0	92.8	82.9
150	100.0	99.4	98.3
200	100.0	100.0	99.8

Table 3.2: The imitator learning bias. The table shows the percentage of imitator networks which succeed acquiring languages of the various classifications, according to e, the number of exposures to the system. For imitator networks, fully ambiguous systems are easier to learn.

e	System Type		
	Fully Ambiguous	Partially Ambiguous	Unambiguous
1	0.1	0.2	0.3
2	0.0	0.3	0.3
3	0.1	0.3	0.5
4	0.0	0.5	0.5
5	0.1	0.4	0.7
10	0.0	0.9	1.6
25	0.0	3.7	7.8
50	0.0	10.1	26.2
100	0.0	14.9	49.1
150	0.0	15.2	53.8
200	0.0	15.2	54.8

Table 3.3: The obverter learning bias. The table shows the percentage of obverter networks which succeed acquiring languages of the various classifications, according to e, the number of exposures to the system. Obverter networks find unambiguous systems easier to learn.

mappings. Learning proceeds via the backpropagation process, with the same learning rate as used in the ILM. A network was judged to have learned a system successfully if the observed system could be reproduced in production — for every meaning-signal pair $\langle m_i, s_j \rangle$ production of the signal associated with m_i resulted in s_j being produced. The results are summarised in Tables 3.2 and 3.3 by communication system type.

As can be seen from Table 3.2, for imitator agents systems exhibiting a higher degree of ambiguity are easier to acquire than systems exhibiting a lower degree of ambiguity, for

System Type	% population
Unambiguous	2
Partially Ambiguous	25
Fully Ambiguous	73

Table 3.4: The behaviour of imitator agents with random connection weights. The table shows the percentage (based on 1000 test networks) of imitator networks with random connection weights (in the range [-1,1]) who use a communication system of the given type. Random imitator networks tend to produce fully ambiguous systems.

all values of e. Table 3.3 shows that obverter agents have the opposite learning bias — systems exhibiting lower degrees of ambiguity are easier to acquire, for all values of e. Learnability never reaches 100%, even for unambiguous communication systems. It appears that certain unambiguous systems are unlearnable by obverter agents, while certain unambiguous systems are 100% learnable. The key point is that certain unambiguous systems are highly learnable whereas partially ambiguous and fully ambiguous systems are less learnable.

Returning to the results for ILM runs involving imitator populations, for low values of e, no communication system can reliably be learned. Populations essentially behave in a random fashion. The typical random behaviour of imitator agents is shown in Table 3.4. The majority of individuals use fully ambiguous systems, resulting in chance levels of communicative accuracy. As e increases, fully ambiguous systems rapidly become highly learnable, and are always more learnable than less ambiguous systems. Less ambiguous systems are less likely to be successfully learned than fully ambiguous systems, and are unstable over time. The populations therefore converge on fully ambiguous systems, resulting in low levels of communicative accuracy.

In contrast, the communicative accuracy in obverter populations increases as e increases. For low values of e all systems are unlearnable, and individuals use a random system. As shown in Table 3.5, obverter agents with random connection weights tend to use a unambiguous systems. The communicative accuracy of the population is therefore low, as there are a large number of uncoordinated unambiguous systems present. As e increases, the learnability of unambiguous systems increases, and is always higher than the learnability of more ambiguous systems. Unambiguous systems become increasingly stable relative to more ambiguous systems and the populations converge on shared unambiguous communication systems, resulting in high levels of communicative accuracy for high values of e.

System Type	% population
Unambiguous	65
Partially Ambiguous	33
Fully Ambiguous	2

Table 3.5: The behaviour of obverter agents with random connection weights. The table shows the percentage (based on 1000 test networks) of obverter networks with random connection weights (in the range [-1,1]) who use a communication system of the given type. Random obverter networks tend to produce unambiguous systems.

The behaviour of these populations in the ILM is therefore determined by the learning biases of the two network architectures. These learning biases result in direct bias acting on cultural transmission, with the cultural variants favoured by the bias eventually reaching fixation.

3.3.3.2 Learning bias and natural selection

Figures 3.3 and 3.4 show the results for simulation runs for imitator and obverter agents where $t < N$ (only the top t individuals act as cultural parents, and t is less than the population size N, therefore the less able communicators may not act as cultural parents), for various values of t and e (learning exposures), again in the absence of noise ($p_n = 0$). In these simulations there are two pressures operating on the communication systems in the populations:

1. *Selection for learnability*, driven by the agents' learning bias, favouring either more ambiguous communication systems (in the case of imitator agents) or less ambiguous systems (in the case of obverter agents).
2. *Selection for communicative success* driven by natural selection of communication systems, favouring systems which result in successful communication.

In populations of imitator agents pressures 1 and 2 are in conflict, with selection for learnability favouring fully ambiguous systems (as discussed in the previous section), while natural selection favours shared unambiguous systems. For obverter populations these pressures are not in conflict, with both favouring the development of shared unambiguous systems.

The addition of natural selection of cultural variants has little impact on the emergent communication systems — as with the case where there is no natural selection, imitator populations converge on fully ambiguous communication systems and communicative accuracy remains uniformly low, while obverter populations converge on unambiguous

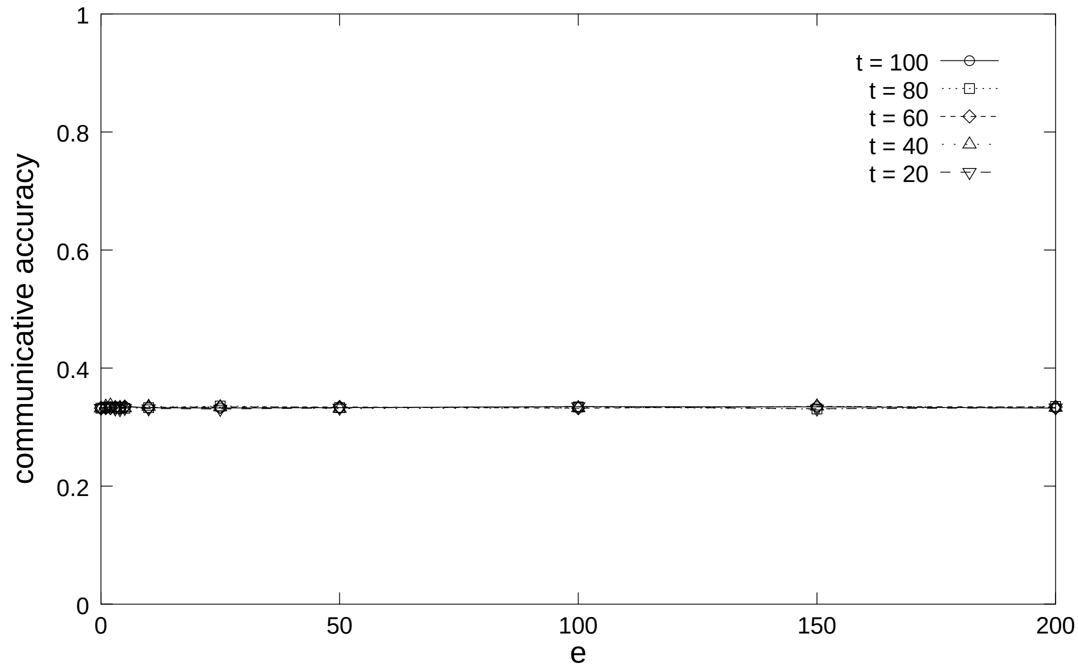

Figure 3.3: The average final communicative accuracy of imitator populations where there is natural selection of cultural variants ($t \leq N$), and no noise on cultural transmission ($p_n = 0$), as a function of e. Natural selection of cultural variants clearly has no impact.

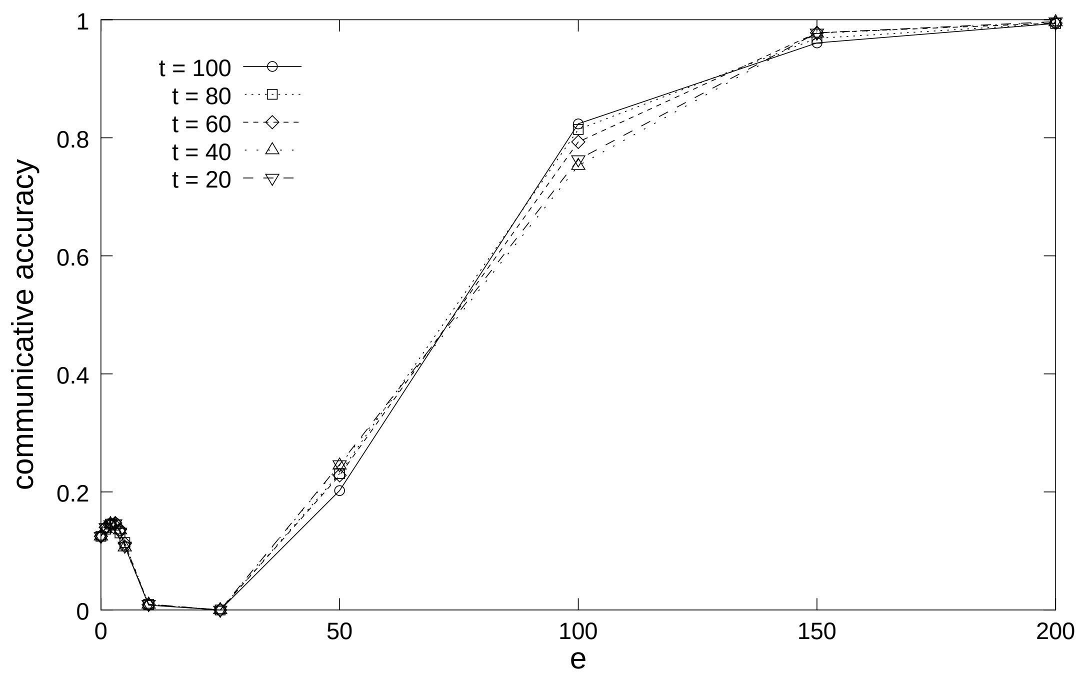

Figure 3.4: The average final communicative accuracy of obverter populations where there is natural selection of cultural variants ($t \leq N$), and no noise on cultural transmission ($p_n = 0$), as a function of e. Natural selection of cultural variants has little impact — there are very slight differences in final communicative accuracy, dependent on t.

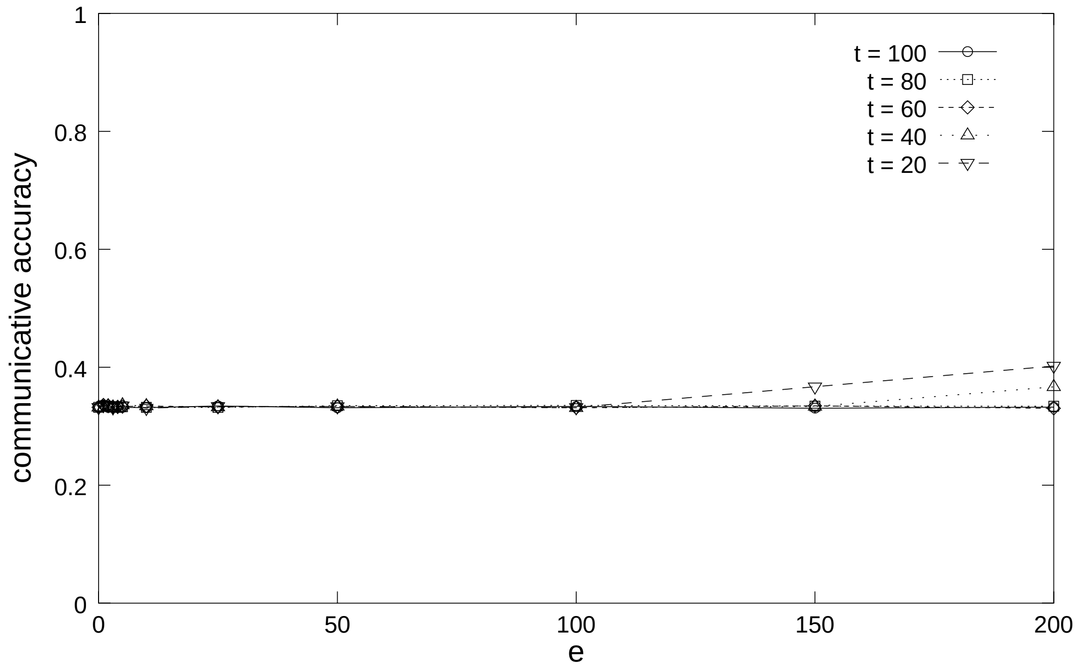

Figure 3.5: The average final communicative accuracy of imitator populations where there is natural selection of cultural variants ($t \leq N$), and noise on cultural transmission ($p_n = 0.05$), as a function of e. Natural selection of cultural variants clearly a slight impact for high e.

communication systems given sufficiently high e, with consequently high communicative accuracy. The behaviour of the populations is still dominated by the intrinsic learning bias of the agents.

B&R highlight the importance of cultural variation in populations where cultural transmission is undergoing natural selection — where there is no variation, natural selection is powerless. It is possible that we are not seeing any impact from natural selection due to a lack of cultural variability in the populations. While the initial populations exhibit variability (see Tables 3.4 and 3.5), which direct bias clearly feeds off, it could be that biased cultural transmission eliminates this variability too quickly, preventing natural selection of cultural variants from functioning. In order to investigate this possibility, the experiments outlined above were repeated with noise on cultural transmission ($p_n = 0.05$). This noise will potentially introduce cultural variation, which natural selection can then feed off. The results are plotted in Figures 3.5 and 3.6.

The introduction of noise has a slight impact. In imitator populations, when e is very high, there is a slight increase in average communicative accuracy for the case where $t = 20$. This is in fact due to one (when $e = 150$) or two (when $e = 200$) of the ten runs converging on partially ambiguous communication systems. This only occurs when

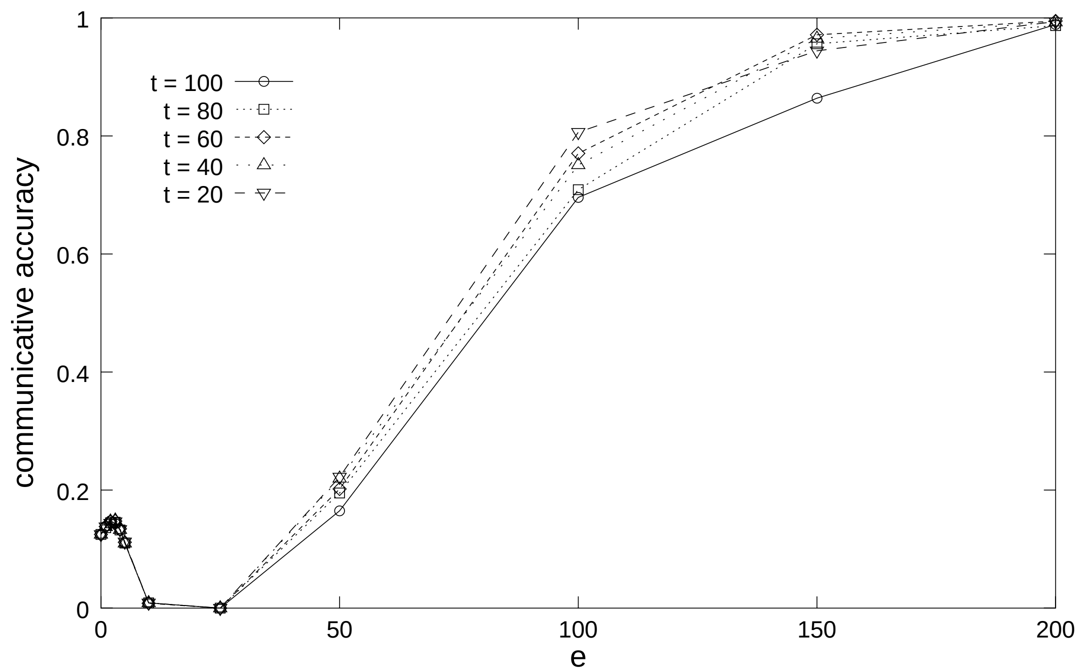

Figure 3.6: The average final communicative accuracy of obverter populations where there is natural selection of cultural variants ($t \leq N$), and noise on cultural transmission ($p_n = 0.05$), as a function of e. Natural selection of cultural variants has a noticeable impact for $e = 100$ or 150.

e is very high as this is the point where the individual's learning bias is weakest — as can be seen from Table 3.2, partially ambiguous systems seem to be as learnable as fully ambiguous systems where $e \geq 150$. However, the natural selection of cultural variants needs to be very severe to produce this slight effect.

In obverter populations there is, similarly, a slight impact, with natural selection of cultural variants improving the populations' communicative accuracy somewhat for certain amounts of learning ($e = 100$ or 150). The effect is still fairly minor.

3.3.4 Summary

The two distinct models of communicative agents have different learning biases, as highlighted by the acquisition tests outlined in Section 3.3.3.1. Placing these agents within an ILM allows the consequences of the iterated application of these learning biases to be explored. In the case where the direct bias on cultural transmission introduced by the agents' learning bias is the sole factor at play, the populations converge on the type of communication system favoured by that bias, as predicted by B+R. Natural selection of cultural variants in addition to this biased transmission has a very minor impact, and even then only when noise is injected into the system to provide variation.

Why do imitator and obverter agents have different biases and how can this bias best be described? This issue will be returned to in Section 3.5. However, it is clear from the relatively simple experiments outlined in this section that the behaviour of populations of individuals is strongly determined by their learning bias, which can override other pressures acting on cultural transmission, such as natural selection. Comparison of more than two alternative learning biases remains desirable, and a model allowing the comparison of a much wider range of biases is outlined in the next section. The feedforward neural network model is returned to in Chapter 4, where it is used to investigate the interactions between genetic and cultural transmission of communication.

3.4 Model 2: an associative network model

The feedforward network model described above is limited in the sense that the learning bias of individual agents is a consequence of their network architecture, and there are only two such possible architectures — the imitator architecture, and the obverter architecture. Ideally we would like to be able to experiment with a wider range of learning biases, in order to isolate the elements of bias which drive the cultural evolution of symbolic vocabulary.

A promising approach to addressing precisely this question is outlined in Oliphant (1999). Oliphant investigates how different learning rules influence the development of a vocabulary system through cultural processes within a population of associative networks. While the approach described in this paper is promising, its execution suffers from several shortcomings. Firstly, only three possible learning rules are considered. Secondly, while it is shown that certain learning rules result in the emergence of optimal communication, the properties of the learning rules that result in this behaviour are not explicitly identified. Thirdly, the results for those three learning rules are not related to other results in the field.

In this section I introduce a model, based on Oliphant's, which allows a wide range of learning rules, and associated learning biases, to be explored. This exploration allows me, in Section 3.5, to identify the key bias leading to the emergence of communicatively optimal, symbolic vocabulary through cultural processes. This bias can also be identified in the feedforward network model described above, and in most other models of the emergence of vocabulary via cultural evolution.

3.4.1 Communicative agents

The model of communication is as outlined in Section 3.2. Given the nature of the communicative agent model, the probabilistic interpretation of the communicative accuracy function is more natural, and is used in this section.

An associative network is used to model communicative agents. Since this model is less standard than the feedforward network model outlined in Section 3.3 a detailed and somewhat formal description is given here.

3.4.1.1 Representation

Agents are modelled using networks consisting of two sets of nodes $\mathcal{N}_M$ and $\mathcal{N}_S$ and a set of weighted bidirectional connections $\mathcal{W}$ connecting every node in $\mathcal{N}_M$ with every node in $\mathcal{N}_S$.

Patterns of activation over $\mathcal{N}_M$ are considered to represent meanings, whereas patterns of activation over $\mathcal{N}_S$ are considered to be signals. Restricting these patterns of activation to contain a single active unit yields $|\mathcal{N}_M|$ orthogonal meaning representations and $|\mathcal{N}_S|$ orthogonal signal representations, suitable for representing sets of unstructured meanings and unstructured signals such as those described in Section 3.2. If Gi is the ith node from the set $\mathcal{N}_G$ and the activation of node Gi is a_{Gi} then the meaning m_i corresponds to a pattern of activation over $\mathcal{N}_M$ where $a_{Mi} = 1$ and $a_{M(j \neq i)} = 0$. Similarly, the signal s_i corresponds to a pattern of activation over $\mathcal{N}_S$ where $a_{Si} = 1$ and $a_{S(j \neq i)} = 0$. This representational scheme is illustrated in Figure 3.7.

3.4.1.2 Production and reception

Patterns are retrieved from the network using a k-winners-take-all strategy. In order to retrieve a pattern of activation over nodes in $\mathcal{N}_S$ based on an input pattern of activation over nodes in $\mathcal{N}_M$ the weighted sum of inputs to node Si, q_{Si}, for each $Si \in \mathcal{N}_S$ is calculated according to the formula:

$$q_{Si} = \sum_{j=1}^{j=|\mathcal{N}_M|} a_{Mj} \cdot w_{Mj,Si}$$

where $w_{a,b} \in \mathcal{W}$ is the weight of the connection between nodes a and b. The k nodes in $\mathcal{N}_S$ with the highest values of q then have their activations set to 1, while all other nodes in $\mathcal{N}_S$ have their activations set to 0. If several nodes have equal q a random winner is selected from among them. Patterns of activation over the nodes in $\mathcal{N}_M$ are retrieved

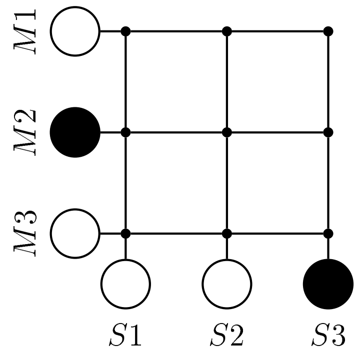

Figure 3.7: A neural network where $|\mathcal{N}_M| = |\mathcal{N}_S| = 3$. Large filled circles represent nodes with activation of 1, large empty circles represent nodes with activation of 0. The pattern of activation over $\mathcal{N}_M$ therefore represents the meaning m_2 ($a_{M2} = 1$, $a_{M1} = a_{M3} = 0$). Similarly, the pattern of activation over $\mathcal{N}_S$ represents the signal s_3

based on input patterns of activation over $\mathcal{N}_S$ in exactly the same way. For all simulations outlined in this paper, $k = 1$ — retrieved patterns of activation only ever consist of a single active node and $(|\mathcal{N}| - 1)$ non-active nodes. This ensures that retrieved patterns of activation conform to our representation of meanings and signals outlined above. This retrieval process is illustrated in Figure 3.8.

Retrieving a pattern of activation over $\mathcal{N}_S$ given an input pattern of activation over $\mathcal{N}_M$ corresponds to retrieving the signal associated with a given meaning — *production* of a signal associated with a given meaning. Retrieving a pattern of activation over $\mathcal{N}_M$ given an input pattern of activation over $\mathcal{N}_S$ corresponds to retrieving the meaning associated with a given signal — *reception* of a given signal and interpretation of that signal to yield a meaning. Note that the production and reception behaviour of such networks are not necessarily closely related — for example, the network in Figure 3.8 would produce $s2$ when prompted with $m2$, but would interpret $s2$ as meaning $m3$. Using a single network for both production and reception, as opposed to two separate networks, does however allow the possibility of a coupling of production and reception.

3.4.1.3 Learning

In order to store the association between patterns of activation over $\mathcal{N}_M$ and $\mathcal{N}_S$ the activations of the nodes in $\mathcal{N}_M$ and $\mathcal{N}_S$ are set to the required values and the weights of the connections in $\mathcal{W}$ are adjusted according to some weight-update rule W. If we assume that W must only adjust connection weights based on local information and that all patterns of activation will be binary, W can be specified by the 4-tuple $(\alpha\ \beta\ \gamma\ \delta)$,

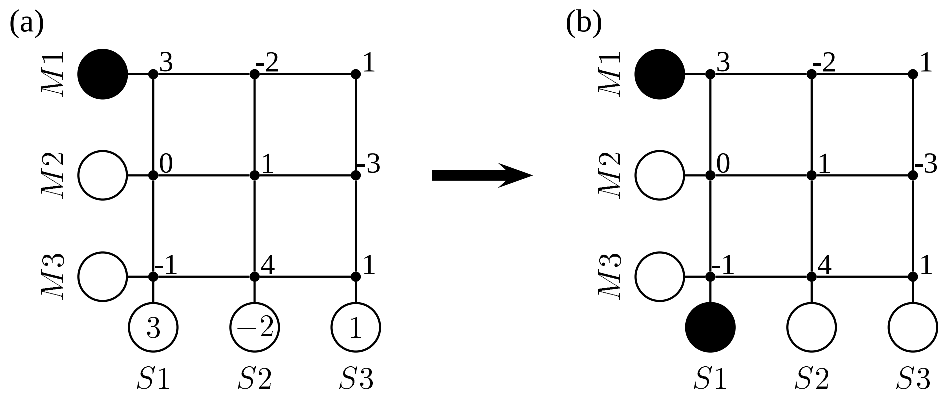

Figure 3.8: Retrieval of a pattern of activation over $\mathcal{N}_S$ based on a pattern of activation over $\mathcal{N}_M$. As before, large filled circles represent nodes with activation of 1. Connections between nodes are represented by the intersections of connecting lines and have an associated weight. In (a), the nodes in $\mathcal{N}_M$ have been set to a pattern of activation, resulting in a pattern of weighted sums of inputs over the nodes in $\mathcal{N}_S$ — the q values for those nodes The numbers in the centre of the nodes in $\mathcal{N}_S$ represent the weighted sums to those nodes. In (b) the result of the application of the winner-take-all process is shown — q_{S1} is greater than q_{S2} or q_{S3}, therefore node $S1$ has its activation set to 1 while nodes $S2$ and $S3$ have their activations set to 0.

where the value in α specifies how the weight of connection $w_{i,j}$ should be adjusted when $a_i = a_j = 1$, the value in β specifies how $w_{i,j}$ should be adjusted when $a_i = 1$ and $a_j = 0$, the value in γ specifies how $w_{i,j}$ should be adjusted when $a_i = 0$ and $a_j = 1$ and the value in δ specifies how $w_{i,j}$ should be adjusted when $a_i = a_j = 0$. While weights could be adjusted in many ways we will restrict ourselves here to the simplest case where α, β, γ and δ must take integer values in the range $[-1, 1]$. This yields a range of $3^4 = 81$ possible weight-update rules.

Given our interpretations of patterns of activations of $\mathcal{N}_M$ and $\mathcal{N}_S$ this storage process represents the process of learning the association between a meaning and a signal in a meaning-signal pair $\langle m, s\rangle$ according to some rule W. The learning process is illustrated in Figure 3.9.

3.4.2 Acquisition of an optimal system

We now have a model of communication, a model of an agent and processes of production, reception and learning. The feedforward neural network model highlighted the importance of the learning biases of agents when accounting for the behaviour of populations of such agents. The first question to be addressed here is therefore to ask whether individual agents, in isolation, can acquire an optimal communication system. To this end an unambiguous set of meaning-signal pairs $\mathcal{A} = \{\langle m_1, s_1\rangle, \langle m_2, s_2\rangle \ldots \langle m_{10}, s_{10}\rangle\}$

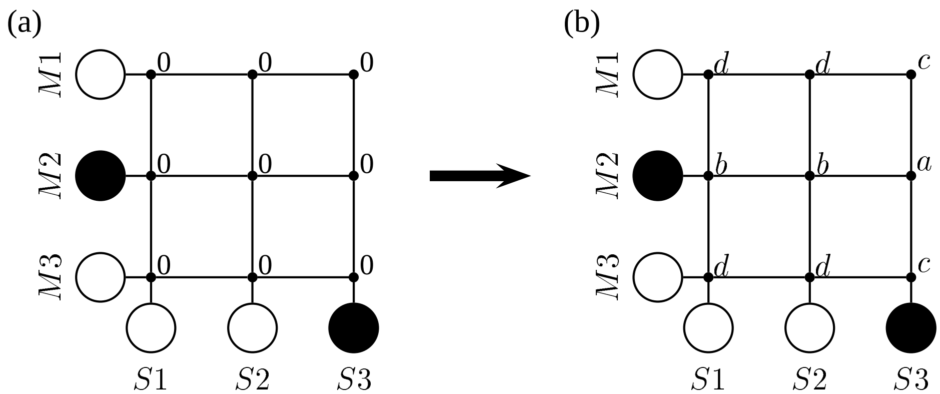

Figure 3.9: Storage of the meaning-signal pair $\langle m_2, s_3 \rangle$ using the weight-update rule $W = (a\ b\ c\ d)$. In (a), the nodes in $\mathcal{N}_M$ and $\mathcal{N}_S$ have been set to the patterns of activation representing m_2 and s_3. All connections have weight 0. In (b) the result of the application of the storage process is shown — all connections now have weights of a, b, c or d, depending on the activations of the nodes they connect.

was constructed. Agents using each of the 81 possible weight-update rules were then trained on $\mathcal{A}$, by storing each meaning-signal pair in $\mathcal{A}$ in their network. The agents were then evaluated to see if they had successfully acquired an optimal communication system based on exposure to the unambiguous set of meaning-signal pairs $\mathcal{A}$. Agents are judged to have acquired an optimal system, if, for every $\langle m_i, s_i \rangle \in \mathcal{A}$ both:

1. Production of the signal associated with m_i always[5] results in s_i being produced, i.e. $\langle m_i, s_i \rangle$ can be reproduced in production *and*
2. reception of s_i always results in the interpretation m_i, i.e. $\langle m_i, s_i \rangle$ can be reproduced in reception, meaning that the agent would communicate optimally with itself or another agent using the same weight-update rule exposed to $\mathcal{A}$.

The 81 weight-update rules can therefore be classified according to a [$\pm$learner] feature. 31 of the 81 possible weight-update rules were judged to be capable of acquiring the optimal communication system and were classified as [$+$learner]. The remaining 50 weight-update rules were classified [$-$learner].

[5]The term "always" has to be introduced to account for the stochastic nature of the behaviour of some networks, resulting from multiple nodes in the network receiving the same weighted sum of inputs on presentation of a pattern. In practice, "always" was reduced to "for every one of 1000 trials".

3.4.3 The Iterated Learning Model

As with the feedforward network model discussed above in Section 3.3, this model of a communicative agent can be slotted into an Iterated Learning Model to evaluate how the different weight-update rules influence the development of communication over time in a population. Unlike the feedforward network model, a *gradual*, rather than generational, population turnover model is used. The gradual population turnover model was preferred to counter the possibility in the new model of "inverting" learners, who learn the opposite communication system to their cultural parents. In a generational ILM populations of such agents would score highly on the intra-generational communicate accuracy measures, but could not be said to have learned the communication system of their cultural parents.

In the gradual population turnover model (see Figure 2.6 (b) in Chapter 2) at every time-step a single agent is selected at random and removed from the population. The remaining members of the population produce some observable behaviour, in the case of this model sets of meaning-signal pairs. A new individual arrives and learns based on observations of the population's observable behaviour, then enters the population. The process then repeats.

More formally, the ILM consists of an initialisation process and an iteration process:

Initialisation Create a population of N agents[6], each using the weight-update rule W and possessing communication system L.

Iteration

1. Select an agent at random from the population and remove it.
2. For every remaining member of the population, generate a set of meaning-signal pairs by applying the network production process to every $m \in \mathcal{M}$. Noise is added to each meaning-signal pair[7] with probability p_n.
3. Create a new agent with connection weights of 0 who uses weight-update rule W.

[6] $N = 100$ for all ILMs outlined in this section.

[7] In order to add noise to a meaning-signal pair $\langle m_i, s_j \rangle$, s_j is replaced with a randomly-selected $s_k \in \mathcal{S}$, where $k \neq j$.

4. The new agent receives e exposures to the population's observable behaviour. During each of these e exposures the new agent observes the complete set of meaning-signal pairs of a randomly selected member of the population and updates their connection weights according to the observed meaning-signal pairs and their weight-update rule W.
5. The new agent joins the population. Return to 1.

Each pass through the iteration process will be termed a *cohort*. Note that the random removal of agents from the population means there is no selection based on communicative ability. As with the feedforward network ILM, the fact that every individual begins its life with a weight-update rule and initial set of connection weights suggests some kind of innate endowment of these components. In the simulations outlined in this section populations are homogeneous with respect to this endowment, and we restrict ourselves to investigating the impact of cultural transmission factors on the emergent communication systems. In Chapter 4 the biological evolution of these innate endowments in a genetically heterogeneous population will be investigated.

3.4.4 Maintenance of an optimal system

The first question to be addressed using the ILM is whether a population of agents possessing a weight-update rule W can maintain an optimal system over time in the presence of a small degree of noise. Recall from the description of the ILM given above that the agents in the initial population use some predefined communication system L. For the experiments outlined in this section, the initial population's set of weights $\mathcal{W}$ were constructed such that the $p(m)$ of the initial L generates the set of meaning-signal pairs $\mathcal{L} = \{\langle m_1, s_1\rangle, \langle m_2, s_2\rangle \ldots \langle m_{10}, s_{10}\rangle\}$ — the initial population shares an unambiguous meaning-signal mapping. ILMs were run with each of the 81 possible learning rules, with noise introduced with probability $p_n = 0.05$ and each individual receiving exposures to the communication systems of three randomly-selected members of the population ($e = 3$). Populations were defined as having *maintained* the initial optimal system if the population's communicative accuracy remained above 0.95 for every cohort of a run.[8] Weight-update rules were classified as [+maintainer] if the optimal system was maintained for each of ten 2000-cohort runs.

The populations exhibited four typical patterns of behaviour, illustrated in Figure 3.10. Populations (a), (b) and (c) in Figure 3.10 have failed to maintain the optimal system

[8] The population's communicative accuracy was estimated by evaluating every individual's average communicative accuracy as both producer and receiver with two randomly selected partners according to the measure $ca(P, R)$ given in Section 3.2, averaging over all individuals in the population.

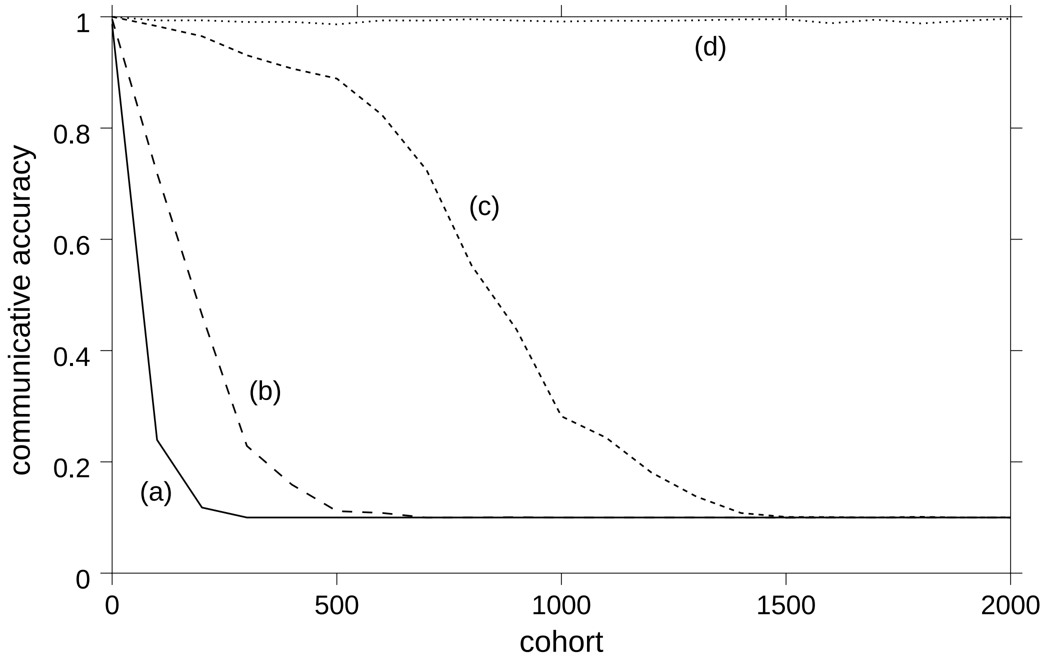

Figure 3.10: Populations of agents using the 81 learning rules exhibit four patterns of behaviour when attempting to maintain an optimal system. This figure plots the communicative accuracy over time of single populations exhibiting these patterns of behaviour: rapid collapse to chance levels of communicative accuracy, as in (a); less rapid collapse to chance levels of communicative accuracy, as in (b) and (c); maintenance of the optimal system, as in (d).

and can therefore be classified as [$-$maintainer], although population (a) in Figure 3.10 exhibits a more rapid decrease in communicative accuracy than populations (b) and (c). Unsurprisingly, all 50 populations using weight-update rules with the [$-$learner] feature followed the pattern of (a) and can therefore be classified [$-$learner, $-$maintainer]. Of the remaining 31 weight-update rules, 13 resulted in the type of pattern exemplified by populations (b) and (c) and can be classified as [$+$learner, $-$maintainer] and 18 resulted in patterns similar to that of population (d) in Figure 3.10 and can be classified as [$+$learner, $+$maintainer].

3.4.5 Construction of an optimal system

Finally, the 81 weight-update rules were examined to see whether they resulted in the emergence of optimal communication systems from random behaviour when placed in the context of the ILM. In the previous section the initial population's communication system, L, was optimal. In the models outlined in this section L has maximum entropy — every $m \in \mathcal{M}$ is associated with every $s \in \mathcal{S}$ with equal probability, $|\mathcal{M}| = |\mathcal{S}| = 10$. This was achieved by setting the connection weights of every individual in the initial population to 0. Unlike in the previous section, cultural transmission is noise-free —

$p_n = 0$ (although results show that similar behaviour occurs with $p_n > 0$). Simulations were run for each of the 81 possible learning rules. A population was defined as having *constructed* an optimal system if the population's communicative accuracy reached 1.0. Weight-update rules were classified [+constructor] if optimal systems were constructed in each of ten 5000-cohort runs.

The populations exhibit three typical patterns of behaviour, of which populations (a), (b) and (c) in Figure 3.11 are representative examples. The populations which fit the pattern exemplified by (a) in Figure 3.11 have clearly failed to construct an optimal system and in fact persist at the random level of performance for $|\mathcal{M}| = |\mathcal{S}| = 10$. All of the weight-update rules which were classified as [$-$maintainer] follow this pattern and can be classified as [$-$constructor].

Populations behaving similarly to population (b) in Figure 3.11 are performing above the random level, but have not constructed an optimal system as defined above. In fact, as suggested for a more limited case by Oliphant (1999), the level of communicative accuracy in these populations hovers around the level we would expect given a random assignment of signals from $\mathcal{S}$ to meanings from $\mathcal{M}$ with replacement:

$$\text{communicative accuracy} \approx 1 - \left(1 - \frac{1}{|\mathcal{S}|}\right)^{|\mathcal{M}|}$$

The reason for this level of performance will be made clear in section 3.5.1. Nine of the 18 weight-update rules which were classified [+maintainer] fit this pattern and can be classified as [$-$constructor].

Populations fitting the pattern exemplified by population (c) in Figure 3.11 have succeeded in constructing an optimal system from random behaviour and can be classified as [+constructor]. Nine of the 18 weight-update rules which were classified as [+maintainer] fit this pattern.

3.4.6 Summary: The classification hierarchy

The three tests outlined above divide the 81 weight-update rules into four groups, summarised in Table 3.6.

The fact that all weight-update rules which are [+constructor] are [+maintainer] and all rules which are [+maintainer] are [+learner] suggests a hierarchy of weight-update rules, summarised in Figure 3.12.

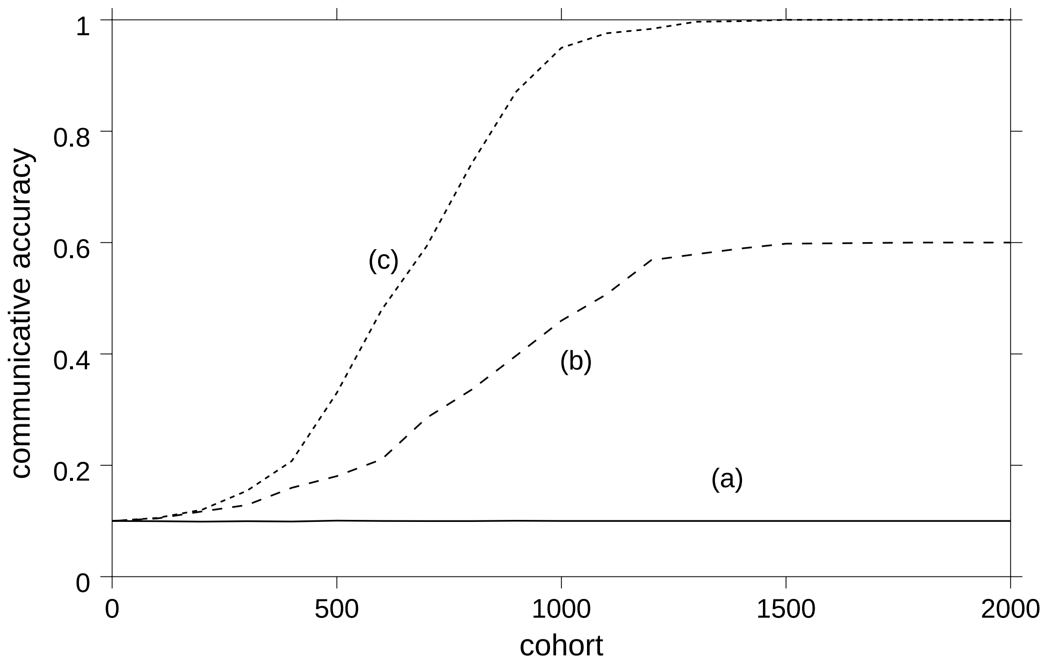

Figure 3.11: Populations of agents using the 81 learning rules exhibit three patterns of behaviour when attempting to construct an optimal system: failure to construct an optimal system and chance-level communicative accuracy, as in (a); failure to construct an optimal system, but levels of communicative accuracy significantly above chance, as in (b); construction of an optimal system, as in (c).

Classification	Number
[−learner, −maintainer, −constructor]	50
[+learner, −maintainer, −constructor]	13
[+learner, +maintainer, −constructor]	9
[+learner, +maintainer, +constructor]	9

Table 3.6: The number of weight-update rules of each particular complete classification, from the sample of 81.

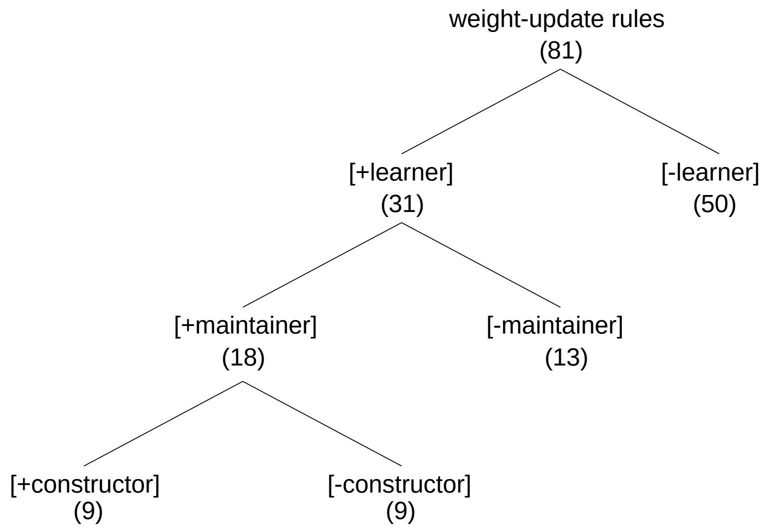

Figure 3.12: The hierarchy of weight-update rules. Read from the top, each node places additional restrictions on the properties of the weight-update rules. The numbers possessing each feature are given in parentheses at each point in the tree.

3.5 The Key Bias

Why are obverter networks in the feedforward network model described in Section 3.3 biased in favour of acquiring unambiguous systems but imitator agents, with a slightly different architecture, are biased in favour of acquiring fully ambiguous systems? Similarly, what is it about the particular assignment of -1s, 0s and 1s to the four conditions α, β, γ and δ in the associative network model[9] (described in the previous Section) that makes one weight-update rule incapable of learning an optimal communication system whereas another weight-update rule is capable of constructing such a system from random behaviour in the context of iterated cultural transmission?

The learning biases of the different network architectures or weight-update rules are best described in terms of the one-to-one nature of mappings between meanings and signals. As defined in Section 3.2, in an optimal communication system $r(p(m)) = m$ for all $m \in \mathcal{M}$. This requires that:

1. Each $m \in \mathcal{M}$ should be expressed by a distinct $s \in \mathcal{S}$, i.e. $p(m)$ should be a one-to-one function.
2. Each $s \in \mathcal{S}$ should map back to a single $m \in \mathcal{M}$ such that $p(m) = s$, i.e. $r(s)$ should be a superset of the inverse of $p(m)$.

[9]See Section 3.4.1.3. To recap, the value in α specifies how to change the connection weight between coactive units, β specifies how to change the connection weight between an active meaning node and an inactive signal node, γ specifies how to change the connection weight between an inactive meaning node and an active signal node, and δ specifies how to change the connection weight between two inactive units.

3.5.1 The key bias in the associative network model

There is a clear pattern relating the properties of weight-update rules to the assignment of actions to values in the $(\alpha \;\; \beta \;\; \gamma \;\; \delta)$ 4-tuple. Given the (approximately) bidirectional nature of the networks and assuming $|\mathcal{S}| \geq |\mathcal{M}|$, point 1 above ($p\,(m)$ should be a one-to-one function) proves to be crucial in determining which weight-update rules are [+constructor], which are [+maintainer, $-$constructor] and which are [+learner, $-$maintainer, $-$constructor]. Weight-update rules which are [+constructor] are biased in favour of a one-to-one $p(m)$, those which are [+maintainer, $-$constructor] are neutral with respect to the one-to-one nature of $p(m)$ and those which are [+learner, $-$maintainer, $-$constructor] are biased in favour of a many-to-one $p(m)$.

3.5.1.1 The [+constructor] bias

Is there any pattern of assignment of values to conditions in the weight-update rule specification $(\alpha\ \beta\ \gamma\ \delta)$ that characterises rules which are [+constructor] but not rules which are [$-$constructor]? Yes.

> A weight-update rule is [+constructor] if $\alpha > \beta \wedge \delta > \gamma$

Why does this pattern of weight changes result in the construction of optimal systems from random behaviour? Consider a network where $|\mathcal{N}_M| = |\mathcal{N}_S| = 2$ using the weight-update rule $(a\ b\ c\ d)$. Prior to learning, all the connection weights in $\mathcal{W}$ are 0. If we represent $\mathcal{W}$ as a matrix with the value in row i and column j representing the weight of the connection between nodes Mi and Sj then its initial weights will be:

$$\begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$

If this network is exposed once to the meaning m_1 (recall from Section 3.4.1.1 that for this meaning $a_{M1} = 1$, $a_{M2} = 0$), paired with the signal s_1 (similarly, $a_{S1} = 1$, $a_{S2} = 0$), its weight matrix will be:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

For rules which are [+constructor] $a > b$. This means that if our simple network uses a [+constructor] rule it will correctly produce s_1 to communicate m_1, due to the winner-take-all retrieval procedure.

For [+constructor] rules, $d > c$. In the context of our simple network, this means that if the network uses a constructor rule it will automatically prefer to use the signal s_2 to communicate meaning m_2, despite the fact it has only been trained to associate m_1 with s_1. This is the crucial property of [+constructor] rules — they are biased in favour of acquiring one-to-one mappings between meanings and signals. What consequences does this bias have in the context of iterated cultural transmission?

Only communication systems which conform completely to the biases of learners will be stable over iterated cultural transmission — communication systems which partially conform to learner biases will be less likely to be acquired than systems which conform more fully to the learner biases, and will therefore be filtered out of the population over time. This differential retention of communication systems resulting from learner biases results in direct bias on cultural transmission, as defined by B&R. The [+constructor] bias in favour of one-to-one mappings between meanings and signals results in many-to-one mappings being filtered out of the population. Eventually, through the process of iterated learning, the population converges on a shared one-to-one mapping between meanings and signals — an optimal communication system is constructed.

3.5.1.2 The [+maintainer] bias

Can the [+maintainer] property also be explained in terms of allocations of actions to the $(\alpha\ \beta\ \gamma\ \delta)$ weight-update rule specification? First, is there any pattern which uniquely identifies the [+maintainer, $-$constructor] rules? Yes.

> A weight-update rule is [+maintainer, $-$constructor] if $\alpha > \beta \wedge \delta = \gamma$

Once again consider a network where $|\mathcal{N}_M| = |\mathcal{N}_S| = 2$ using the rule $(a\ b\ c\ d)$ exposed once to m_1 paired with s_1. As before, the resultant weight matrix is:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

As for [+constructor] rules, for [+maintainer, $-$constructor] rules $a > b$. This means that if our simple network uses a [+maintainer, $-$constructor] rule it will correctly produce s_1 to communicate m_1.

For [+maintainer, $-$constructor] rules $d = c$. This means that, unlike [+constructor] rules, the network using a [+maintainer, $-$constructor] rule will be equally likely to express m_2 using s_1 or s_2, due to their equal weights in the network. [+maintainer,

$-$constructor] rules are therefore neutral with respect to one-to-one mappings. This explains both the ability of populations of agents using such rules to maintain optimal systems in the context of the ILM and the behaviour of these populations as they attempt to construct optimal systems.

[$+$maintainer, $-$constructor] rules can maintain an optimal system in the presence of noise. The initial optimal system is, by definition, a one-to-one mapping between meanings and signals. Given the neutrality of [$+$maintainer, $-$constructor] rules to the one-to-one nature of mappings, such optimal systems can be acquired in the presence of noise, provided the noise is not sufficient to drown out the one-to-one mapping.

Recall from Section 3.4.5 and Figure 3.11 that, when provided with an initially random system, populations of agents using [$+$maintainer, $-$constructor] rules converge on the level of communicative accuracy one would expect given a random assignment, with replacement, of signals to meanings. This can be explained in terms of the neutrality of [$+$maintainer, $-$constructor] rules to the one-to-one nature of mappings. The initial population's random behaviour, when taken as a whole, will embody a random assignment of signals to meanings. This random assignment will become shared among the population through the process of cultural transmission. While [$+$constructor] agents remove the many-to-one elements of the initial random system, [$+$maintainer, $-$constructor] agents do not — the population's eventual communication system will embody the same number of many-to-one mappings as the initial random behaviour.

What then of the [$+$maintainer] property in isolation from the [$\pm$constructor] feature? This can be captured thus:

A weight-update rule is [$+$maintainer] if $\alpha > \beta \wedge \delta \geq \gamma$

The fact that rules which are [$+$constructor] are always [$+$maintainer] is captured by this statement, as is the fact that it is possible to be [$+$maintainer, $-$constructor].

3.5.1.3 The [$+$learner] bias

The pattern of assignments of actions to the weight-update rule specification $(\alpha\ \beta\ \gamma\ \delta)$ that characterises rules which are [$+$learner] is:

A weight-update rule is [$+$learner] if $\alpha + \delta > \beta + \gamma$

or, in simple terms, in order to be able to acquire an optimal communication system you must make stronger associations between units which tend to have matching activations

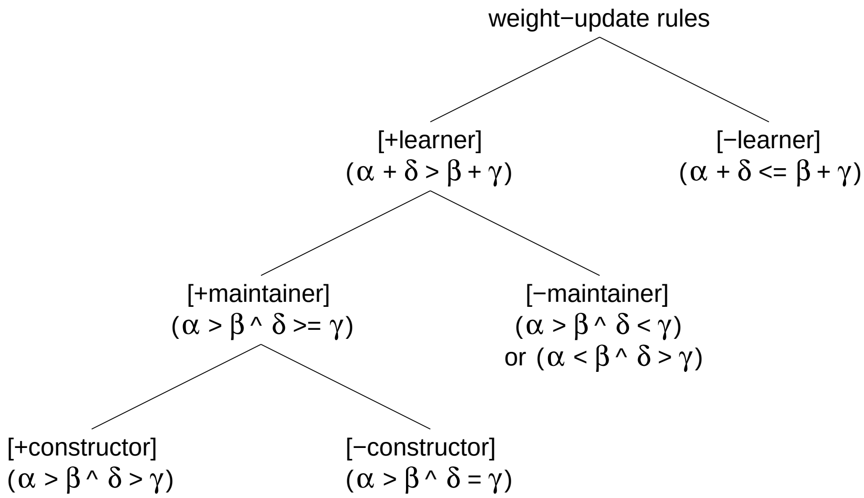

Figure 3.13: The hierarchy given in Figure 3.12, expressed in terms of restrictions on possible values in each condition of weight-update rules.

than between units which tend to have conflicting activations. Note that the $\alpha > \beta \wedge \delta \geq \gamma$ constraint on [+maintainer] rules guarantees that all such rules are also [+learner].

Why are rules which are [+learner, −maintainer, −constructor] unable to maintain or construct optimal communication systems? As we might expect, such weight-update rules are biased *against* one-to-one mappings between meanings and signals and in favour of many-to-one mappings. This immediately rules out construction of the one-to-one mappings characterising optimal systems, and also maintenance of such systems. Any many-to-one mappings introduced by noise will be preferentially acquired by [+learner, −maintainer, −constructor] agents and will spread through populations of such agents, resulting in the type of decrease in communicative accuracy seen in Figure 3.10.

3.5.1.4 Summary of the key bias in the associative network model

The weight-update rule hierarchy given in Figure 3.12 is re-presented in Figure 3.13 in terms of the constraints on the values of the weight-update rules. Each terminal node of the tree has a bias, summarised in Table 3.7.

3.5.2 The key bias in the feedforward network model

In the feedforward network model both imitator and obverter agents learn using the back-propagation algorithm. The bias is therefore introduced by the architecture of these networks, rather than the particular learning rule used. Imitator networks map from input

Classification	Bias
[$-$learner, $-$maintainer, $-$constructor]	NA
[$+$learner, $-$maintainer, $-$constructor]	favours many-to-one mappings
[$+$learner, $+$maintainer, $-$constructor]	neutral
[$+$learner, $+$maintainer, $+$constructor]	favours one-to-one mappings

Table 3.7: A summary of the learning biases of each particular combination of features. Weight-update rules which are classified as [$-$learner, $-$maintainer, $-$constructor] cannot be said to have a learning bias as they cannot learn.

meanings to output signals, whereas obverter networks map from input signals to output meanings. This turns out to be crucial in understanding the bias of these networks.

Feedforward neural networks learn many-to-one functions. Due to the deterministic nature of the feedforward propagation of activation values they cannot learn one-to-many mappings. The easiest function for a network to acquire is an all-to-one mapping from inputs to outputs, the hardest learnable function is an injective (one-to-one) function and one-to-many mappings are unlearnable. The reversal process used to model reception behaviour for imitators and production behaviour for obverters is similarly biased — it generates a function, which may be injective or many-to-one, based on the function the feedforward network has acquired. In general, if the network has acquired a function $f(x)$ which has a range y, then the reversal process ensures that element $y_i \in y$ will map onto a single element $x_i \in x$ such that $f(x_i) = y_i$ — in simple terms, the reversal process deterministically reverses the function acquired by the network.

In imitator agents the feedforward network learns functions from meanings to signals — it learns $p(m)$. Since it is a feedforward network it will be biased towards acquiring a many-to-one or all-to-one $p(m)$. As illustrated in Figure 3.14 and discussed in the caption, the maximally stable $p(m)$ for imitator agents is therefore an all-to-one fully ambiguous function. Imitators are therefore biased against one-to-one mappings from meanings to signals. Reception in imitators will be based on their acquired $p(m)$ — as shown in Figure 3.14, in the case of an all-to-one $p(m)$, in $r(s)$ the signal s_i that constitutes the range of $p(m)$ will map onto a single element from m. Therefore a population of imitators agents will tend to produce the same signal for every meaning and interpret the ambiguous signal as communicating one arbitrary selected meaning. This situation results in performance equivalent to random guessing.

In obverter agents the feedforward network learns functions from signals to meanings — it learns $r(s)$. As illustrated in Figures 3.15 and 3.16 the only culturally stable system has a one-to-one $p(m)$ and an $r(s)$ which includes at least the inverse of $p(m)$. Obverter

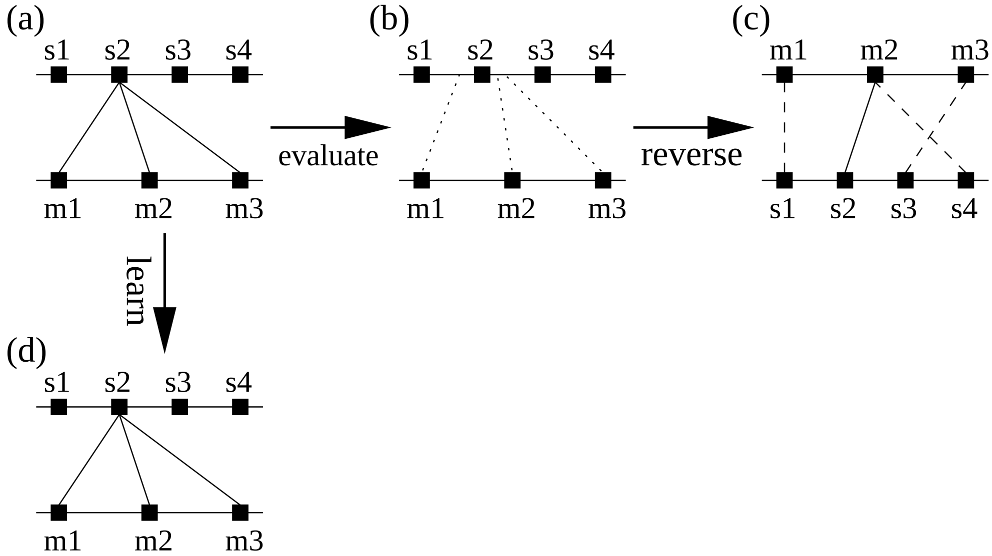

Figure 3.14: (a) is a representation of an imitator agent's feedforward network encoding an all-to-one $p(m)$ mapping three meanings onto a single signal, $s2$. The function from a domain of real numbers (input unit activations) to a codomain of real numbers (output unit activations) is represented by two lines, the lower line representing the domain, the upper representing the codomain. Squares represent particular points on the line corresponding to binary meanings or signals. Associations are shown with solid lines between elements in the domain and elements in the codomain. (b) represents the confidence-measuring step of the reversal process for the network underlying (a). In order to decide $r(s2)$, the real-number values of $p(m1)$, $p(m2)$ and $p(m3)$ are calculated. These real-numbered mappings are represented by dotted lines in (b). (c) represents the $r(s)$ derived from applying the reversal process to (a). $r(s2) = m2$ because $m2$ mapped closer to $s2$ than any other m in (b). The other associations are effectively random. The random nature of these mappings is represented by dashed lines. (d) represents the function acquired by an imitator network exposed to behaviour generated by (a) — as it is an all-to-one function between meanings and signals it is easily learned by imitator agents. This is in fact the only stable function for imitators.

agents are therefore strongly biased in favour of acquiring systems with the properties of optimal communication systems.

How can we relate these feedforward network biases to the classification hierarchy developed for the associative network weight-update rules? Obverter networks, biased in favour of one-to-one mappings between meanings and signals, should clearly be classified as [+constructor]. The classification of imitator agents is less clear. Imitator agents are capable, given sufficiently high e, of acquiring an optimal, unambiguous communication systems, and should therefore be classified as at least [+learner]. Populations of such agents cannot construct an optimal system, and should therefore be classified [$-$constructor]. Their status with respect to the [$\pm$maintainer] feature is less clear. We would expect, given their bias in favour of many-to-one functions, that they should be classified as [$-$maintainer]. However, simulation runs were carried out to measure the

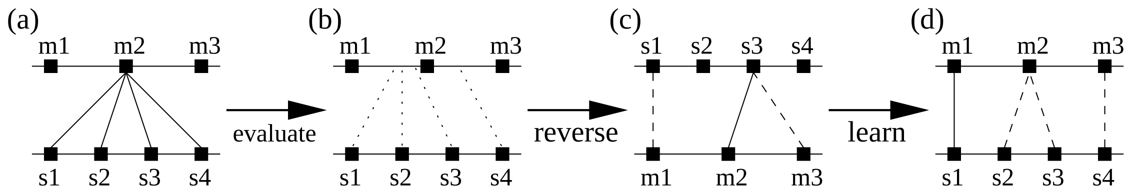

Figure 3.15: (a) represents an all-to-one $r(s)$ encoded in an obverter agent's feedforward network. As obverters map from signals to meanings this is the most learnable $r(s)$. (b) represents the confidence-measuring step of reversing this $r(s)$ to generate a $p(m)$ — as before, real-number mappings are shown as dotted lines. (c) shows the $p(m)$ derived from (a). $p(m2) = s3$ as $s3$ mapped closest to $m2$ in (b). The other associations are essentially random. The $p(m)$ in (c) produces the meaning-signal pairs $\{(m1, s1), (m2, s3), (m3, s3), \}$. Meanings and signals in these pairs are transposed (yielding $\{(s1, m1), (s3, m2), (s3, m3), \}$) to train the next generation of obverter networks. (d) shows the $r(s)$ resulting from training an obverter network on the signal-meaning pairs $\{(s1, m1), (s3, m2), (s3, m3), \}$. $r(s1) = m1$, as expected. However, feedforward networks cannot learn one-to-many mappings so $r(s3)$ is effectively randomly assigned to a signal, in this case $m2$. As $s2$ and $s4$ are not represented in the training set they are effectively randomly assigned mappings. Notice that the mapping in (a) has been destroyed in (d) — the many-to-one mapping in (a) is not culturally stable.

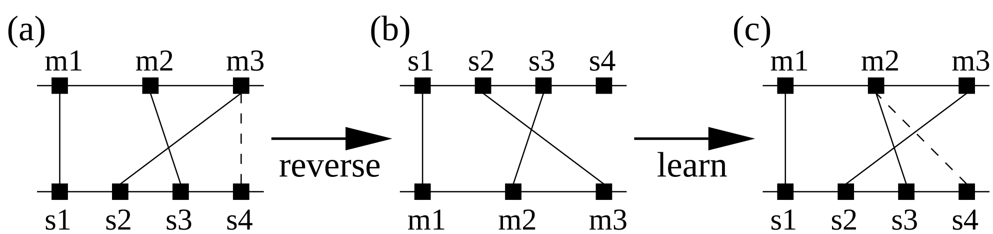

Figure 3.16: Only an unambiguous $p(m)$ is stable for obverter agents. (a) represents an obverter agent's $r(s)$. (b) is the $p(m)$ derived from reversal of (a) — it is a one-to-one function. (c) illustrates that the $r(s)$ resulting from training the next generation of agents on data produced by (b) is effectively similar to (a) and will therefore lead to (b) once again — (a) and (b) are culturally stable. The only unstable aspect is the floating synonym $s4$. This synonym is highly unlikely to interfere with the mapping in (b) and the floating synonym phenomenon can be observed in other obverter models.

ability of populations of such agents to maintain an optimal system in the presence of noise ($e = 200$, $p_n = 0.05$) and no runs were found for which they failed to do so. However, this appears to be due to the large value of e, which reduces the impact of noise. As we will see in Chapter 4, injection of a slightly different form of noise does result in failure to maintain an optimal system. We will therefore classify imitator networks as [$+$learner,$-$maintainer,$-$constructor].

3.5.3 The key bias in other models

Can we understand the behaviour of other models of the cultural evolution of vocabulary in terms of this key bias? Specifically, in the models where cultural evolution is driven by direct bias, does the direct bias result from a learning bias similar to that of [$+$constructor] agents? Dealing with other neural network models first, does this key bias appear in the neural network models of Hutchins & Hazelhurst (1995), Batali (1998) and Livingstone & Fyfe (1999) (discussed in Section 3.1) and Kvasnička & Pospíchal (1999) (which will be discussed in Chapter 4), and Hare & Elman (1995) and Kirby & Hurford (2002) (discussed in Chapter 5)?

Hutchins & Hazelhurst's (1995) model can be treated separately from the other models, which all share a common model of a learner. Hutchins and Hazelhurst use autoassociator networks to model communicative agents, with patterns of activation over the hidden layer being interpreted as signals. Autoassociator networks must develop a distinct pattern of activation over the hidden layer for every input-output pair (input-output pairs are equivalent to meanings as defined here) in order to succeed in the autoassociator task. Interpreting the hidden-layer patterns of activation as signals therefore builds in a one-to-one bias of the type identified as crucial for developing an optimal communication system.

Batali (1998), Kvasnička & Pospíchal (1999), Livingstone & Fyfe (1999) and Kirby & Hurford (2002) all use the obverter feedforward network configuration, with networks mapping from input signals to output meanings. As discussed in Section 3.5.2, such a configuration results in a learning bias in favour of one-to-one meaning-signal mappings. The obverter network configuration, which is quite common in the literature, therefore builds in a strong bias in favour of optimal communication systems.

Hare & Elman's (1995) model of morphological change deserves a brief mention here. This network maps from semantic representations of verbs to representations of the phonological realisation of those verbs, and Hare & Elman observe a simplification of the phonological system, with increasing numbers of verbs being expressed with similar affixes. While this pattern exhibits some complex interactions between phonological regularity and frequency of tokens, the general pattern of convergence to many-to-one mappings is what we should expect to see from an imitator network architecture. We can speculate that, had Hare & Elman allowed their simulations to continue for several hundred generations, all verbs would end up being expressed with a single phonological form. The learning bias of the imitator network would essentially destroy the morphological system. Contrast this with Batali's (1998) results, where the obverter network architecture results in the emergence of a morphological system.

Non-neural network models of the evolution of vocabulary, where that evolution is driven by direct bias, are actually rather scarce, the only clear examples being the models of Hurford (1989) and Oliphant & Batali (1997). As discussed in Section 3.1, Hurford considers three learning strategies — Calculators, Imitators and Saussureans. Populations of individuals using the first of these strategies cannot maintain optimal systems over time, even when there is no noise on cultural transmission, while Imitator and Saussurean populations can construct communication systems which yield intermediate levels of communicative accuracy, with Saussureans being somewhat more successful than Imitators in this respect.

Based on these results, we would expect that Calculators can be classified as [$\pm$learner, $-$maintainer, $-$constructor] and Imitators and Saussureans can be classified as [$+$learner, $+$maintainer, $-$constructor]. Saussureans have the additional advantage over Imitators that their production and reception behaviour are necessarily closely coupled — while Imitator learners acquire their production and reception matrixes completely independently from one another, Saussureans construct their reception matrix on the basis of their own production matrix.

Do these learning strategies have the biases we would expect with respect to the one-to-one nature of the meaning-signal mapping? In other words, are Imitators and Saussureans neutral with respect to the one-to-one nature of mappings, and how are Calculators biased with respect to this property?

It should be fairly obvious that Imitators and Saussureans *are* neutral with respect to many-to-one mappings from meanings to signals — they acquire their production matrix straightforwardly, by memorising the meaning-signal pairs they observe being produced. The story with Calculators is somewhat more complicated. Consider a Calculator trying to learn a system involving 2 meanings and 3 signals. Assume that the learner observes a system where $m1$ is communicated using $s1$ and $m2$ is communicated using $s2$, with

$s3$ being unused. This is a one-to-one system. On the basis of this observed production behaviour the Calculator will arrive at the reception matrix

R	m1	m2
s1	1	0
s2	0	1
s3	0.5	0.5

$s3$ is interpreted as meaning $m1$ and $m2$ with equal probability, given that these signals are unused in the observed production system. This individual's reception behaviour will now be used by the next generation of learners to form their production behaviour. What happens?

The Calculator above will produce two possible sets of reception behaviour. In both, $s1$ is interpreted as $m1$, $s2$ interpreted as $m2$, with $s3$ being interpreted as $m1$ or $m2$ at random. What consequence does this have for the production matrices of the next generation? If the learner at the next generation observes $s3$ being interpreted as $m1$ they will arrive at the production matrix

P	s1	s2	s3
m1	0.5	0	0.5
m2	0	1	0

In other words, they will now produce either $s1$ or $s3$ to communicate $m1$. If the learner instead observes $s3$ being interpreted as $m2$ they will arrive at the production matrix where $s2$ and $s3$ are produced with equal probability for $m2$. In other words, the spare signal leads to the creation of one-to-many mappings between meanings and signals. The inability of Calculators to acquire an optimal system (albeit over the course of two learning episodes) therefore indicates they should be classified as [$-$learner]. This classification explains the fact that populations of such individuals immediately lose an initially perfect system in an ILM where there is no noise on cultural transmission.

As discussed in Section 3.1, Oliphant & Batali contrast two learning strategies (Imitate-Choose, henceforth imitator, and obverter). Populations of obverter learners can construct an optimal system, whereas populations of imitators cannot. Oliphant & Batali attribute this to the fact that obverter agents base their production behaviour on the population's reception behaviour, therefore explicitly designing their communication systems so as to be understood, whereas imitator agents do not. However, the [+constructor] agents described in Section 3.4 base their production behaviour on production behaviour, yet still arrive at an optimal communication system. Do the obverter agents described by

Oliphant & Batali arrive at an optimal system in a different way, or can the behaviour of populations of such individuals be better explained in terms of a learning bias in favour of one-to-one mappings between meanings and signals?

Figure 3.17 analyses how a population's production (P) and reception (R) functions change over time, for two initial starting conditions — an initial one-to-one mapping, and an initial many-to-one mapping. Both these initial mappings are learnable by imitator agents. For obverter agents, the one-to-one mapping remains stable over time — the one-to-one P matrix leads to a one-to-one R matrix, which in turn leads back to the one-to-one P matrix. In contrast, the many-to-one P matrix is unstable. The obverter procedure attempts to derive an R matrix from this P matrix by finding the meaning m_x for which s_1 and s_2 is at a maximum. Since m_1 and m_2 are equally likely in both these contexts, a random choice is made, which leads to four possible R matrixes, all equally probable. For the sake of convenience, only two are shown in Figure 3.17. One of these is a one-to-one R matrix, which leads to a one-to-one P matrix, which, as we have seen, is stable. The other is a many-to-one R matrix, which leads, again through random selection, to 4 possible P matrixes, two of which are shown. Only the one-to-one P matrix is stable — the other, as we have seen, is unstable. The learning bias of the obverter agents favours one-to-one mappings between meanings and signals.

3.6 Biases in vocabulary acquisition in humans and non-humans

Is there any evidence that language acquisition in humans is guided by biases in favour of one-to-one mappings between meanings and signals? If so, then the result of the computational models shown here would suggest that optimal, or at least communicatively useful, communication systems could arise in human populations through purely cultural processes.

The one-to-one bias described here is typically broken down into two subcomponents when talking about vocabulary acquisition in humans. The one-to-one bias consists of both a bias against homonymy (many-to-one mappings from meanings to signals) and a bias against synonymy (one-to-many mappings from meanings to signals). In Sections 3.6.1 and 3.6.2 I will present evidence from the language acquisition literature that suggests that children apply both these biases to the learning of vocabulary. I will start with the proposed bias against synonymy, as this is perhaps slightly less controversial, then move on to homonymy. Finally, in Section 3.6.3 I will briefly review some evidence from ape 'language' learning experiments which suggest that apes may not possess similar learning biases to human infants. The postulated uniqueness of these learning biases

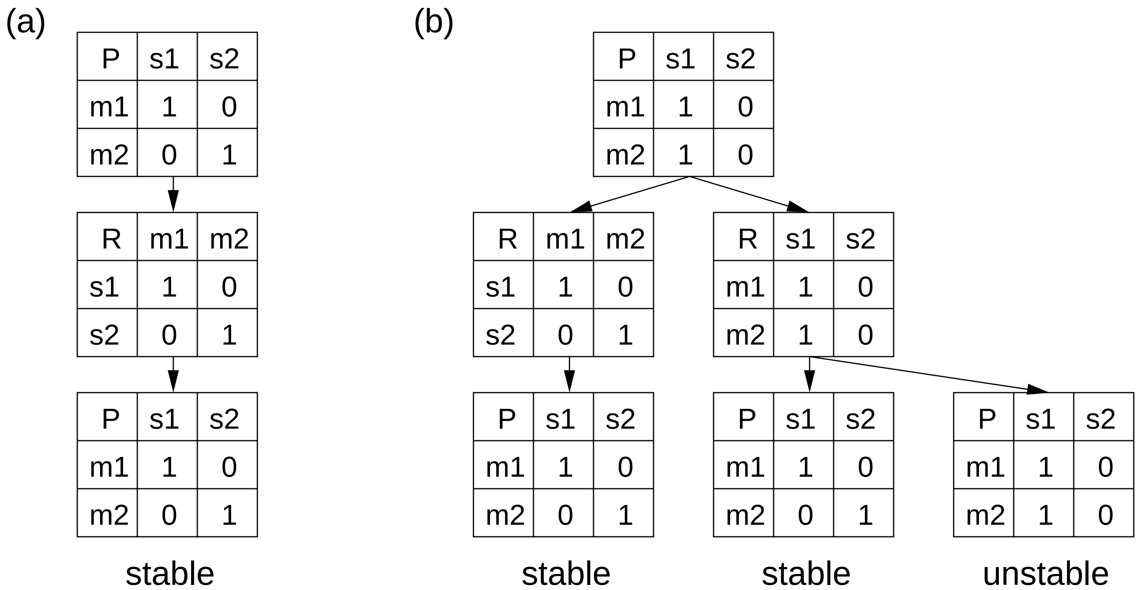

Figure 3.17: The learning bias of Oliphant & Batali's obverter learner. As illustrated in (a), one-to-one mappings are stable over time — a one-to-one production matrix (marked by a P) leads to a one-to-one reception matrix (R), which leads back to a one-to-one production matrix. Many-to-one mappings are unstable, as illustrated in (b). The many-to-one production matrix leads to four possible reception matrixes (only two are shown here). The one-to-one reception matrix leads to a one-to-one production matrix which, as shown in (a), is stable. The many-to-one reception matrix leads either to a one-to-one production matrix, which is stable, or a many-to-one production matrix, which as we have seen is unstable. In other words, only one-to-one matrixes are stable over time in populations of such learners.

to humans forms the motivation for Chapter 4, where investigations into the biological evolution of the one-to-one learning bias are discussed. I will return to the role of one-to-one biases in the acquisition of linguistic structure (both syntactic and morphological, as opposed to essentially unstructured vocabulary which I discuss here) in Chapter 5.

3.6.1 Biases against synonymy in humans

Eve Clark and Ellen Markman have proposed that children have word-learning biases which make synonyms (mappings from one meaning to several distinct words) either unlearnable or difficult to learn. Both authors claim that these biases help children in the rapid acquisition of vocabulary — the acquisition process in children has been termed "Fast Mapping" (see Bloom (2000), Chapter 2, for review). While their conception of this bias is in fact rather different, both base their theories on a set of experimental studies carried out by Kagan (1981), replicated by, among others, Markman & Wachtel (1988).

In Kagan's original study, children were shown three objects, two of which were familiar (a doll and a dog) and one of which was unfamiliar (for example, a lemon zester). Children were allowed to play with the objects, and were then asked to "Give me the zoob" (or some other nonsense word) by the experimenter. The children showed a strong preference for giving the novel item.

In Markman & Wachtel's (1988) replication, children were shown a single familiar object (for example, a plate) and an unfamiliar object (e.g. a radish rosette maker) and asked by a puppet frog to "Show me the fendle" or some other nonsense word. Children reliably respond by giving or showing the unfamiliar object. Results from a control group study, where children were asked simply to "Show me one", indicated that this preference was not due to a preference on the part of children to respond with the unfamiliar object — children only exhibit such a preference when prompted with a novel word.

Clark (Clark 1988; Clark 1990; Clark 1993) proposes two pragmatic principles guiding vocabulary acquisition. The first, which she terms the Principle of Contrast (henceforth Contrast), states that "different words have different meanings". The second, the Principle of Conventionality (henceforth Conventionality), states that "for certain meanings, there is a form that speakers expect to be used in the language community".

How does this relate to the child's behaviour in the experiment outlined above? The child knows, through Contrast, that contrasting words contrast in meaning. The child knows, through Conventionality, that established words have priority. It is assumed that the child already knows the established word for the familiar object. The child then deduces, via Conventionality, that if the experimenter wished to refer to the familiar object they would use the conventional word for that object. However, notes the child, the experimenter used a novel word. The child reasons that that word cannot refer to the familiar object, because different words have different meanings. The child therefore concludes that the new word must refer to the unfamiliar object, responds appropriately by giving the experimenter the unfamiliar object, and learns the (nonsense) name for the unfamiliar object.

Contrast rigidly rules out synonyms — Clark emphasises that any difference in the form of a word indicates a difference in meaning. To put it another way, according to Clark synonyms do not exist. Clark is keen to point out that it does not rule out homonymy (many-to-one mappings from meanings to signals), a point which I will return to below.

Markman (Markman & Wachtel 1988; Markman 1989; Markman 1992) proposes a Mutual Exclusivity (ME) bias in children — "children should be biased to assume, especially at first, that terms [words] are mutually exclusive" (Markman 1989:188) and "each object will have only one label". Note that, unlike Contrast, this is not an inviolable principle, but a tendency or bias that can be overridden given sufficient evidence. Like Contrast, ME discourages synonymy — each object ideally has only one label, therefore there should be no one-to-many mappings between meanings and signals. Markman is less clear on the status of ME with respect to homonymy (for example, the term does not appear in the index to her 1989 book), although she makes frequent reference to "one-to-one" biases in vocabulary acquisition. Like Contrast, ME enables the child in the experiment above to deduce that the novel word should refer to the unfamiliar object and learn that labelling — ME dictates that the novel word cannot refer to the object which is already labelled, therefore the novel word must refer to the unfamiliar object.

The main difference between Clark's position and Markman's is in the severity of the bias. Markman (1989) presents what she considers to be evidence that children *can* learn to violate ME. Specifically, they can learn super-ordinate terms ("poodle" but also "dog" and "animal"). However, children find such super-ordinate terms difficult to learn, with children as old as 14 making errors. Typically, the error involves mistaking a super-ordinate term for a term expressing a collection of the subordinate items. Macnamara (1972) gives the example of a child who will accept that a particular plaything is a "truck" or a "train", but will simultaneously deny that it is a "toy", with the term "toy" being reserved for a group of trucks, trains, Teddy bears and so on. This type of evidence leads Markman to conclude that ME is violable. In contrast, Clark would say that super-ordinate terms have a different meaning from their subordinate terms, therefore Contrast is preserved, although this offers no explanation as to why children find super-ordinates hard to learn. However, regardless of the debate between Contrast and ME, the experimental evidence suggests that children have a bias against acquiring synonyms, one-to-many mappings from meanings (objects in the experiment) to signals.

3.6.2 Biases against homonymy in humans

The status of biases with respect to homonymy is somewhat problematic. One does not have to look very far through any language to find homonyms, mappings from several meanings to a single surface form. At first blush, this seems to indicate that children do not have any bias against acquiring many-to-one mappings from meanings to signals. Indeed, Clark is keen to point out that Contrast and Conventionality do not in any way bias against homonymy. Markman remains silent on the subject of homonymy, although in several places she refers to "one-to-one" biases in vocabulary acquisition.

Briefly considering a possible experiment should, however, serve to cast doubt on the first intuition that children are unbiased with respect to the acquisition of homonyms. Imagine a slight variation in the experiment outlined above, where, rather than the experimenter asking the child to "show me the zoob", the experimenter asks the child "Show me the shoe" or whatever the familiar object was. I suspect that the child would respond by showing the shoe, at least with the same level of reliability as children prompted with the novel word would respond by showing the unfamiliar object. The fact that nobody, to my knowledge, has carried out this experiment indicates that this is probably not a very controversial hypothesis — our everyday experience indicates that people know the names for things and if you request those things then they don't assume you are talking about something else.

However, this experiment, were it to proceed as I expect, would illustrate that children must be biased against homonymy. If children are unbiased with respect to homonymy then under such experimental circumstances they should show the shoe or the unfamiliar object with equal probability — if many-to-one mappings between meanings and signals are as possible as one-to-one mappings, then the child cannot tell whether "shoe" means shoe_1, the familiar sense of shoe, or shoe_2, a new use of homophonous "shoe" to refer to the unfamiliar object. It could be argued that the child would prefer the familiar sense of "shoe" due to the fact that they have frequently encountered this use of "shoe" (this is essentially Clark's Conventionality principle). However, if we agree that children prefer not to learn new meanings for established words then we are accepting a bias against homonymy. It could be argued that this bias against homonymy only comes into play once a well-established convention is in place — for example, once the child has experienced several hundred utterances of "shoe", all with the same reference, they will be resistant to learning a new meaning for "shoe". However, it cannot simply be sheer weight of numbers which performs this function. Bloom & Markson (1998) describe an experiment where children are presented with two novel objects, given a single nonsense name for one object ("bem", for example), then asked to "Show me the jop". Under these circumstances children still reliably show the unnamed novel object. Given that neither object has been encountered before, and the word "bem" is also novel, it seems that a single exposure to a word paired with an object biases children against interpreting or acquiring that word as conveying a different meaning.

It in fact seems that, with a learner unbiased with respect to homonymy, word learning would become all but impossible — every possible utterance of a familiar word could refer to any object at all. If we accept Contrast or ME, then the problem is fractionally reduced — any word can refer to any object which has not already been labelled with a different word. However, Contrast/ME cannot really resolve the issue as learning even a single label becomes an intractable task. This problem becomes worse when we consider that children can learn labels for subparts of objects. Joint attention might narrow down the focus of possible objects, ruling out other whole objects as the referent of “shoe”, but when I say “Show me the shoe”, do I mean shoe$_1$, the whole object, or shoe$_3$, the string that ties the shoes up, or shoe$_4$, the man-made fibre which the shoe uppers are made of?

I believe that we are forced to conclude, on logical grounds, that children must have some bias against acquiring homonyms, many-to-one mappings from meanings to words. Without such a bias, word learning would become impossible or at best extremely arduous. This is in fact not a terribly new position, although the argument given above may be a novel one. McMahon (1994) briefly discusses work carried out by Jules Gilliéron in France from 1896. Gilliéron, with the aid of several fieldworkers, compiled the Linguistic Atlas of France. One of Gilliéron’s primary concerns was to construct “phonetic etymologies”, by comparing the expected forms of modern French words (based on a set of hypothesised changes occurring between Latin and modern French) to the words actually attested. Where the predictions and the data did not match up Gilliéron attempted to explain the discrepancy.

Gilliéron’s theory predicted that the modern French word for cockerel should be derived from the Latin “gallus”. In the Gascony region these hypothesised changes should have lead to the form “gat”. However, this is also the predicted form for the word for cat, “cattus” in Latin. Gilliéron’s survey revealed that “gat” in Gascony does in fact refer to cat, with cockerel being referred to by another word. Gilliéron appealed to an avoidance of homonymy to explain this mismatch — cats and cockerels are both farmyard animals (or presumably were in turn-of-the-century France), and homophony involving meanings from the same semantic field is avoided, therefore “gat” was restricted to meaning cat and an alternative form was employed for cockerel.

This early example does not indicate where the bias against homonymy resides — is it in the language learner or the language user? Martinet (1972) proposes the second of these alternatives. Martinet’s primary concern is an account of phonemic change, but he works within a functionalist framework: “The basic assumption of functionalists in such matters is that sound shifts do not proceed irrespective of communicative needs, and that one of the factors which may determine their direction and even their appearance is the basic necessity of securing mutual understanding” (Martinet 1972:144). The imperative to preserve mutual understanding should discourage, among other things, homonymy — homonymy leads to ambiguity.

Lass (1980) provides a concrete example of an irregular phonemic change which appears to result in the avoidance of homonymy. During the change from Old English to modern English, the Old English vowel /y/ appears to change in two distinct ways. In the first, regular path, Old English /y/ changes to modern English /ɪ/, via Middle English /i/. However, in some lexical items Old English /y/ appears to change to Middle English /u/ and then to modern English /ʌ/. The vowel of modern English "shut" proceeded according to the less frequent, irregular path. Had it proceeded according to the more regular path, the result would have been "shit", as Lass puts it "[t]his particular homophony would be, I would think, about as 'pernicious' as any" (Lass 1980:76). Those subscribing to Martinet's line of reasoning would perhaps argue that "shut" avoided the more regular change in order to avoid homonymy, which would lead to a decrease in communicative function, although Lass argues strongly against this interpretation, as we will see below.

Turning briefly to cross-categorial homonymy, Macnamara (1982) reports two pieces of evidence, based on the acquisition behaviour of two young subjects, that children prefer not to acquire homophonous terms which refer both to an action and an object. Macnamara's first subject preferred to use ambiguous terms such as "comb" to refer to either the action or object, but not both, even when the child's parents used the term as both noun and verb in the child's presence. Macnamara's second subject, his son, went so far as to invent a new word to avoid this type of homonymy.

To summarise, the logical argument and proposed experiment outlined at the beginning of this Section suggest that children must be biased to some extent against acquiring many-to-one mappings from meanings to signals. This is supported by some concrete examples of where this type of homonymy avoidance might be observed empirically. The empirical evidence at this point is somewhat weak, and will be considerably strengthened in Chapter 5, where more evidence for a bias against homonymy is presented. However, for now it is time to return to the two main arguments against a bias with respect to homonymy — the fact that languages contain numerous homonymous mappings, and that teleological mechanisms do not exist for homonymy avoidance.

Firstly, if children do need to be biased against homonymy, why is homonymy so frequent in language? Doesn't this prove that children are in fact not biased against homonymy? There are two possible responses to this position, both similar to Clark and Markman's respective defences of their proposed biases relating to synonymy.

We could imitate Clark's rather rigid line of argument, and insist that, just as difference in signal reflects difference in meaning, congruence of signal represents similarity of meaning. This line, in a rather strong form, is pursued by Haiman (1980). While this argument can probably be used to deflect some cases of homonymy, such as polysemous uses of “mouth” (mouth of an animal, mouth of a cave, mouth of a river etc), it perhaps does not do to push it too far.

Alternatively, we could appeal to the kind of explanation that Markman makes to explain violations of ME — perhaps homonyms are simply somewhat harder to learn than non-homonyms, but still learnable. This testable hypothesis therefore allows that we should indeed expect to see homonymy in language, particularly when we consider that a learner bias against homonymy is not the only pressure acting on language and language acquisition — not only do phonological shifts and borrowing continually bring homonymy into a language, but there are other possible pressures:

> “[These other pressures might] pertain to the number of fixed expressions, patterns, and locutions that a speaker must master, remember, and manipulate in language use. The impracticality of having a separate lexical item for every conceivable object, event, or situation a speaker is likely to encounter is of course a truism. Languages never provide a lexical inventory that is vast enough to label with uniqueness and precision the elements of every conceivable contingency; rather they depend on the speaker to use creatively a more restricted inventory of lexical units in conjunction with the resources of the grammatical system.” (Langacker 1977:114)

It is interesting that this tendency to minimise the number of lexical items which have to be learned impacts differently on synonyms and homonyms. Synonyms will be disfavoured — memorising two words for a single object increases (perhaps unnecessarily) the learning burden. The preference for smaller vocabularies then reinforces the child’s bias against synonyms. However, a pressure to minimise vocabulary size *favours* homonyms — if a single word can be used for two meanings, then the total number of words which must be learned is reduced. The tendency to minimise vocabulary size therefore fights against the postulated learner bias against homonyms. This is one possible explanation for the apparently contradictory fact that there are few (if any) true synonyms in language and numerous homonyms, while children are biased against both synonyms and homonyms.

The second objection to proposed anti-homonymy biases is the suggested location of such biases. Lass (1980) deals rather forcefully with functionalist explanations. Perhaps his most telling criticism is that, according to the functionalists, “it seems that speakers avoid homophones by prolepsis, i.e. by taking avoiding action in advance . . . [b]ut this is surely absurd . . . the only mechanism left is for speakers actually to produce the offending articles, and then, having discovered what they've done, to remove them ('My God, I've just said "please *shit* the door"; better change it to *shut*')" (Lass 1980:79). If we accept this line of argument, this leaves us with the homonymy avoidance residing in the language learner.

Croft (2000) deals with teleological explanations for homonymy avoidance:

> "perhaps the greatest objection is that there is no plausible theory motivating a teleological mechanism [whereby language users change the linguistic system for the sake of the linguistic system] . . . innovations must be brought about ultimately as a result of actions by speakers. Yet there is no obvious motivation for speakers to innovate to make the grammatical system more symmetrical, or to preserve distinctions for the sake of preserving distinctions. Speakers have many goals when they use language, but changing the linguistic system is not one of them" (Croft 2000:70)

I could not agree more. Locating a bias against homonymy in the language learner avoids the distasteful aspects of teleological and functional explanations — children learning language avoid homonyms not because they're worried that they'll say "Please shit the door mother", but simply because they can't help it — the bias against homonymy is a component of the innate device which determines the way children acquire language. Of course functionalist explanations of the origin of this language learning bias still have to be explained, which will be the role of Chapter 4.

3.6.3 Biases in non-human animals?

Clark and Markman present evidence that children are biased against acquiring one-to-many mappings between meanings and signals. I have also presented an argument that children must be biased against acquiring many-to-one mappings between meanings and signals. We should therefore, if we accept both these factors, expect children to be biased in favour of one-to-one mappings between meanings and signals. This is precisely the bias that we found to be necessary for the cultural evolution of functional communication systems. This provides partial support for our theory that only humans have the mental capacity to support learned symbolic vocabulary — we have established that humans have the necessary bias. The second supporting strut for this argument would be to fill in the “only” piece — to show that no non-human animals have a learning bias which favours one-to-one mappings between meanings and signals.

Do non-human animals have a bias in favour of one-to-one mappings between meanings and signals? Some evidence on this front comes from the ape language-learning experiments, but this evidence is somewhat sketchy. This is largely due to the fact that Kanzi, possibly the most prominent and successful ape language learning subject, essentially learned what he learned without the researchers noticing — while the focus had been on teaching Kanzi’s mother to communicate “Kanzi had been keeping a secret. He had been learning these words all along . . . We thought he did not know how to talk with the keyboard, but he did.” (Savage-Rumbaugh *et al.* 1998:22). While this set of circumstances may shed some light on whether or not apes need explicit reinforcement to learn a communication system, it is rather disappointing from our current perspective — nobody noticed how Kanzi went about learning lexical items.

There is some empirical evidence on this point, however. David Premack reports (in the discussion section following Premack (1983)) on an experiment where chimpanzees requested a previously unnamed object using a “new word”, a newly introduced piece of plastic. This suggests a bias against homonymy — if the apes were unbiased with respect to homonymy, they could happily refer to the unnamed object using a known word, therefore introducing a many-to-one mapping between meanings and signals.

Other experimental work casts doubt on this conclusion, however. Lyn & Savage-Rumbaugh (2000) describe a fairly rigorous set of experiments into the ability of two pygmy chimpanzees (Kanzi and his younger half-sister Panbanisha) to learn new words for novel items. Their overall experimental setup is inspired by the experimental setup used to investigate the Contrast/ME principle in human infants.

Lyn & Savage-Rumbaugh tested the ability of the two apes to acquire new words (lexigrams) for ten novel items. Before an ape was tested on an item, it received pre-exposure to that item being named by two human experimenters in a naturalistic dialogue. These dialogues took place outside the ape’s living enclosure, but in clear view of the ape. During each presentation session, the two experimenters played with and discussed the item, naming it between six and 19 times. Presentation sessions were grouped in threes, with the experimenters discussing one novel item, going away, returning with another novel item, and so on.

Within one hour of these groups of three presentation sessions the apes were tested on their ability to name the novel items. During each test session the ape was presented with five familiar items and the three novel items it had just seen named. The apes were allowed to play with all the items until their initial curiosity waned. 11 tests were then conducted, during which the ape was asked to give a particular item to the experimenter.

During the first three such tests the apes were instructed to hand over a familiar item. Both apes responded correctly in at least two thirds of these initial tests. For the remaining eight tests the apes were told to give the experimenter each of the eight items, in random order. If the ape made a mistake during any one of these eight tests, the request was repeated a single time. During testing, food and praise were given freely to the apes, but "[i]ndication of the incorrectness of a response was kept to a minimum" (Lyn & Savage-Rumbaugh 2000:261), notwithstanding the repetition of the request.

If an ape failed to correctly name all three novel items the presentation and test sessions were repeated within 48 hours. The repeat presentation sessions proceeded exactly as before, with all three novel items being named. However, during the test sessions, the apes were only prompted to give the novel items which they had earlier failed on.

Lyn & Savage-Rumbaugh report that Kanzi required a mean of 2 presentation/test sessions before he correctly responded to the name of a novel item, which amounted to a mean 22.8 exposures to that novel item being named by experimenters. Panbanisha fared rather less well, requiring a mean of 4.1 presentation/test sequences and 42.5 exposures per item. Furthermore, on 70% of the test incidents when the apes were asked to give a novel item they failed to do so, either choosing a familiar item, choosing more than one item or refusing to answer. As Lyn & Savage-Rumbaugh point out, this is a below-chance level of performance, indicating a preference by the apes to select familiar items when asked for a novel item. Finally, while Kanzi's level of performance remained fairly constant over the series of experiments, Panbanisha's performance gradually improved — it took her a greater number of cycles to learn the name for the first novel item than it did for the last novel item.

How should we interpret these results with respect to homonymy and synonymy biases in apes? Firstly, it should be noted that apes clearly have a rather more difficult time with the task than children, requiring pre-exposure and multiple tests to select 'novel' items when presented with 'novel' words. Secondly, the apes' performance at some aspects of the task indicate a lack of a bias against homonymy, or at least a very weak bias. The apes sometimes failed when asked to give a familiar item during the first three trials of each test session, although not more than 33% of the time. This suggests that they are not strongly biased against associating a familiar word with a novel object or the wrong familiar object (a many-to-one mapping from objects to words). Their performance on requests for familiar items during the remaining eight tests is not reported. Thirdly, the apes' preference for selecting familiar, already-named items when prompted with a novel word indicates the absence of a bias against synonymy — the apes fail, or at least take time to pass, the Contrast/ME test. Finally, Panbanisha apparently has to learn how to perform the task, while it comes naturally to human infants. Whereas the experimenters interpret this as Panbanisha coming to terms with the test environment (Heidi Lyn, personal communication), it could be taken to indicate that she takes time to come to terms with the idea of naming novel items with novel words. What is clear from this set of experiments is that apes find the acquisition of words much more difficult than humans, and the process, which is known as Fast Mapping in human infants, takes time. The experiments also suggest that apes are either unbiased with respect to homonymy and synonymy, or at least much less biased than human infants, although this conclusion is more a matter of interpretation.

To further muddy the waters, it could be noted that, in a detailed report on a separate set of ape language experiments, Savage-Rumbaugh *et al.* (1986) report that Mulika (Kanzi's younger sister) "began by using the lexigram *milk* for many different things, including requests to be picked up, requests for attention, requests to travel to different places, requests for food and requests for milk" (Savage-Rumbaugh *et al.* 1986:219). Matata, Kanzi and Mulika's mother, "did not develop an adequate concept of one-to-one correspondence between a given symbol and a given referent" (Savage-Rumbaugh *et al.* 1986:215). These (admittedly circumstantial) examples of a lack of any obvious one-to-one bias in ape vocabulary acquisition should throw further doubt on whether the biases humans bring to this task are present in a closely related species.

3.7 Summary of the Chapter

In this Chapter I have presented two models of the cultural evolution of unstructured communication systems — one revolving around a feedforward network model of a learner, the other based on an associative network model. In both models, a learning bias in favour of one-to-one mappings between meanings and signals was found to be key in driving the cultural evolution of communication. In the feedforward network model, obverter agents have this bias whereas imitator agents do not. Consequently, populations of obverter agents converge on optimal, unambiguous communication systems whereas populations of imitator agents do not. These results were found to hold even in the face of fairly strong natural selection of cultural variants. In the associative network model, only certain weight-update rules (those classified as [+learner, +maintainer, +constructor]) possess the one-to-one bias. Only those weight-update rules, when placed in the context of the ILM, result in the emergence of communicatively optimal systems of meaning-signal mappings.

An examination of the learning biases involved in other models of the cultural evolution of communication shows that this one-to-one learning bias is paramount. More significantly from the point of view of understanding language evolution in human populations, one-to-one biases seem to be brought to bear by human infants when acquiring vocabulary. I have presented arguments that human infants are biased against acquiring synonyms and homonyms. Furthermore, non-human primates appear not to have this bias. This suggests that the one-to-one biases applied by humans to the vocabulary learning task may be unique among primates, and may explain the uniqueness (among the primates) of language as a culturally transmitted, symbolic communication system.