CHAPTER 4

The genetic and cultural evolution of communication

In the previous Chapter I equated the properties of the learning bias required for the cultural evolution of optimal communication with properties of the human language acquisition device, in particular the apparent human biases against synonymy and homonymy during vocabulary acquisition. I made the assumption, for the purpose of Chapter 3, that this human capacity must have evolved through natural selection for communication.

It is time to return to this assumption and investigate whether the learning bias required to support a learned symbolic vocabulary can, in fact, evolve under selection pressure for communicative success. If this is the case then we can form an argument that this aspect of the human cognitive apparatus evolved specifically to support vocabulary acquisition.

Chapters 2 and 3 focussed on models of cultural transmission. In Section 4.1 I introduce a simple mathematical model of genetic transmission. This is synthesised with the general model of cultural transmission discussed in Chapter 2 to form B&R's *dual transmission model* (Section 4.2). The Iterated Learning equivalent of the dual transmission model is the Evolutionary Iterated Learning Model, outlined in Section 4.3. In Sections 4.4 and 4.5 I describe how the two ILMs which form the basis of Chapter 3 can be expanded to Evolutionary Iterated Learning Models, to investigate the possibility of coevolution between a culturally-transmitted communication system and the genetically-encoded learning bias facilitating the acquisition of such a system.

4.1 Modelling genetic transmission

The most basic models of genetic transmission are based around three processes:

$$\begin{tikzcd} G_t \arrow[r, "\text{ontogeny}"] & \begin{pmatrix} F_t \\ G_t \end{pmatrix} \arrow[r, "\text{selection}"] & \begin{pmatrix} F'_t \\ G'_t \end{pmatrix} \arrow[r, "\text{transmission}"] & G_{t+1} \end{tikzcd}$$

Figure 4.1: A simple model of genetic transmission. G_t gives the distribution of genotypes at time t. This leads, via developmental processes, to an associated distribution of phenotypes, F_t. The environment then takes its toll, removing some phenotypes and their associated genotypes. This gives a new distribution of phenotypes ${F_t}'$ and genotypes G'_t. The remaining genotypes then lead, via reproduction, to the distribution of genotypes at time $t+1$.

- Ontogeny, where the genotype[1] is translated, in the context of the environment, into the phenotype. This depends on 1) the distribution of genotypes population, G_t, and 2) the process of ontogeny.
- Selection, where the environment takes its toll on the population by removing individuals based on the performance of their phenotypes. This depends on 1) the distribution of phenotypes, F_t, and 2) the selection procedure for removing phenotypes from the population.
- Genetic transmission, where the surviving individuals reproduce and transmit their genotypes to their offspring. This depends on 1) the distribution of phenotypes after selection, F'_t, which implicitly gives G'_t, the distribution of genotypes of surviving individuals, and 2) the procedure by which new genotypes are formed during reproduction.

This cycle is illustrated in Figure 4.1.

4.1.1 Natural selection on genetic transmission

The simplest models of natural selection acting on genetic transmission deal with the changes in frequency of alleles of a single gene in asexually-reproducing haploid populations — each individual has a single gene drawn from a set of n alleles and each individual inherits the allele of their single parent. In sexually-reproducing diploid populations the equations are complicated by the fact that each individual has two alleles for each gene and receives one allele from each of their two parents.

---

[1] A clarification of terminology is perhaps useful here:
Genotype: The particular genetic makeup of an individual.
Phenotype: The particular physical, non-genetic makeup of an individual, which will be a consequence of an individual's genotype in interaction with the environment.
Genome: The specification of possible genotypes, which limits the range of possible genotypes.
Phenome: The specification of possible phenotypes, which limits the range of possible phenotypes.

Ontogeny is typically treated in a very simplistic manner in mathematical models of population genetics. In the haploid organism, single gene case there are n distinct alleles and therefore n distinct genotypes $G_1 \ldots G_n$. It is typically assumed that there are n distinct phenotypes $F_1 \ldots F_n$ and ontogeny maps genotype G_i onto phenotype F_i. Selection then acts on the phenotype, but since there is a one-to-one correspondence between genotypes and phenotypes we can talk of selection acting on genotypes and effectively ignore ontogeny.

The *fitness* of such a genotype is the average probability of individuals with that genotype surviving to reproductive age, multiplied by the average number of offspring that each genotype of reproductive age produces. Now consider a population with two genotypes G_a and G_b with fitness f_a and f_b respectively. Evolution by natural selection takes place in such a population where the two genotypes do not reproduce at equal rates — $f_a \neq f_b$. As shown in Section A.2 in Appendix A, if $f_a > f_b$ then genotype G_a will increase in frequency, and if $f_a < f_b$ then it will decrease in frequency — the fitter genotype comes to dominate the population. The rate of change is at a maximum when genetic diversity is at a maximum, which occurs when both genotypes occur with equal frequency — in other words, natural selection depends on genetic diversity, and the rate of evolution is higher when the population exhibits more diversity.

4.1.2 *Models of the genetic transmission of communication*

Models of the genetic evolution of communication can be roughly divided into two main groups — those which address the question "when should we expect to see communication?" and those which ask "what structure should we expect communication to exhibit?".

4.1.2.1 *When should we expect to see communication?*

The first question has typically been addressed by theoretical biologists but has recently been tackled by researchers using agent-based modelling techniques. Researchers in this area are concerned with the interlocking issues of signal honesty, altruism and signal costs.

Much of this work has been inspired by, or relates to, Krebs & Dawkins's (1984) conception of signalling as manipulation. Krebs & Dawkins define signals as "means by which one animal makes use of another animal's muscle power" (Krebs & Dawkins 1984:382) — a signal emitted by one animal causes another animal to act in a certain way. This is bound up with their definition of communication (given in Dawkins & Krebs (1978)) that communication systems are systems where the signaller has evolved to induce a response in the receiver which benefits the signaller in terms of inclusive fitness.

Where the interests of the signaller and the receiver are coincident (both signaller and receiver get some benefit), the evolution of signalling seems fairly unsurprising — individuals who signal will be selected for, as will individuals who respond to signalling. One common thread in the formal modelling approach revolves around identifying common interests between signallers and receivers, particularly in situations where no common interest is obvious. For example, alarm call systems obviously offer a benefit to receivers of alarm calls (they can evade a predator they may not have noticed), but the signallers receive no obvious benefit (they have already seen the predator), and may in fact incur a cost (their signalling behaviour may attract the predator). So why do alarm call systems exist? Kin selection, whereby signallers receive an indirect payoff via the genes they share with receivers, has been advanced as an explanation for this apparently altruistic behaviour, based on both empirical evidence from studies of alarm-calling species (Sherman 1977) and evidence from computational modelling (Ackley & Littman (1994) and Oliphant (1996), but see Noble (1999) for a negative result).

Krebs & Dawkins (1984) make the subsidiary prediction that, where the interest of signallers and receivers coincide, there should be coevolution between signaller and receiver so that the signal is cheap for the signaller to produce and the receiver is sensitive enough to pick up on these “conspiratorial whispers”. Analytic and simulation models (Noble 1998) support a slightly modified version of Krebs & Dawkins’s (1984) conspiratorial whispers theory — where both sender and receiver receive a clear payoff then signals should be cheap, whereas if the signaller can expect to receive a large payoff but the receiver only receives a marginal benefit, receivers should be resistant to cheap signals and signallers should therefore evolve to use costly signals.

Matters are more complicated where there is a direct conflict of interest between signaller and receiver, which cannot be resolved by appeals to mechanisms such as kin selection. Consider the case of males who produce signals to indicate to females their quality as a mate. Males, typically, will wish to mate as often as possible whereas females should be choosy, preferring to bear offspring only for high-quality males. In this scenario there is a direct conflict of interest and we would predict that signalling would be useless — any population of “honest” signallers, where the signals of males reliably reflected their quality and the females took the males at their word, would be prone to invasion by “liars”, males who exaggerated their own quality.

One potential solution to this problem is the handicap principle (Zahavi 1975; Zahavi 1977). If the signals males produce are costly, with a cost proportional to their claims of quality, then females can trust the signals — only a male of high quality could afford to bear the cost of producing the signal indicating high quality, therefore individuals producing the high quality signal must be telling the truth. This result has been verified by analytic and simulation models, with one or two qualifications (see, for example, Grafen (1990) and Bullock (1997)).

Alternatively, Krebs & Dawkins (1984) suggest that signalling could persist in the face of a conflict of interests, but result in an evolutionary arms race between signallers and receivers — signallers will be selected to produce increasingly effective manipulative signals, and receivers will be selected for increased resistance to these manipulative signals. This position has received less support from formal models — Noble (1999) presents results from both analytic and simulation models which suggest that communication is never more than a transitory phenomenon when a conflict of interests exists.

While these issues are obviously important, they are somewhat outwith the main area of inquiry of this thesis — I assume that signallers and receivers have coincident interests, and both receive a payoff from successful communication. In part this is due to my focus on the *structure* of communication systems and the learning bias required to bring it about, but also to the observation that issues of honesty, altruism and so on don't seem to have a great impact on the structure of language — people tell the truth and tell lies, and altruistically give and selfishly withhold information all the time, and it's not clear if this has any lasting structural consequences.

4.1.2.2 How should communication be structured?

There are several models which show that communicatively-optimal, innate communication systems (i.e. shared one-to-one mappings) can evolve under natural selection for communicative success (e.g. Werner & Dyer (1992), MacLennan & Burghardt (1993), Levin (1995)). Such models typically assume a pre-existing dedicated communication channel, although it has been shown recently that communication can evolve in the absence of such a dedicated channel, through evolutionary reappropriation of initially non-communicative behaviours (Quinn 2001). Functional communication systems have also been shown to emerge as a consequence of the evolution of internal representations (Cangelosi & Parisi 1998).

Turkel (2002)[2] presents an interesting model, based on Hinton & Nowlan's (1987) influential paper, which demonstrates the occurrence of the Baldwin effect. Turkel interprets his model as a model of language evolution. This is perhaps the weakest point of the model, for reasons which will become obvious. However, the results show interesting behaviour with respect to the evolution of coordinated behaviour, and as such it is relevant to models of the evolution of communication.

In the model, an individual's genotype is a string of 1s, 0s and ?s. The process of genotype-phenotype mapping simply directly maps this genotype to a phenotype string of 1s, 0s and ?s. Individuals then participate in pairwise coordination interactions, which could be considered communicative interactions, where coordination is simply phenotype matching. During each interaction the two individuals involved temporarily[3] replace all the ?s in their phenotype with a 0 or a 1, selected at random. If the two phenotypes then match then the interaction is considered to be a success. Otherwise, a different temporary replacement of ?s with 1s and 0s is attempted. This process repeats until coordination is established or a certain prespecified number of replacement trials is reached, in which case the interaction is considered to be a failure. The level of success of an episode of interaction is inversely proportional to the number of replacement trials required to achieve a phenotype match.

There are several points to note here. Firstly, if, at any given point on the phenotype, one individual has a 1 and the other individual has a 0 then coordination is impossible. Secondly, individuals with large numbers of ?s are likely to take longer to successfully match their phenotypes. Finally, Turkel explicitly penalises individuals who require 0 trials to match their phenotypes (i.e. both individuals have no ?s in their phenotypes and their phenotypes match) — this apparently capricious step is taken to ensure that ?s never disappear from the population, a consequence of Turkel's interpretation of the model as a model of P&P UG.

Turkel's main contribution is to show that aspects of phenotypes which were originally plastic (?s) become fixed (to 1s or 0s), via the Baldwin effect, although plasticity is never fully eliminated, in part due to Turkel's quirky fitness function. The Baldwin effect can be summarised thus: behaviour which is initially plastic becomes innately fixed when plasticity has some cost. In Turkel's model, excess ?s have a cost as they reduce the

---

[2] This paper has in fact been in circulation, and receiving citations, since 1994. However, for a variety of reasons, it has only recently been published. This explains the fact that, for example, Kirby & Hurford present an extension of a model in 1997 that was not published till 2002.

[3] The replacement is only temporary, therefore there is no true cultural transmission in the population — an individual's phenotype after a series of interactions is exactly the same as its phenotype before those interactions.

coordination payoff individuals can expect to receive. Parts of the genotype which began as ?s therefore tend to be replaced by 1s or 0s, depending on the genetic makeup of the rest of the population. This model therefore demonstrates that the Baldwin effect can potentially lead to the transfer from a learned coordinated system of behaviour (perhaps a communication system) to a largely innate, coordinated system of behaviour. Yamauchi (2001) demonstrates that this can only happen when the relationship between genotypes and phenotypes is fairly simple — when a single phenotypic trait is determined by several genes, or when a gene determines several phenotypic traits, the impact of the Baldwin effect is reduced.

4.2 The dual inheritance model

How can models of genetic transmission be combined with models of cultural transmission to produce a unified model of transmission? B&R propose the *dual inheritance model*, sketched in Figure 4.2.

The population at time t is defined by F_t, a distribution over phenotypes, and G_t, a distribution over genotypes. G_t is derived from G'_{t-1}, the distribution of genotypes of individuals of reproductive age at the previous generation. The process of ontogeny and cultural transmission then give F_t, the phenotype distribution, based on:

- the mechanisms of ontogeny,
- G_t, the distribution of genotypes in the population,
- the mechanisms of cultural transmission and
- F'_{t-1}, the distribution of phenotypes of individuals of reproductive age at the previous generation.

In a dual transmission model the distribution of genotypes in the population is relevant because genetic variants are taken to have an influence on the process of cultural transmission — for example, individuals with a particular genotype may not take part in cultural transmission, or may prefer to acquire a particular cultural trait.

The dual inheritance model then proceeds in a similar fashion to the general model of cultural transmission given in Chapter 2. F_t^l gives the distribution of phenotypes in the population once individual learning has been taken into account. This depends on two factors: 1) F_t, the distribution of phenotypes prior to individual learning; 2) the process of individual learning, by which individuals change their phenotype in response to the

$$\begin{tikzcd}
\begin{pmatrix} F'_{t-1} \\ G'_{t-1} \end{pmatrix} \arrow[rr, "\text{cultural transmission}"] \arrow[dr, "\text{mating}"'] & & \begin{pmatrix} F_t \\ G_t \end{pmatrix} \arrow[r, "\text{individual learning}"] & \begin{pmatrix} F^l_t \\ G_t \end{pmatrix} \arrow[r, "\text{selection}"] & \begin{pmatrix} F'_t \\ G'_t \end{pmatrix} \\
& G_t \arrow[ur] & & &
\end{tikzcd}$$

Figure 4.2: The dual transmission model. The genotype distribution of the population at time t, G_t is determined by the the reproduction process and the distribution of genotypes in the population at time $t-1$. Cultural transmission, constrained by the genetic makeup of the population (G_t) then yields the distribution of cultural variants F_t. This process is dependent on the mechanisms of cultural transmission, the influence of genes on the cultural transmission process, and the distribution of cultural variants in the population at time $t-1$.

environment. F'_t and G'_t give the distribution of phenotypes in the population after removal of individuals through death has been taken into account. These depend on: 1) F^l_t and G_t, the distribution of phenotypes and genotypes prior to death; 2) the process of differential retention, by which some individuals survive and some individuals are removed due to environmental factors. Note that the distribution of genotypes remains unaffected by cultural transmission and individual learning.

4.2.1 The genetic transmission of direct bias

B&R consider the circumstances under which a biological capacity for individual learning and biased and unbiased cultural transmission will be favoured by natural selection. In general, their technique is to assume that the capacity for the behaviour of interest is genetically encoded, and that there are two possible alleles. Individuals with one allele will participate in the relevant behaviour, whereas individuals with the other allele will not. B&R construct equations to see under what circumstances the allele for the particular behaviour will be favoured by natural selection on genetic transmission.

For the purpose of this thesis it is sufficient to review their model of the genetic evolution of direct bias. Recall from Chapter 2 that direct bias on cultural transmission will increase the frequency of the favoured variant in a population, with the rate of increase depending on the strength of the direct bias and the cultural variance in the population. B&R expand this model, following their general technique outlined above, to consider the case where an individual's genotype determines their preference for cultural variants.

Let us assume that there are two cultural variants, c and d, and two genetic variants, e and f. e is the unbiased genotype, and f is the biased genotype. In other words, individuals with genotype e will happily acquire cultural variants c or d, whereas individuals with genotype f will prefer to acquire one particular cultural variant, let's say c.

As shown in Section A.3 of Appendix A, if genetic parents are selected at random (there is no natural selection acting on genetic transmission) then, as we would expect, the frequency distribution of genotypes in the population remains unchanged. Among individuals with the biased allele f cultural variant c increases in frequency according to the strength of the bias and the cultural variance in the parent population, and variant d decreases by a similar factor. Among individuals with the unbiased allele e the frequency of the two cultural variants remains unchanged. In other words, individuals with the unbiased genotype acquire either cultural variant, whereas individuals with the biased genotype preferentially acquire variant c at the expense of variant d.

B&R then go on to add natural selection to the model. Natural selection weeds individuals out after cultural transmission and prior to breeding. Let us assume that individuals in possession of cultural variant c receive some fitness payoff s (and are therefore more likely to reproduce genetically and act as cultural parents), and individuals with the biased genotype f incur some cost z for that bias (and are therefore penalised if $z > 0$).

Assume for a moment that the proportion of individuals with cultural variant c, is fixed at some arbitrary value. What happens to the frequency of individuals with the biased genotype? As shown in Section A.3 of Appendix A:

- if the biased genotype has no associated cost ($z = 0$):
  - if the population exhibits no cultural variation then the biased allele has no fitness advantage over the unbiased allele and does not change in frequency.
  - if the population exhibits cultural variation then the biased allele will increase in frequency.
- if the biased genotype has a cost ($z > 0$):
  - if the population exhibits no cultural variation then the biased genotype will decrease in frequency — the biased allele will suffer a fitness penalty due to its cost and no fitness benefit over the unbiased allele due to the lack of cultural variation.

- if the population exhibits cultural variation then the biased genotype will either be selected for or against, dependent on the relative strength of the direct bias, the cultural makeup of the population, the advantage of possessing cultural variant c and the cost associated with the bias.

To summarise, in a population which is completely converged culturally (on either variant) the frequency of the biased variant should either remain constant (if biased learning is costless relative to unbiased learning), or decrease (if biased learning has a cost).

What can we predict about the frequency of cultural variant c? As discussed in Section A.3 of Appendix A, variant c is always favoured by selection, and by biased transmission when there are individuals with the biased genotype in the population. Therefore variant c will increase in frequency until the population completely converges on c. As discussed above, at this equilibrium state the biased genotype either has no advantage over the unbiased genotype or is at a disadvantage (where $z > 0$). Therefore, at equilibrium we should expect selection to either be neutral with respect to bias, or to see only the unbiased allele — directly biased transmission pushes the population to converge on the favoured cultural variant, at which point selection pressure on the population's genotypes either stops, or acts to reduce the frequency of the biased allele which drove cultural convergence in the first place.

4.3 The Evolutionary Iterated Learning Model

The Evolutionary Iterated Learning Model (EILM) is an extension of the ILM to model the dual transmission of cultural and biological traits — it is the Iterated Learning equivalent of B&R's dual transmission model. As in the ILM, populations consist of collections of individuals, who acquire their linguistic competence on the basis of the observed behaviour of other individuals. Individuals also inherit their genetic endowment from other members of the population, their biological parent(s). As with the ILM, population turnover can either be generational, as depicted in Figure 4.3, or gradual, as in Figure 4.4.

4.3.1 The evolutionary iterated learning of vocabulary

Hurford (1989) presents the original EILM. Hurford assumes three possible genetically-encoded learning strategies, which he calls Imitators, Calculators and Saussureans. As discussed in Chapter 3, Section 3.5.3, based on experiments outlined in Hurford (1989) and an analysis of these three learning strategies, Calculators can be classified as [$-$learner, $-$maintainer, $-$constructor] and Imitators and Saussureans can be classified as [$+$learner,

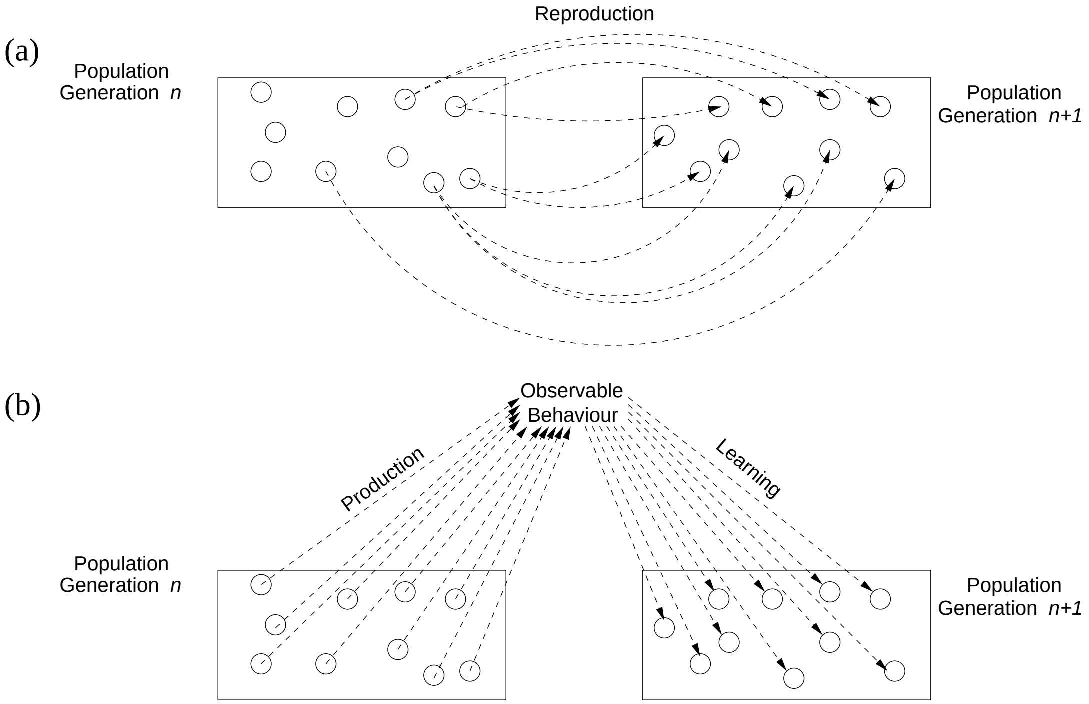

Figure 4.3: The generational EILM. (a) illustrates the process of genetic transmission — individuals at generation $n + 1$ (represented by circles) inherit their genotypes via genetic transmission (arcs) from individual at generation n. (b) illustrates the process of cultural transmission — the generation n individuals produce a set of observable behaviour which is observed and learned from by generation $n + 1$ individuals.

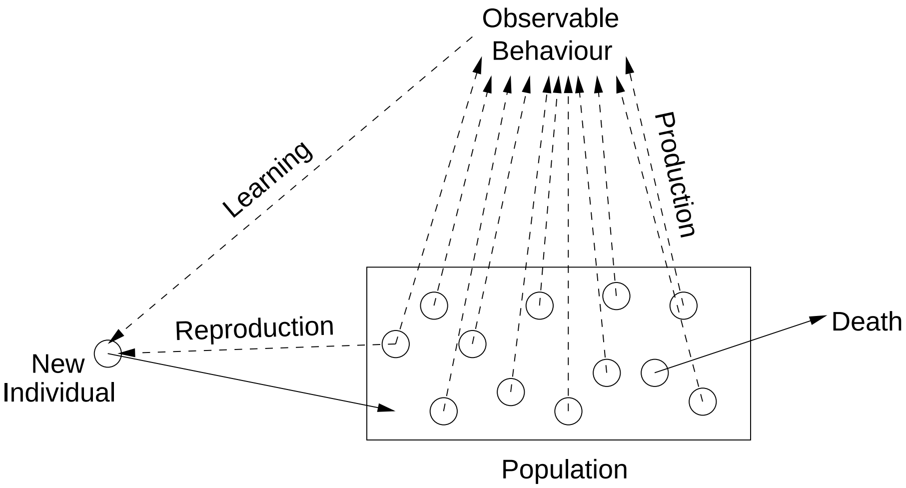

Figure 4.4: The gradual EILM. A single individual is removed from the population, representing death. The remaining individuals then reproduce to produce a new individual, who inherits its genotype from its parent(s). The new individual then learns based on the observable behaviour produced by the population, and enters the population.

+maintainer, −constructor]. Saussureans have the additional advantage over Imitators that their production and reception behaviour are necessarily closely coupled.

Hurford conducts three sets of experiments, the first two of which have been discussed in Section 3.5.3 and will not concern us here. In the third set of experiments, using a generational model of population turnover, Hurford investigates how natural selection during genetic transmission of learning strategies affects genetically heterogeneous populations. Hurford's simulation runs start with each meaning–signal pair being equiprobable, with an equal distribution of the three genetically-encoded learning strategies. There is strong selection during genetic transmission, with individuals who are more successful at communicating with other members of the population being more likely to breed. Under such circumstances, the Saussurean learning strategy comes to dominate each of 20 simulated populations — natural selection during genetic transmission leads to the evolution of the 'optimal' learning bias. However, Hurford (1989) does not report the frequencies of the various genotypes over the course of the runs, and does not retain the original data (James Hurford, personal communication). This is a pity, as we will see in Section 4.5 that the manner in which the population arrives at the final state is in fact more telling than the final state itself.

Kvasnička & Pospíchal (1999) present an EILM similar to the model which I will describe in Section 4.4. Briefly, Kvasnička & Pospíchal use a feedforward network model of a communicative agent, with the obverter architecture. Individuals acquire their communication system based on the observable behaviour produced by the previous generation of the population. The connection weights in an individual's network, together with some limited aspects of network topology, are specified genetically, and opportunity to breed is determined by communicative accuracy. Kvasnička & Pospíchal report that optimal communication emerges from initially random behaviour, and conclude that this is a consequence of gene-culture coevolution. However, it is not clear from their model that there is actually any gene-culture coevolution occurring. While the network topology may indeed evolve to facilitate communication (a similar result is reported by Livingstone & Fyfe (2000)), it is not clear to what extent the transmission of initial connection weights plays a role in biasing the learners towards acquiring optimal communication systems. As discussed in Chapter 3, the obverter network architecture builds in a bias in favour of one-to-one mappings from meanings to signals — optimal communication can emerge in populations of obverter networks in the absence of any genetic transmission of connection weights. The relative importance of pressures acting on genetic and cultural transmission is not adequately decoupled in Kvasnička & Pospíchal's (1999) model.

EILMs have also been used to investigate gene-culture interactions in the evolution of structured communication — for example, the models described in Kirby & Hurford (1997) and Briscoe (2000b). Discussion of these models will be delayed until Chapter 6.

4.4 Model 1: Adding genetic transmission to the feedforward network model

In this Section, I describe an extension of the feedforward network model outlined in Section 3.3 which can be used to investigate possible interactions between the genetic transmission of learning bias and the cultural evolution of communication. This research appears in Smith (in press).

Feedforward neural networks are sensitive to their initial connection weights (Kolen & Pollack 1990) — certain combinations of initial connection weights can make certain input-output mappings unlearnable, whereas other initial weights can reduce the number of exposures required to learn particular mappings. The initial connection weights effectively bias a feedforward neural network towards acquiring certain input-output mappings. This suggests a first step towards investigating the dual transmission of vocabulary systems — if we assume that the initial connection weights in an individual's neural network are genetically determined, then we have a model where genetic information affects that individual's learning bias.

The feedforward network model outlined in Section 3.3 can be straightforwardly extended to include the genetic transmission of initial connection weights, in order to investigate how differential genetic transmission of learning bias impacts on the evolution of unstructured communication. Some of the assumptions arising from this treatment are somewhat dubious, in particular the assumptions that there is a one-to-one correspondence between genotype and phenotype and that genetic information merely provides a starting point for unconstrained learning in the phenotype. Two points can be made in defence of this model. Firstly, the technique of evolving initial connection weights for neural networks is not a new one, and has been adopted for non-linguistic problem domains with some success by, for example, Montana & Davis (1989), Belew *et al.* (1992) and Nolfi *et al.* (1994). Batali (1994) adopts a similar technique to a study of the learnability of structured languages. Secondly, the feedforward network model merely provides a starting point — while it yields some useful results, it will be superseded by the more plausible model outlined in Section 4.5.

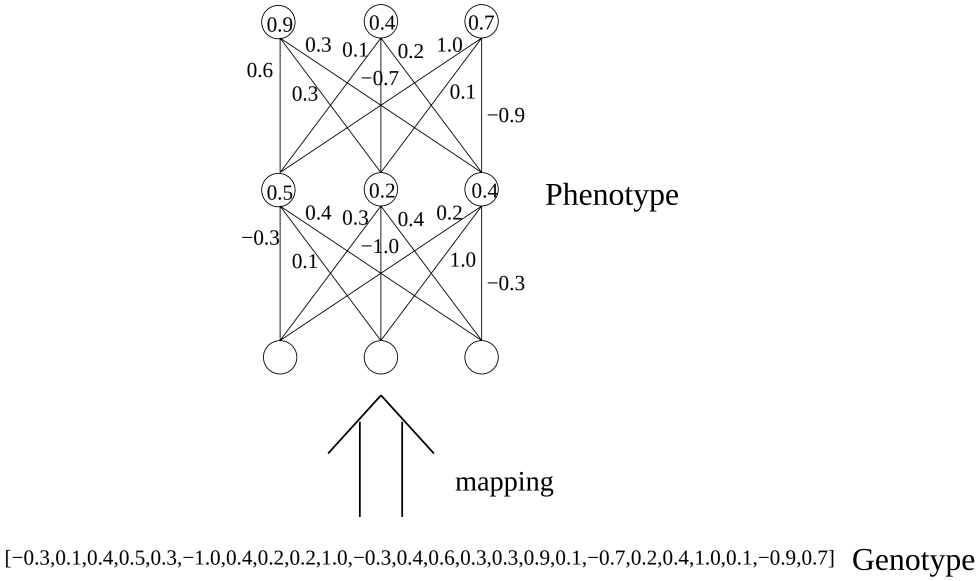

Figure 4.5: Genotype-phenotype mapping, the process of translating a genotype into a phenotype. Each locus on the genome specifies the weight of a connection in the phenome. In the genotype-phenotype mapping process, the allele at each locus in an individual's genotype specifies the initial weight of the associated connection in the individual's phenotype network, as illustrated here. The weights from bias nodes to nodes in the hidden and output layer are shown in the centre of the associated node.

4.4.1 Genotypes, phenotypes and the genotype-phenotype mapping

The EILM requires a model of the genome and a model of the phenome. The phenome model is simply the feedforward network model of a communicative agent outlined in Section 3.3. The genome is a chromosome with 24 loci. The possible alleles for each locus are the set of real numbers.

Ontogeny, in the terms outlined in Section 4.1, involves the mapping from a genotype to a phenotype. Each individual consists of a genotype-phenotype pair. Each locus in the genome corresponds to a connection in the phenome network. In the process of genotype-phenotype mapping the allele at each locus in an individual's genotype specifies the initial weight of the associated connection in the individual's phenotype network. This mapping process is illustrated in Figure 4.5.

4.4.2 Reproduction

Individuals inherit their genotype from their parents. In this model the genome is haploid, so asexual reproduction would be a reasonable and simple option. However, previous research suggests that asexual reproduction in small populations tends to result in a rapid loss of genetic diversity and a corresponding weakening of evolutionary pressures (Mitchell 1996). A sexual model of reproduction is therefore used, to attempt to maintain genetic diversity in the population.

Each individual has two parents. The parental genotypes are combined to form a single offspring genotype using one-point crossover[4]. Mutation[5] during reproduction is also included, once again to maintain genetic diversity in the population.

4.4.3 The EILM algorithm

As with the ILM described in Section 3.3, in the feedforward network EILM a generational model of population turnover is used. Populations consist exclusively of either imitator or obverter networks. The EILM consists of an initialisation process and an iteration process.

Initialisation Create a population $population_{g=0}$ of N agents[6]. Each agent is either an imitator or obverter, with populations being homogeneous in this respect. Each initial agent has a random genotype, with the allele at each locus selected randomly from the range $[-1, 1]$. The population is heterogeneous in this respect. Each initial agent's phenotype is determined by their genotype and the genotype-phenotype mapping process discussed in Section 4.4.1 .

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 $ca\,(P, R, m_i)$, for every communicatively relevant meaning $m_i \in \mathcal{M}_{CRS}$.
2. For every member of $population_g$, generate a set of meaning-signal pairs by applying the network production process to every $m \in \mathcal{M}_{CRS}$.

---

[4] Crossover occurs with probability p_x ($p_x = 0.95$ in all simulations outlined in this section). When crossover takes place, the alleles for the first n loci of the offspring's genotype are inherited from the first parent and the remaining alleles are inherited from the second parent. n is randomly selected to be between 1 and $l_g - 1$, where l_g is the length of the genome. When crossover does not take place the whole genotype is inherited from the first parent.

[5] Point mutations occur on the newly-formed genotype with probability p_m ($p_m = \frac{0.1}{l_g}$ for all simulations outlined in this section). Mutation involves replacing the allele at the mutated locus, which has a value a, with an allele with the value $a+r$, where r is a random number in the range $[-1, 1]$. This mutation operator, in conjunction with the unrestricted range of alleles, allows the possibility of the emergence of extremely large-valued alleles. However, in practice such alleles do not occur. In the simulations outlined in this section all alleles remain within the range $[-5.41, 5.29]$.

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

3. Create a new population $population_{g+1}$ of N agents of the same type (imitator or obverter) as $population_g$. Each member of $population_{g+1}$ inherits their genotype from two parents from $population_g$, via the reproduction process outlined above. Parents are selected randomly from amongst the fittest b members of $population_g$.
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[7].
5. $population_g$ is removed and replaced with $population_{g+1}$. Return to 1.

In this model there are three possible pressures operating on the population:

1. *Selection for communicative success operating on genetic transmission* (when $b < N$), driven by natural selection, favouring genes whose phenotype realizations are successful communicators.
2. *Selection for communicative success operating on cultural transmission* (when $t < N$), driven by natural selection of communication systems, favouring systems which result in successful communication.
3. *Selection for learnability operating on cultural transmission*, 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).

How do these pressures interact? The behaviour of populations attempting to construct and maintain an optimal system, for various simulation parameter settings, are discussed in the next two sections.

4.4.4 Emergence of a communication system

Simulation runs were carried out to evaluate the behaviour of imitator and obverter populations with initially random communication systems (given by the random initial weights in the generation 0 population), for various levels of natural selection pressure acting on genetic transmission (dependent on b) and cultural transmission (dependent on t). Runs were allowed to proceed for 1000 generations, with 10 runs being carried out for each set of simulation parameters. In the figures that follow, the measurement of interest is

---

[7] As before, a learning rate of 0.5 is used.

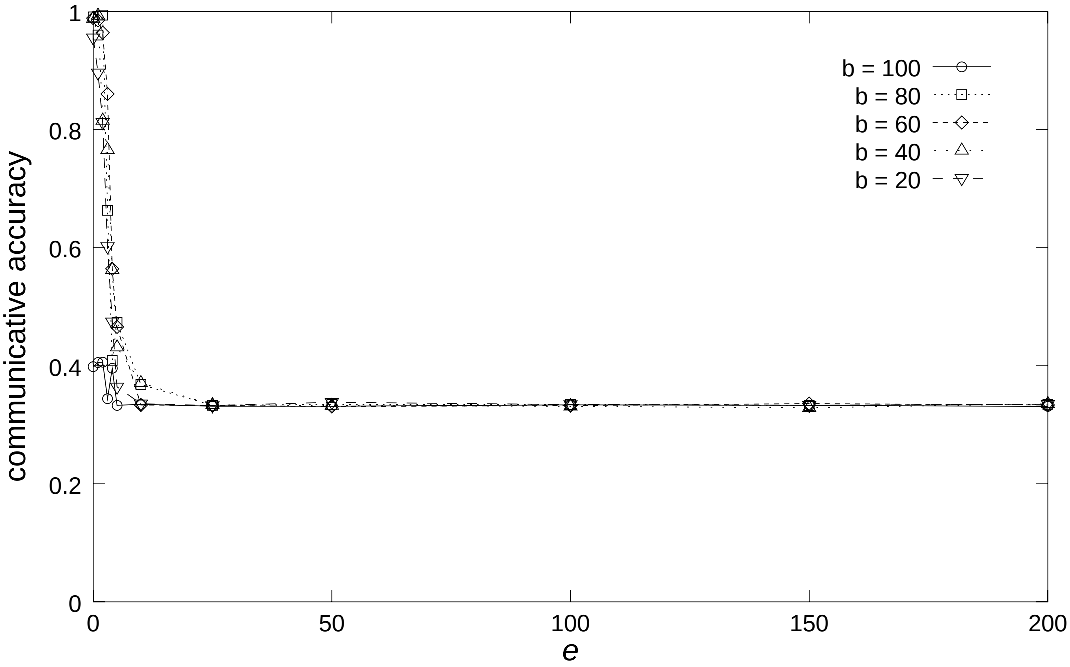

Figure 4.6: The final communicative accuracy of imitator populations in the EILM, in the case where there are varying degrees of natural selection acting on genetic transmission ($b \leq N$) and no natural selection of cultural variants ($t = N$), as a function of e. Communicative accuracy drops rapidly as e increases — cultural transmission clearly has a detrimental effect.

the final average communicative accuracy of the populations. This was obtained by measuring the communicative accuracy of each individual in the final 10 generations of the population, according to the process outlined above, averaging over the population then averaging over the 10 simulation runs carried out with that experimental setting.

4.4.4.1 Imitator populations

Figure 4.6 shows the communicative accuracy of the final systems in imitator populations, for the case where $b \leq N$ and $t = N$, for various values of e (numbers of learning exposures). Results are not shown for the case where $t < N$ — as we might expect given the results discussed in Section 3.3, differential access to cultural parent roles had virtually no impact on the behaviour of the populations.

For the case where $b = N = 100$ (breeding is independent of fitness), the behaviour of the populations is essentially as it was when there was no genetic transmission of initial connection weights — communicative accuracy remains uniformly low. As can be seen from Figure 4.7, for very low values of e and $b < N$ (natural selection on genetic transmission) communicative accuracy does rise significantly above chance levels. When

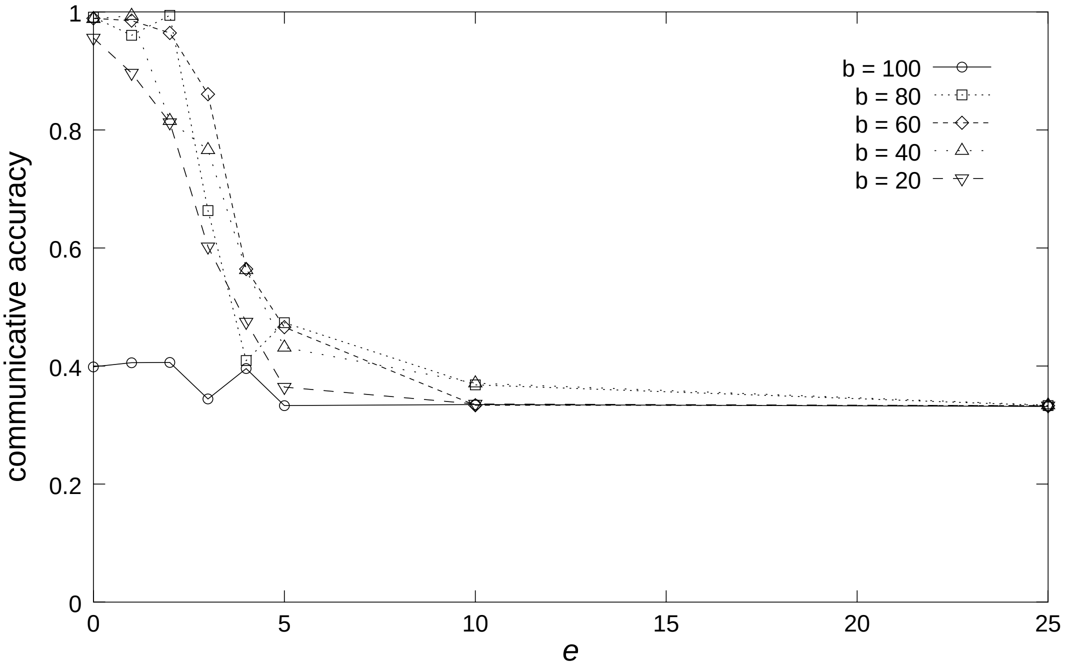

Figure 4.7: The final communicative accuracy of imitator populations where $b \leq N$, $t = N$, for low values of e.

$e = 0$ communicative accuracy reaches optimal levels. This is the result we would expect, given the mathematical and computational models of genetic transmission given above — in the absence of learning, natural selection acts to increase the frequency of fitter genotypes. In the model considered here, fitness is dependent on communicative accuracy, and natural selection therefore acts to increase the frequency of genotypes encoding successfully communicating individuals.

However, as e increases communicative accuracy rapidly drops. For $b < N$ (natural selection on genetic transmission) and $e > 10$ the populations converge on maximally ambiguous communication systems, leading to chance levels of communication. The importance of cultural transmission increases as e increases, and consequently the importance of directly biased cultural transmission increases. As discussed in Section 3.3, imitator individuals are biased in favour of acquiring ambiguous communication systems. When $e = 10$ fully ambiguous systems become completely learnable, and remain more learnable than less ambiguous systems. When this value of e is reached, the pressure for genotypes encoding successful communicators is completely overridden by the pressure on cultural transmission introduced by learner bias, and the populations converge on fully ambiguous communication systems. To put it another way, as e increases, cultural transmission begins to *shield* (Ackley & Littman 1992) genetic information from natural selection — learning makes individuals with distinct genetic makeups appear phenotypically identical, therefore selection cannot identify beneficial genotypes. This is also the situation predicted by B&R — when the population exhibits no cultural variation, selection acting on genetic transmission is neutral with respect to the genetically-encoded direct bias.

4.4.4.2 Obverter populations

Figure 4.8 plots the communicative accuracy of obverter populations, for $b \leq N$ and $t = N$, for various values of e. Again, no significant effect was found when natural selection of cultural variants was introduced — setting $t < N$ leads to similar results.

For the case where $b < N$, natural selection on genetic transmission clearly has a positive effect — the average communicative accuracy of the populations increases above the level for $b = N$. Natural selection fine-tunes initial connection weights so as to increase the learnability of unambiguous systems. For low to intermediate values of e a direct bias on cultural transmission has clearly evolved. However, for high e the obverter network learning bias alone leads to optimal communication. Under these circumstances, there is no selection pressure in favour of an additional direct bias introduced by the networks' initial connection weights — the population's initial connection weights undergo genetic drift. As discussed above, this situation matches the predictions of both B&R and Ackley & Littman (1992) — when the population exhibits no cultural variation, selection acting on genetic transmission is neutral with respect to the genetically-encoded direct bias.

4.4.5 Maintenance of an optimal system

In Chapter 3, I briefly alluded to the behaviour of imitator populations when attempting to maintain an optimal system. I predicted that imitator agents would be unable to do so, due to their bias in favour of many-to-one mappings between meanings and signals. However, no results from the ILM were found to support this prediction. It turns out that some support for this position can be found by considering populations of imitator agents attempting to maintain an optimal system in the EILM.

Figure 4.9 shows the average communicative accuracy of populations of imitator agents who start out with a shared, optimal, innate communication system — all the agents in the initial population have a set of genes which encode an unambiguous communication system. $e = 200$ for all simulations in Figure 4.9. Various amounts of selection pressure (b and t) are used on genetic and cultural transmission. As can be seen from Figure 4.9,

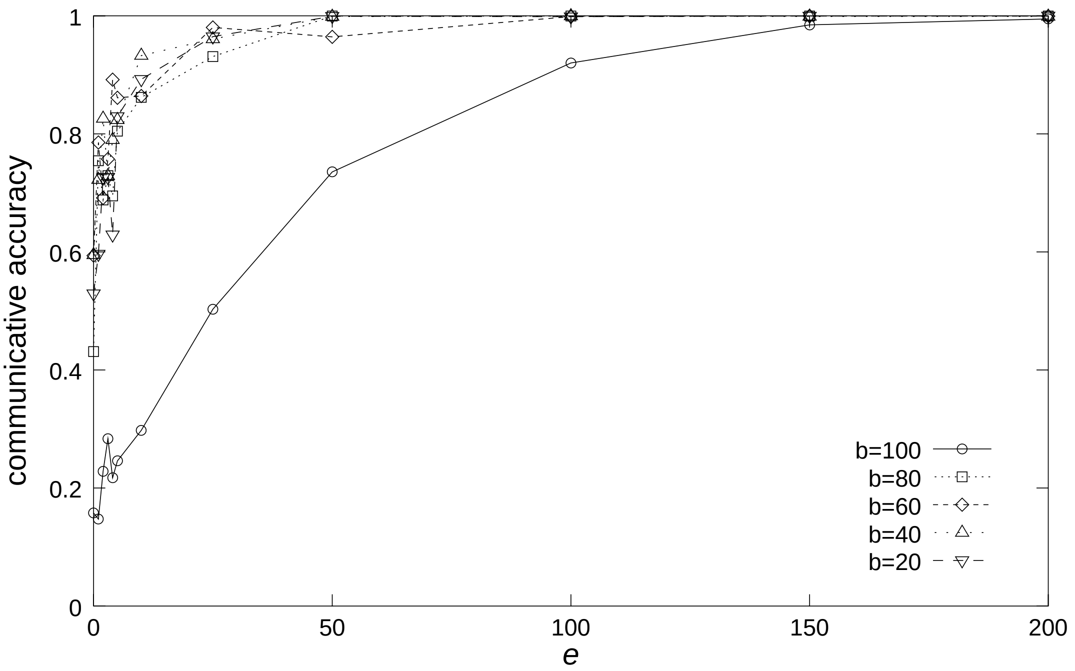

Figure 4.8: The final communicative accuracy of obverter populations where $b \leq N$, $t = N$, for various values of e. Communicative accuracy increases as e increases — cultural transmission is beneficial. Overall levels of communicative accuracy are also higher when $b < N$ — natural selection acting on genetic transmission is beneficial, although the benefit decreases as e increases.

all populations collapse from using an unambiguous communication system to using a partially or fully ambiguous communication system within 25000 generations.

As discussed above, learning in the phenotype almost completely masks an agent's genes but there are certain combinations of genes which make learning a particular communication system impossible. Given the absence of cultural variation in populations with an optimal initial system, genetic transmission occurs at random — there is no selection pressure in favour of a direct bias which results in successful communication. The population's genotypes, and therefore the population's initial connection weights, will undergo drift. In each simulation shown in Figure 4.9 an agent is eventually born whose genes are so bad that they cannot learn the unambiguous communication system in use by the rest of the population. This individual learns a partially ambiguous or fully ambiguous communication system instead. Such agents are unlikely to breed, given that their fitness is usually lower than other agents in the population. Suboptimal communicators do have a negative effect on the fitness of optimally-communicating agents, given that those optimally-communicating agents suffer a penalty for not understanding or being understood by suboptimal communicators, although this will not usually depress the population's fitness enough to allow a suboptimal communicator to breed. However,

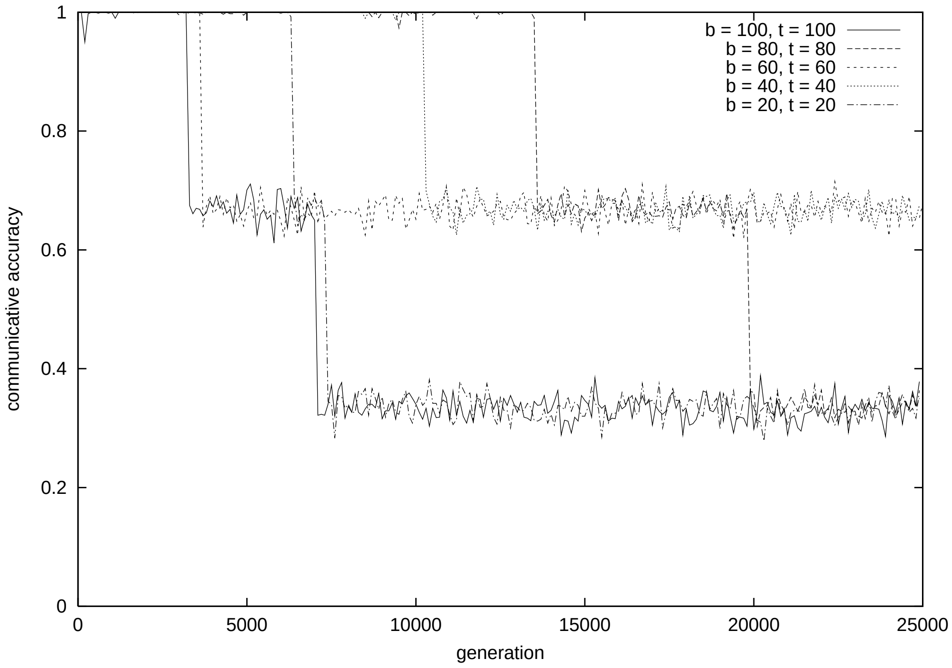

Figure 4.9: Slumping from optimal to suboptimal communication systems in imitator populations in the EILM.

while such individuals are unlikely to breed, their communication systems may be observed and learned from by agents in the next generation, depending on t. Genetic drift therefore introduces cultural variation into the population. This cultural variation then results in the reintroduction of directly biased cultural transmission, resulting from learning bias, which leads to the spread of ambiguous communication systems.

Note that $e = 200$ represents the best-case scenario for imitator agents. $e = 200$ results in the highest level of learnability for unambiguous communication systems and also ensures that, at the early stages of collapse, suboptimal communication systems will constitute only a small part of a learner's observations. In populations where $e \leq 2$ the collapse phenomenon does not occur, but as discussed above the behaviour of these populations is entirely determined by natural selection — they cannot truly be called learning populations.

4.4.6 Summary

The addition of natural selection acting on the genetic transmission of initial connection weights has rather different effects in imitator and obverter populations. In both cases, the learning bias inherent in the two network architectures is still the best predictor of the populations' eventual communication system. In imitator populations, this learning bias favours many-to-one mappings from meanings to signals. The imitator populations therefore converge on fully ambiguous mappings, unless cultural transmission has very little impact (e is low). Furthermore, imitator populations cannot maintain optimal systems in this scenario, due to phenotypic noise introduced by genetic drift.

In obverter populations, in contrast, the addition of natural selection acting on the genetic transmission of initial connection weights has positive effects. When cultural transmission is weak, the addition of natural selection leads to the emergence of initial connection weightings which help establish an optimal, shared communication system. However, this effect diminishes as the strength of cultural transmission increases. Furthermore, comparison of the behaviour of imitator and obverter populations reveals that the correct overall learning bias has to be in place before such beneficial effects are observed.

4.5 Model 2: Adding genetic transmission to the associative network model

While the results outlined in the previous section give an interesting insight into the interaction between directly biased cultural transmission and genetic transmission, they are somewhat unsatisfactory for a number of reasons. There are two forms of direct bias in the feedforward network EILM — one introduced by the network architecture (imitator or obverter), which is externally determined, and one introduced by the genetic transmission of initial connection weights. This second, and much weaker, bias is the only element manipulable by natural selection. We would like the primary bias to be genetically transmitted, to investigate whether biases favouring optimal communication can evolve.

To this end, the associative network model outlined in Section 3.4 is extended to include the genetic transmission of weight-update rules. A small part of this research (Section 4.5.5 in particular) appears in Smith (2001b). Rolls & Stringer (2000) take a similar approach to investigating the evolution of learning, although in their model learning is individual, rather than cultural.

4.5.1 Genotypes, phenotypes and the genotype-phenotype mapping

The phenome is the associative network model of a communicative agent outlined in Section 3.4 of Chapter 3, with $|\mathcal{N}_M| = |\mathcal{N}_S| = 10$. This phenome is defined by the pair $\langle \mathcal{W}, W \rangle$, where $\mathcal{W}$ is an initial set of connection weights and W is a weight-update rule. The genome is a 4-locus chromosome. A genotype is specified by the 4-tuple $(a_\alpha\ a_\beta\ a_\gamma\ a_\delta)$ where a_x is an allele drawn from the set $\{-1, 0, 1\}$.

The process of mapping from a genotype to a phenotype involves converting such a 4-locus chromosome into a $\langle \mathcal{W}, W \rangle$ phenotype. As discussed in Section 3.4.1.3, each weight-update rule W is specified by a 4-tuple $(\alpha\ \beta\ \gamma\ \delta)$. During genotype-phenotype mapping α is set to the value of allele a_α, β is set to the value of allele a_β and so on. The genotype therefore specifies the phenotype's weight-update rule. All $w_{i,j} \in \mathcal{W}$ are set to 0 — every agent has all their initial connection weights set to 0. It would be possible to specify $\mathcal{W}$ in the genotype, as was done in the feedforward network model. However, that option was not explored — the results from the feedforward network model suggest that the genetic evolution of initial connection weights will be of fairly minor importance in comparison to the evolution of the learning bias itself.

To recap, there are 81 possible genotypes, which encode the 81 possible weight-update rules discussed in Chapter 3, Section 3.4. Of these 81 rules, 50 are classified as [$-$learner, $-$maintainer, $-$constructor], 13 are classified as [$+$learner, $-$maintainer, $-$constructor], 9 are classified as [$+$learner, $+$maintainer, $-$constructor] and 9 are classified as [$+$learner, $+$maintainer, $+$constructor].

4.5.2 Reproduction

Individuals inherit their genes from their parents. As with the feedforward network model, organisms are haploid but sexual recombination (involving crossover, in an identical fashion to that outlined for the previous model) is used. Newly-formed genotypes are also subject to mutation.[8] Given the short chromosome it is not clear whether sexual reproduction is necessary to maintain genetic diversity in the population, or whether mutation is sufficient. Simulation runs similar to those described here were carried out with asexual reproduction (each individual has a single genetic parent) and the results were found to be qualitatively similar.

4.5.3 The EILM algorithm

As with the associative network ILM discussed in Chapter 3, a gradual population turnover model is used in the associative network EILM. The EILM consists of initialisation and iteration processes.

---

[8] Point mutations occur on the newly-formed genotype with probability p_m ($p_m = \frac{0.04}{l_g}$ for all simulations outlined in this section, where l_g is the length of the genome.) Mutation involves replacing the allele a_i at the mutated locus with another allele $a_{j \neq i}$, where a_j is selected from the set of possible alleles.

Initialisation Create a population of N agents[9]. Each initial agent has a random genotype, with the allele at each locus selected randomly from the range of possible alleles. Each initial individual's phenotype is determined by their genotype and the genotype-phenotype mapping.

Iteration

1. Select an individual from the population according to the death procedure outlined below 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 meaning $m \in \mathcal{M}$.
3. Create a new agent. The new agent inherits its genotype from its parents, who are selected from the population according to the reproduction procedure outlined below.
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.

The iteration process requires procedures for selecting which individuals die and which individuals reproduce. In the ILM version of the model death was random, and reproduction was effectively ignored. In the feedforward network EILM reproduction occurred with equal probability among the fittest b members of the population. This rank selection procedure would be acceptable for the gradual EILM described here. However, for reasons of computational efficiency [10] a *tournament selection* procedure was used. During a tournament, T individuals are selected from the population at random and evaluated. Each individual is scored according to their average communicative accuracy when acting as both producer and receiver with two randomly selected partners according to the measure $ca\left(P, R\right)$ given in Section 3.2 in Chapter 3. During selection to decide death, the individual with the *lowest* communicative accuracy from among the T selected individuals 'wins' the tournament and is removed from the population. During selection to decide reproduction, the individual with the *highest* communicative accuracy wins the tournament and reproduces.

---

[9] $N = 100$ for all models outlined in this section

[10] As a new individual is introduced at every cohort, communicative accuracy for the entire population would have to be reevaluated at every cohort if the rank procedure was used.

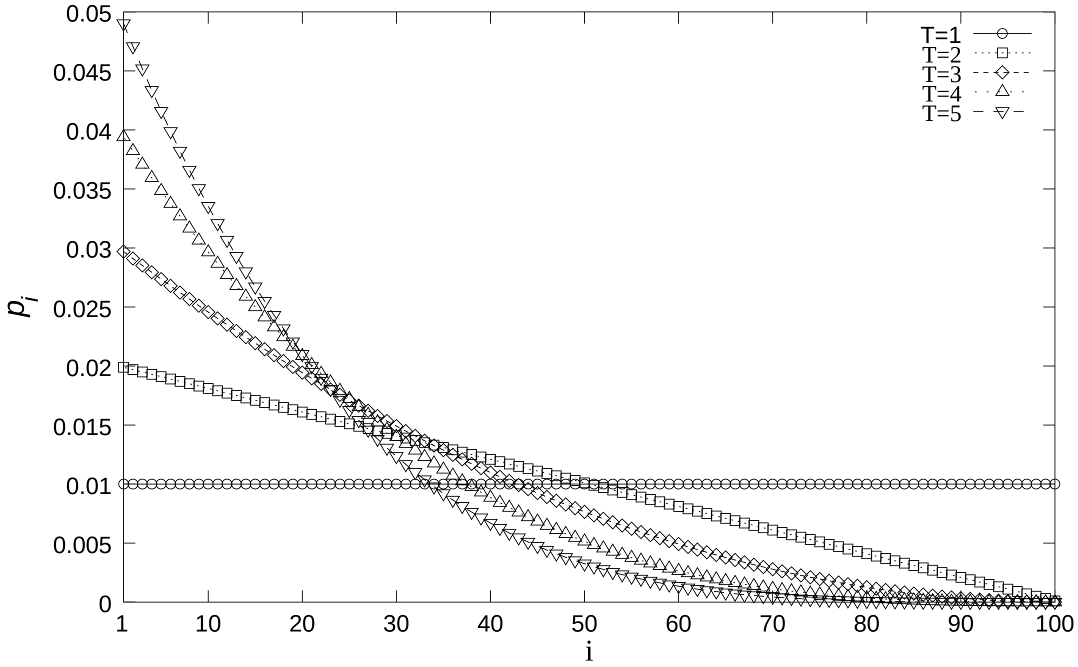

Figure 4.10: Probability of breeding or dying according to rank fitness, for various values of T. i gives the rank fitness of an individual and p_i gives the probability of that individual winning a tournament. For death tournaments, the *least* fit individual has rank 1 and the most fit individual has rank 100. For breeding tournaments, the *most* fit individual has rank 1. Based on the equation from Bäck (1994): $p_i = N^{-T}\left((N-i+1)^T - (N-i)^T\right)$

The strength of selection pressure in the population therefore depends on T. Figure 4.10 shows the probability of individuals of different rank fitness reproducing for various values of T. $T = 3$ was selected for the simulations outlined in this section — it offers a reasonable selection pressure, while still allowing less fit individuals to reproduce, therefore helping to maintain genetic diversity in the population.

One note should be made here regarding the natural selection of cultural variants. The individual who is removed from the population at each cohort does not produce observable behaviour. This was due to reasons of continuity with previous work (e.g. Oliphant (1999)) and simplicity of coding. However, given that removal is dependent on communicative accuracy in the EILM, this introduces natural selection of cultural variants during cultural transmission — less successful communicators are more likely to be removed prior to cultural transmission, and therefore unsuccessful communication systems are selectively removed from the population's culture. It turns out that this does not affect the results reported here. Firstly, it is only a very weak selection pressure (see Figure 4.11). Secondly, as we saw in Chapter 3 with respect to the feedforward network model,

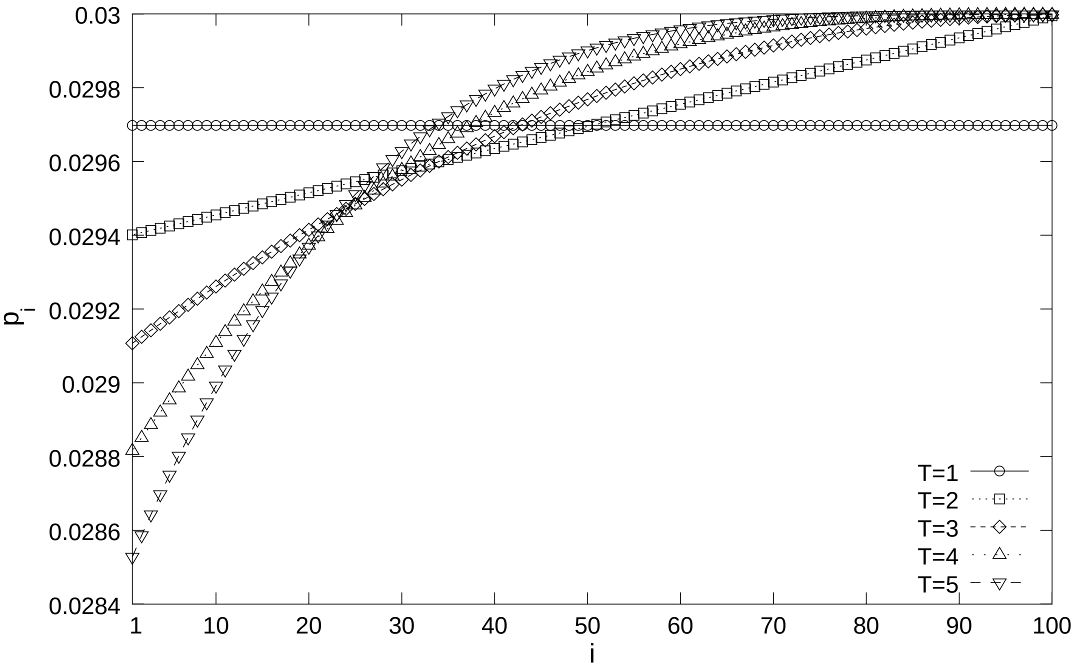

Figure 4.11: Probability of acting as a teacher according to rank fitness, for various values of T. Note the narrow range of the y-axis. i gives the rank fitness of an individual and p_i gives the probability of that acting as a cultural parent. Here the *least* fit individual has rank 1 and the most fit individual has rank 100. The probability of an individual acting as a cultural parent is given by the probability of them not dying multiplied by the probability of being selected to act as a cultural parent — $p_i = \left(1 - \left(N^{-T}\left((N-i+1)^T - (N-i)^T\right)\right)\right) \cdot \left(1 - \left(1 - \frac{1}{N-1}\right)^e\right)$. For the plots shown here, $N = 100$, $e = 3$. There is little difference between the probabilities of the least fit and most fit individuals acting as cultural parents — there is very little natural selection of cultural variants.

natural selection of cultural variants tends to be drowned out by the pressures introduced by learning bias.

4.5.4 Main result: optimal communication rarely emerges

100 runs of the EILM were carried out for $e = 3$ (the same amount of exposures as used in the ILM runs outlined in Section 3.4). Each run proceeded for 5000 cohorts. Figure 4.12 shows the progress of a run of the simulation where the population constructed an optimal[11] communication system. This is a typical example of a successful run. However, only 3 of 100 runs were successful in this respect. Natural selection does not reliably lead

---

[11] The definition of "optimal" was weakened in comparison to the ILM construction tests, to allow for the fact that populations would be genetically heterogeneous and therefore might suffer reduced communicative accuracy due to the presence of non-learners. A population was considered as having converged on an optimal system if communicative accuracy exceeded 0.95 within the 5000 cohorts, and once communicative accuracy exceeded this threshold it remained above it. Weaker definitions of optimal systems yield qualitatively similar results.

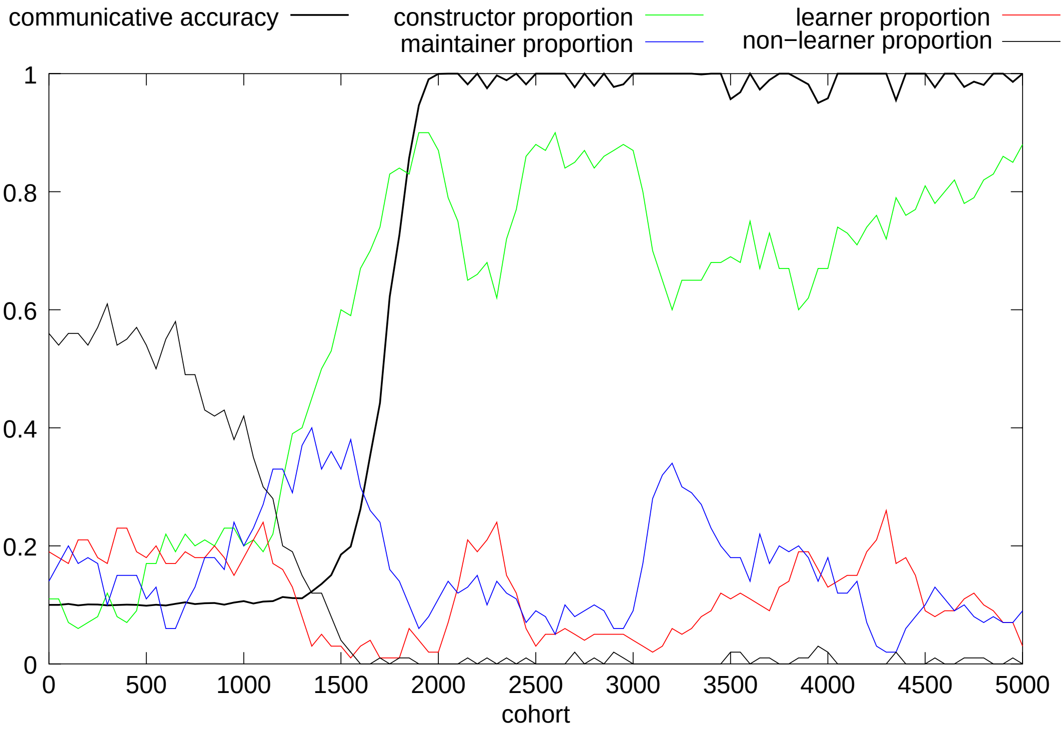

Figure 4.12: The evolution of learning biases leading to optimal communication. Proportions of all the classes of genotypes are given (for example, the black line gives the proportion of the population who use one of the 50 weight-update rules classified as [−learner]). This is a plot of a successful run of the EILM. However, such successful runs are rare.

to the evolution of optimal communication — optimal systems were found to be unlikely to emerge.

Why does natural selection during genetic transmission have such difficulty in identifying weight-update rules which lead to optimal communication? The problem is that the cultural construction of a communication system takes time. Individuals with [+constructor][12] weight-update rules gradually converge, over hundreds of cohorts, on increasingly unambiguous areas of communication system space. In the early stages

---

[12] I will be frequently referring to the various classifications of weight-update rules for the remainder of this Chapter, and writing out the full [±learner, ±maintainer, ±constructor] specification becomes somewhat laborious. I will therefore adopt two simplifying conventions, one which exploits the hierarchical arrangement of the classification system, and one which is more intuitive. The conventions are summarised below:

Full specification	Abbreviated specification	Name
−learner, −maintainer, −constructor	−learner	non-learner
+learner, −maintainer, −constructor	+learner, −maintainer	learner
+learner, +maintainer, −constructor	+maintainer, −constructor	maintainer
+learner, +maintainer, +constructor	+constructor	constructor

of this construction process individuals whose genotypes encode [+constructor] weight-update rules will have little fitness advantage over other individuals. As a consequence the genetic transmission process will be essentially random — the population will undergo genetic drift. In successful runs, such as that shown in Figure 4.12, genetic drift preserves genotypes encoding [+constructor] rules, by chance, in sufficient numbers for sufficient time to allow the construction process to get well under way. Individuals with genotypes encoding [+constructor] rules then show increased communicative accuracy, which leads to steady selection for such genotypes, which increase in number. In the more common, unsuccessful runs, drift does not provide constructor agents in sufficient numbers for sufficient time for the construction process to get under way.

Interestingly, when the population's communicative accuracy nears optimal levels, selection for genotypes encoding [+constructor] weight-update rules stops. The population enters a second stage of genetic drift, where constructor numbers fluctuate randomly. This is due to the fact that maintainers are capable of putting the finishing touches on a communication system and maintaining that system once it is established. Learner agents are also able to acquire the population's optimal system once it is established. Constructor agents lose their fitness advantage over maintainers and learners and genetic transmission becomes semi-random once more. However, non-learners never drift back into the population — they are incapable of learning an optimal system and suffer a severe fitness penalty.

This three-stage drift-selection-drift pattern is common over all successful runs and is illustrated schematically in Figure 4.13. In Figures 4.14–4.17 the relative communicative accuracy[13] of the four classes of genotypes are plotted against their change in number, with the population's average communicative accuracy and the proportion of constructor agents plotted for reference. As can be seen from Figure 4.14, the relative communicative accuracy (henceforth *rca*) of constructor agents fluctuates around 1 till approximately 1000 cohorts through the simulation. We might be tempted to think, given that constructor *rca* is sometimes slightly above 1, that constructors are being selected for. However, this is confounded by the fact that increases in constructor numbers brought about by drift lead to an increase in the communicative accuracy of constructor agents — drift is the cause of the fluctuations in constructor *rca*. More reliable indicators are given by the *rca* values of maintainers, learners and non-learners, for whom this effect is weaker

---

[13] The relative communicative accuracy of a particular weight-update rule is the communicative accuracy of individuals with that weight-update rule divided by the average communicative accuracy of the population. The relative communicative accuracy of a classification (constructors, maintainers etc) is the average of the relative communicative accuracies for all weight-update rules which fall into that classification. Relative communicative accuracy of greater than 1 indicates above-average performance.

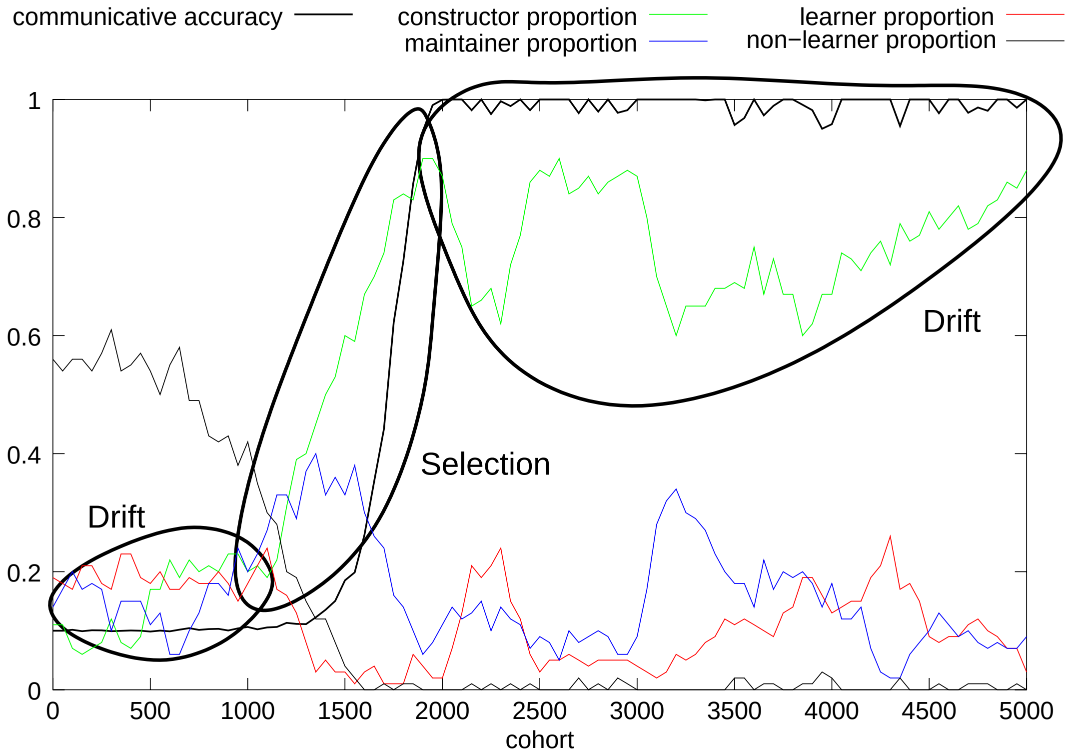

Figure 4.13: The three phases of a successful run. Drift establishes constructor agents in sufficient numbers for the construction process to get underway. Constructor genotypes are then preferentially selected, and communicative accuracy increases rapidly. When an optimal system is established, constructors loose their selective advantage over maintainers (and, to a lesser extent, learners), and a second period of drift occurs.

or non-existent. Learners and non-learners in particular show *rca* of around 1 until 1000 cohorts into the run, by which point the mini-plateau of constructor agents has been established by drift. At this point the *rca* of these individuals plunges, overall communicative accuracy increases markedly and constructor numbers increase due to natural selection for constructor genotypes. After communicative accuracy reaches 1, the *rca*s of learner, maintainer and constructor individuals hover at 1, with non-learners having very low *rca* whenever they are reintroduced by crossover or mutation.

The second period of genetic drift follows the predictions of B&R — once the population reaches cultural convergence, the selection pressure for biases disappears. However, in this model there is still strong selection pressure against [−learner] individuals, who cannot acquire the cultural variant in use in the population.

The first period of drift is not predicted by B&R's mathematical model, which suggests that, given costless learning, the biased allele should increase in frequency when the population is culturally heterogeneous. This turns out not to be the case in this model, for

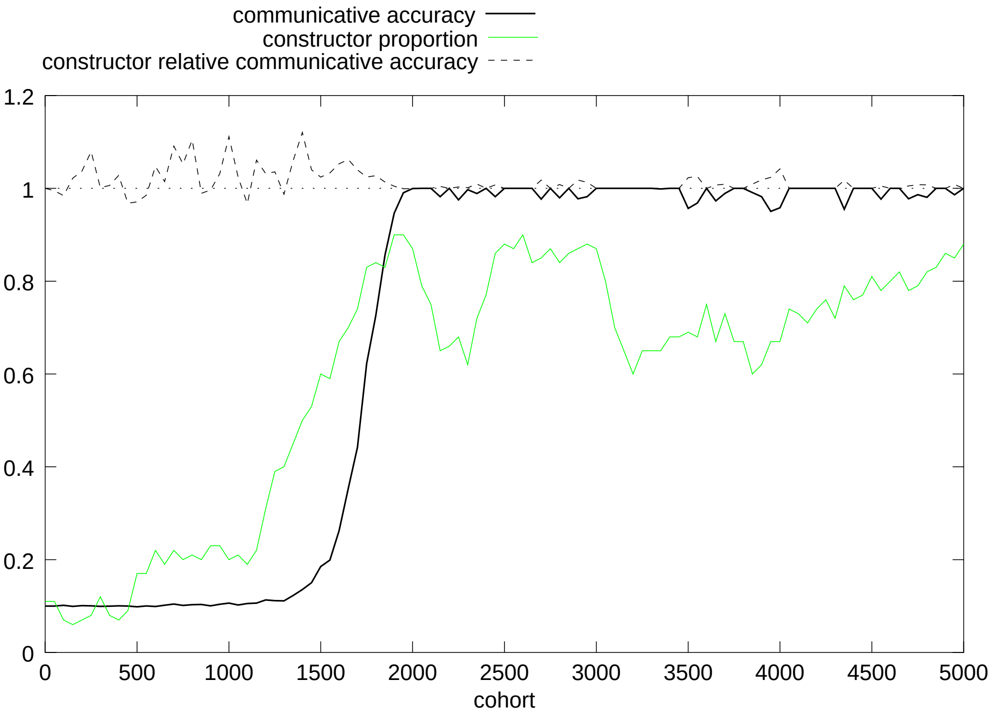

Figure 4.14: The relative communicative accuracy of agents with constructor genotypes. This value fluctuates around 1, although it is most obviously above 1 from 1300 to 1900 cohorts.

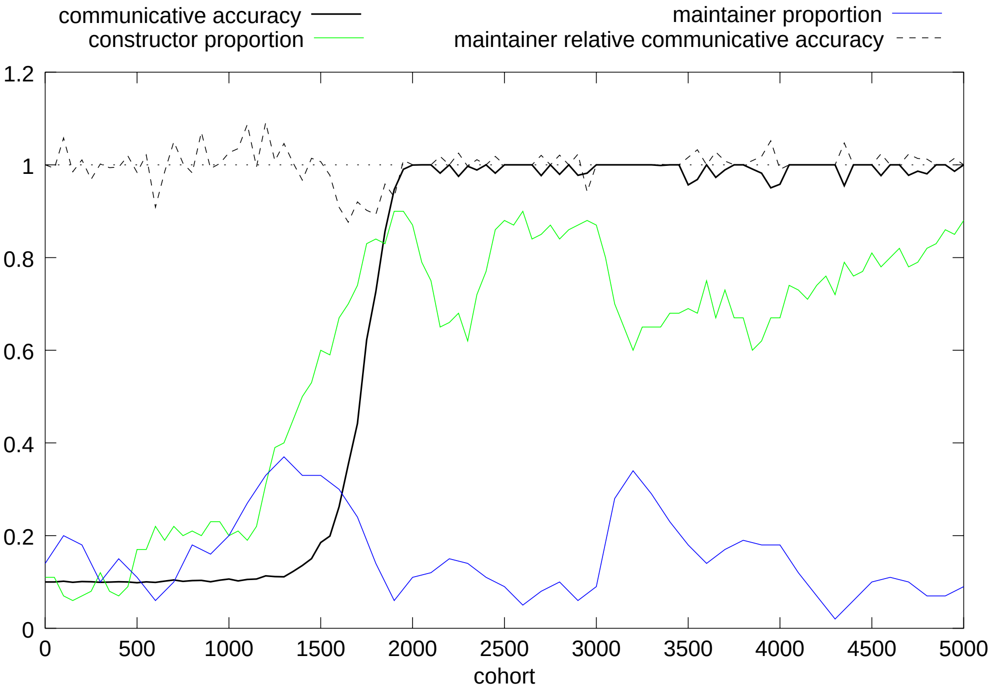

Figure 4.15: The relative communicative accuracy of agents with maintainer genotypes. This value fluctuates around 1, and is obviously below 1 from 1500 to 1900 cohorts, during which time maintainer numbers drop sharply.

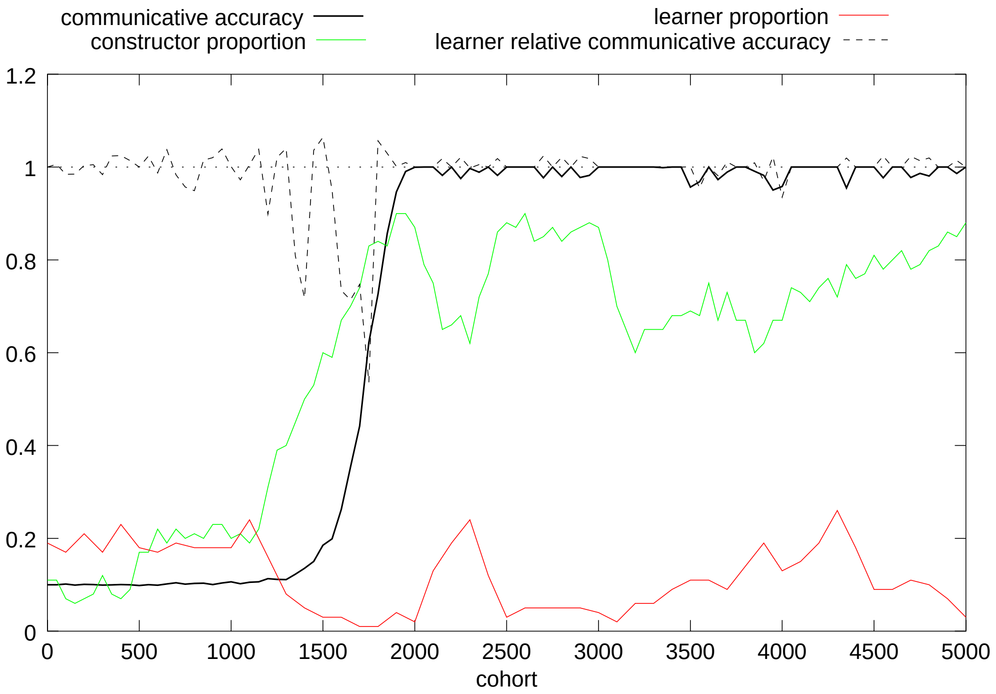

Figure 4.16: The relative communicative accuracy of agents with learner genotypes. This value fluctuates around 1, and drops most markedly below 1 from 1100 to 1800 cohorts, during which time learner numbers drop sharply.

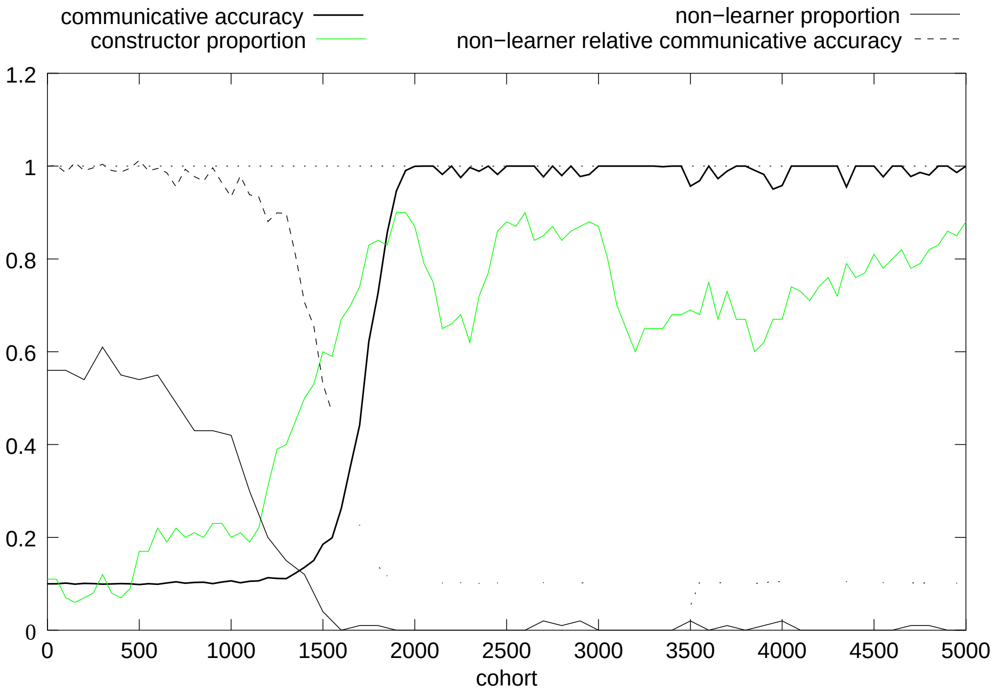

Figure 4.17: The relative communicative accuracy of agents with non-learner genotypes. This value initially fluctuates around 1, and appears to be below 1 from 600 cohorts onwards. By this time, drift has already delivered a significant number of constructor agents. The *rca* for non-learners drops markedly after about 1000 cohorts, at which point constructor numbers shoot up. After the optimal system is established, non-learner numbers and *rca* remain very low.

the initial period until drift establishes [+constructor] individuals in sufficient number. This is due to the fact that the fitness payoff in this model, unlike in B&R's model, is frequency-dependent — if no other individuals have the same cultural variant as you, you receive no payoff.

These simulation results emphasise a fairly basic logical point — if there is no established useful communication system present in the population, natural selection will not favour individuals who are predisposed to learn such a system. In the simulation results shown here, genetic drift is required to break the cycle of non-communication, at which point being biased to acquire an optimal communication system becomes advantageous. However, the point remains that, in a non-communicating population, being able to learn an optimal communication system confers no advantage.

4.5.5 *Varying the speed of convergence by varying e*

The speed of convergence of [+constructor] populations on optimal systems depends on e, the number of exposures learners receive to the observable behaviour of the population. In general, larger values of e lead to more rapid cultural convergence in the ILM. The results outlined above suggest that larger values of e will result in a higher proportion of runs of the EILM converging on optimal systems. For higher values of e there will be a shorter time-lag between emergence of genotypes encoding [+constructor] rules and an increase in communicative accuracy for individuals using such rules, therefore the populations will be less vulnerable to genetic drift during the abbreviated period of construction.

The experiments outlined in the previous section were repeated for all integer values of e in the range $[1, 30]$. The proportion of runs (out of 100) converging on optimal systems are plotted in Figure 4.18, along with the average speed of construction for those values of e in pure ILM runs. As can be seen from this Figure, there is a clear relationship between speed of construction in the ILM and probability of convergence on an optimal system in the EILM — as e increases, speed of construction decreases and probability of convergence in the EILM increases as a consequence. The correlation between time to convergence in the ILM and proportion of convergent runs in the EILM is statistically significant ($r = -0.56$, $p = 0.01$).

4.5.6 *Varying the speed of convergence by varying cultural population size*

Oliphant (1999) reports that population size impacts on the speed of cultural convergence, with larger populations taking longer to converge culturally on optimal systems. More

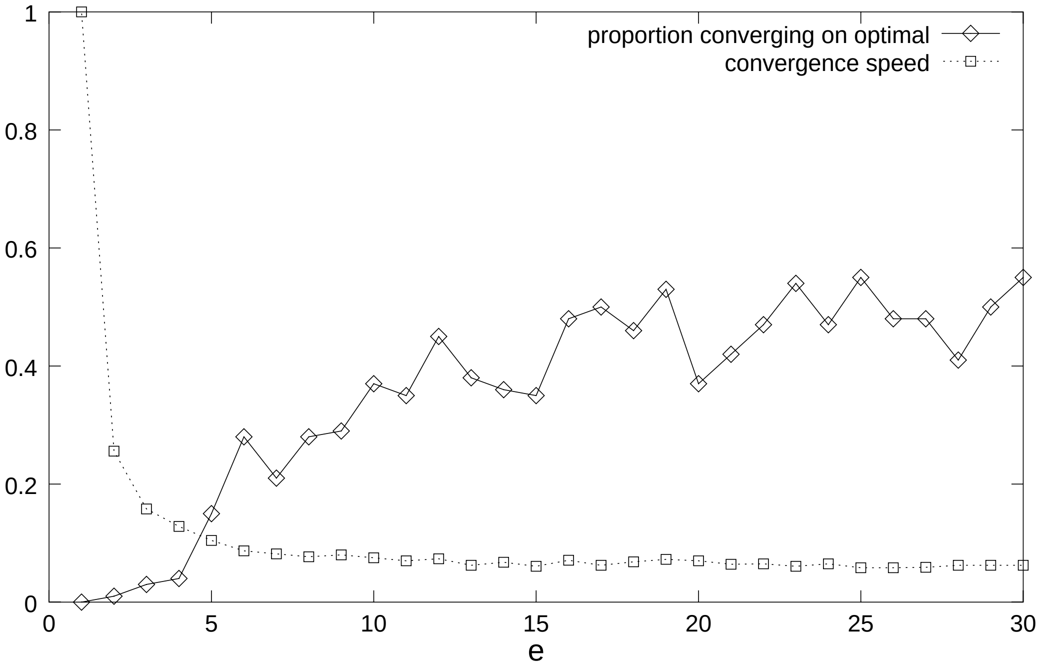

Figure 4.18: The relationship between e, convergence speed in the ILM and probability of convergence on an optimal system in the EILM. As e increases, the time taken by runs of an ILM to converge on a stable system decreases (speed of convergence is averaged over 10 runs and shown as a proportion of the time taken by the slowest converging runs). In the EILM, increased speed of convergence reduce the reliance on genetic drift, and consequently as e increases the proportion of EILM runs converging on an optimal communication system increases.

generally, we would expect cultural innovations to take longer to diffuse through larger populations, assuming that the rate of diffusion is invariant with respect to population size. Given the relationship between speed of convergence in the ILM and probability of converging on an optimal system in the EILM, we might expect smaller population sizes to lead to a higher probability of convergence in the EILM. However, this prediction is complicated by the genetic drift phase of the EILM described above. Genetic drift is more pronounced and its effects more dramatic in smaller populations, which would be a confounding factor in repeating the EILM experiments outlined above for smaller populations.

In order to decouple the effects of genetic population size and cultural population size, a model of spatial organisation was introduced to the EILM. The model of spatial organisation outlined in Oliphant (1997) is adopted for our purposes. The population of N agents is organised in a ring. This allows us to define a distance measure between individuals in the population. In the case where the population is unorganised spatially the probability of picking a particular individual is independent of their position in the population. However, in the spatial organisation case this is a function of their distance from a particular focal position on the ring.

There are four processes which involve picking individuals from the population:

1. picking individuals to participate in death tournaments.
2. picking individuals to participate in breeding tournaments.
3. picking individuals to communicate with.
4. picking individuals to observe and learn from.

Individuals are picked at random to participate in death tournaments, irrespective of their position in the population. Each of the remaining three processes is given a spatialisation parameter, d_b, d_c and d_l. The processes then proceed as follows:

Breeding Tournaments: Breeding produces a single individual to replace the single dead individual, who occupied slot i. The T individuals who will participate in the tournament are selected from the population probabilistically, according to their distance from position i. f is a normal distribution with a mean of 0 and standard deviation d_b. The probability of an individual with distance j from individual i being picked is equal to $f(j)$ — if d_b is low, only individuals close to i will be picked, whereas if d_b is high, individuals who are distant from i may be picked. The offspring produced by breeding occupies slot i.

Communication: Individual i is scored according to their average communicative accuracy when acting as both producer and receiver with two partners according to the measure $ca(P, R)$ given in Section 3.2 in Chapter 3. The partners are selected according to a normal distribution around individual i with standard deviation d_c.

Learning: Individual i receives e exposures to the population's communicative behaviour. During each of these e exposures the learner observes the complete set of meaning-signal pairs of a member of the population, where that member of the population is selected according to a normal distribution around individual i with standard deviation d_l.

For low values of d the population will be highly organised spatially, with individuals interacting only with individuals who are close by in terms of the distance measure. As d increases the population becomes less spatially organised, with the probability of interacting with distant individuals increasing. For $N = 100$ and $d \approx 50$ the population is effectively completely unorganised spatially with respect to the process of interest (breeding, communication or learning).

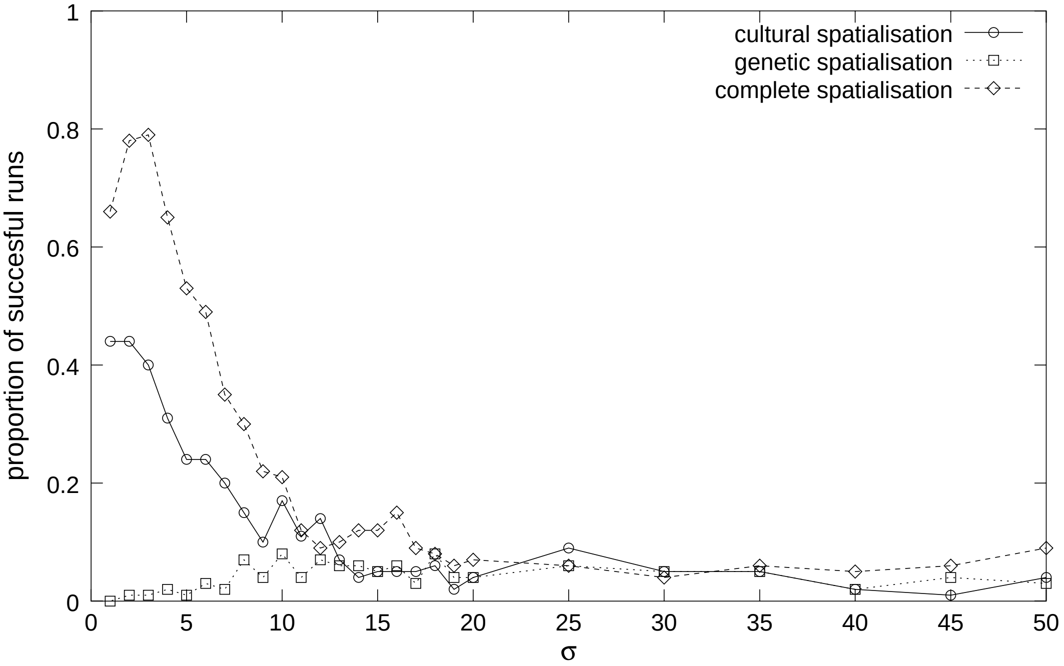

Figure 4.19: The impact of spatial organisation. σ gives the spatial organisation parameter, for those processes which are spatially organised. As σ increases, spatial organisation decreases. When cultural transmission is spatially organised, the proportion of EILM runs converging on an optimal system decreases, due to the increased time to cultural convergence. Conversely, when genetic transmission but not cultural transmission is spatially organised, the increase in σ results in a slight increase in the proportion of successful EILM runs, due to less fierce genetic drift. When genetic and cultural transmission are both spatially organised, a cultural kin-selection effect is observed, which tails off as σ increases.

I will consider three combinations of parameters:

Cultural spatialisation: $d_l = d_c = \sigma$, $\sigma \in [1, 50]$,$d_b = 50$
Genetic spatialisation: $d_b = d_c = \sigma$, $\sigma \in [1, 50]$,$d_l = 50$
Complete spatialisation: $d_l = d_b = d_c = \sigma$, $\sigma \in [1, 50]$

Figure 4.19 plots the proportion (out of 100) of populations which converge on optimal systems[14], for the three parameter conditions for various values of σ (as with the earlier ILMs, $N = 100$, $e = 3$).

Three patterns are clear here. Firstly, the number of convergent runs in the genetic spatialisation condition is very low, and is at its lowest when genetic spatialisation is at its most extreme. This is what we would expect — genetic drift in the sub-populations will

---

[14] Runs were classified as optimal if communicative accuracy reached and subsequently remained above 0.85. This value was lower than the previous classification values to allow for the reduction in communicative accuracy resulting from boundaries between communication systems.

be more pronounced, and cultural convergence will still take a long time. Secondly, the proportion of convergent runs in the cultural spatialisation condition decreases as σ increases — as σ increases, effective population size increases, resulting in slower cultural convergence and hence greater vulnerability to genetic drift.

Finally, the runs for the complete spatialisation case exhibit a similar pattern to the cultural spatialisation case, but with a higher overall proportion of convergence and a hump for very low σ. The hump is due to extreme genetic drift when $d_b = 1$. The overall higher level of convergence suggests a positive interaction between genetic and cultural spatialisation. Neighbouring individuals will typically have similar genotypes, will have learned from similar observable behaviour and will have similar communication systems. If individuals are of the [$+$constructor] classification, this is likely to result in very quick convergence on shared meaning-signal mappings. Sharing meaning-signal mappings leads to increased communicative accuracy, increased probability of breeding, and increased chances of adding yet more constructor agents into the constructor sub-population, who will then learn from similarly biased individuals, and subsequently communicate successfully with them.

Figures 4.20, 4.21 and 4.22 give sample convergent runs in more detail. As can be seen from Figure 4.20, tight cultural spatial organisation leads to the rapid formation of dialects within the population, which yield above-average communicative accuracy. Figure 4.21 shows that these spatially-organised dialects do not form when the population is not spatially organised for cultural transmission. A similar relationship between cultural spatial organisation and dialect formation is reported by Livingstone & Fyfe (1999). Finally, Figure 4.22 shows that, when both genetic and cultural transmission are spatially organised, the emergent, useful dialects are associated with constructor genotypes, which then spread to dominate the population.

4.5.7 Summary

The simulation results outlined in this Section show that natural selection cannot reliably identify weight-update rules which lead to the cultural evolution of optimal communication. This is due to the time-lag between the emergence of such weight-update rules and any communicative payoff to individuals possessing them. In simulation runs which do converge on optimal communication systems, genetic drift provides appropriate weight-update rules in sufficient numbers for the cultural construction process to get under way, at which point constructor weight-update rules are identified and selected for during genetic transmission. Increasing the number of observations each learner makes (increasing

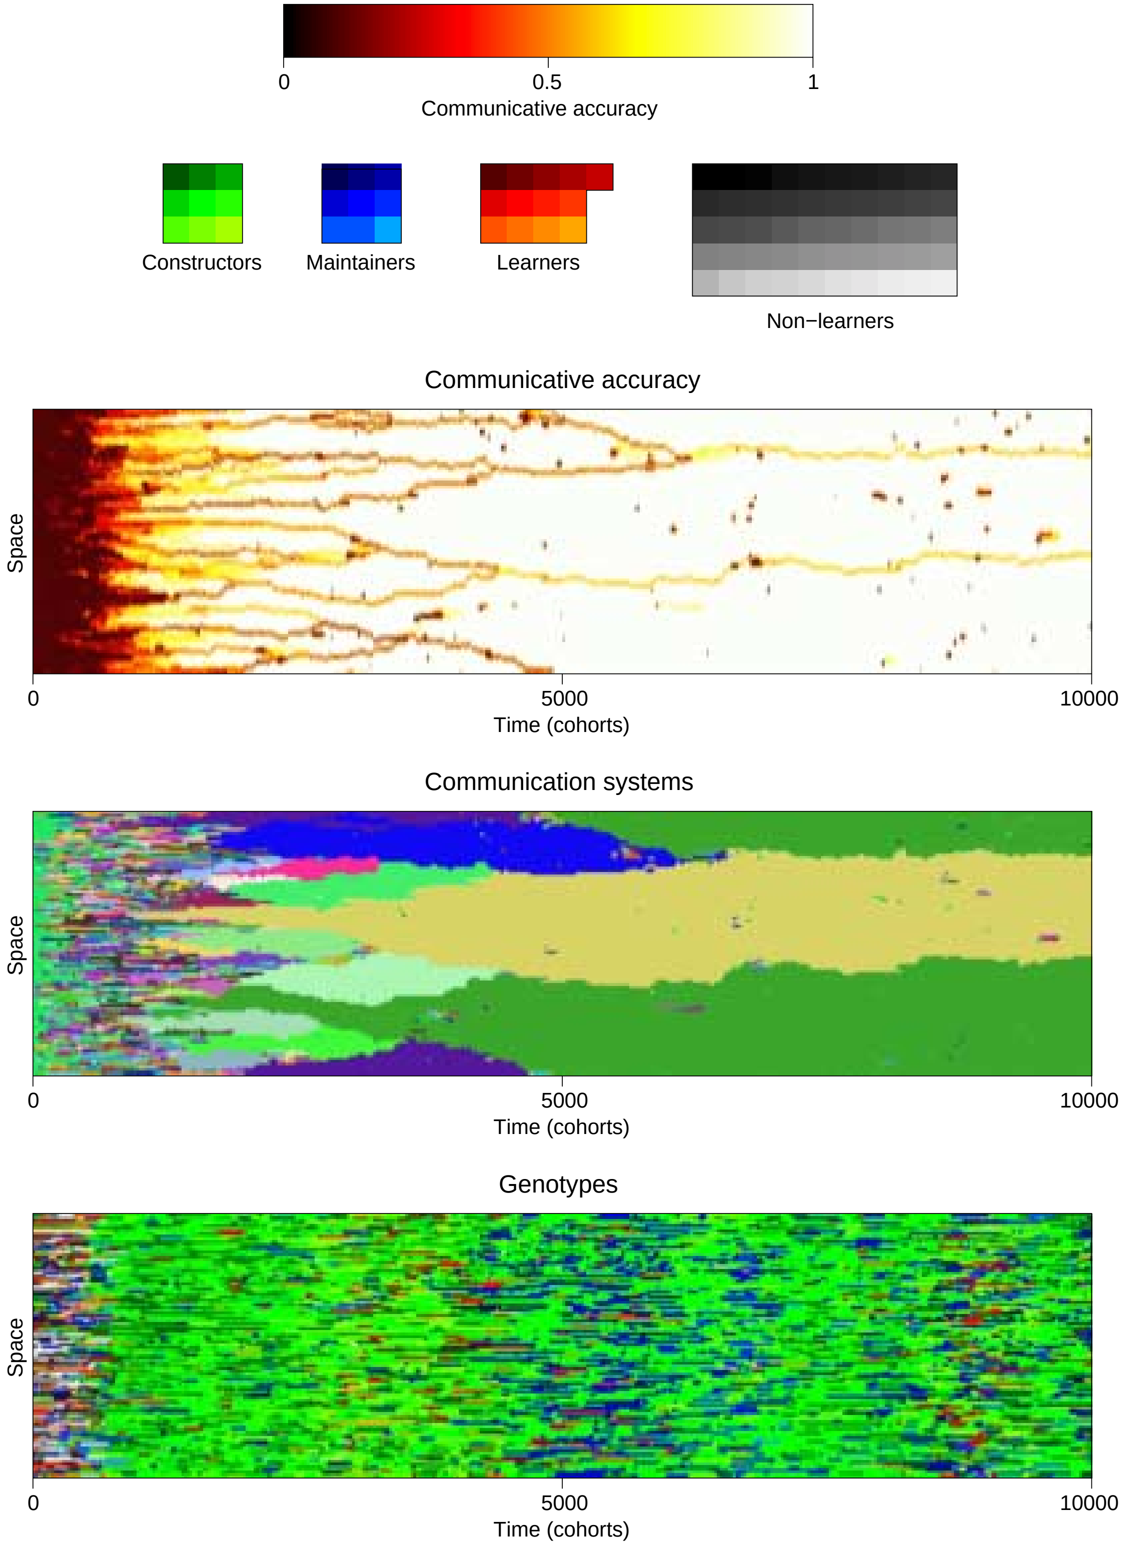

Figure 4.20: Cultural spatialisation only ($d_l = d_c = 1$, $d_b = 50$). Time-space diagrams for communicative accuracy (top), communication systems (middle) and genotypes (bottom). Each vertical strip of each diagram represents the complete population at a particular point in time (in cohorts). In the communicative accuracy diagram, different colours correspond to individuals with different levels of communicative accuracy (see key). In the communication systems diagram, each distinct $p(m)$ is assigned a distinct (random) colour. In the genotypes diagram, different colours correspond to genotypes encoding different weight-update rules (see key).

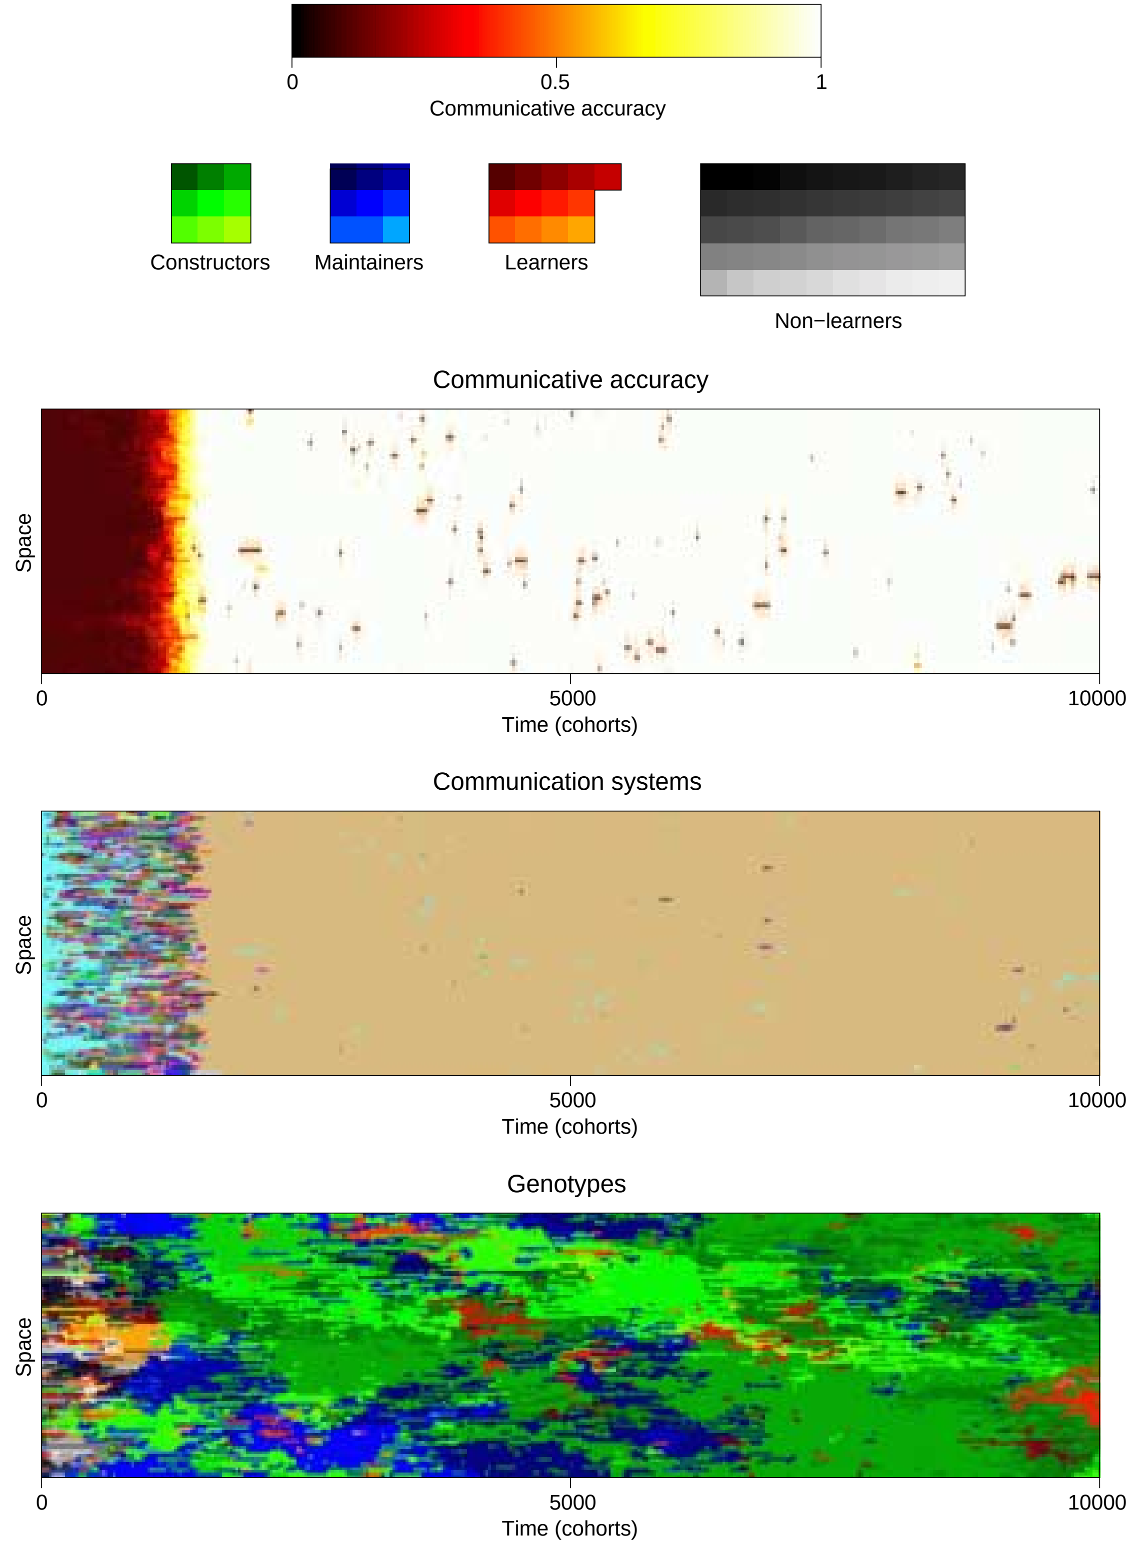

Figure 4.21: Genetic spatialisation only ($d_b = d_c = 3$, $d_l = 50$). While genetic subpopulations exist, diversity of communication systems is rapidly lost.

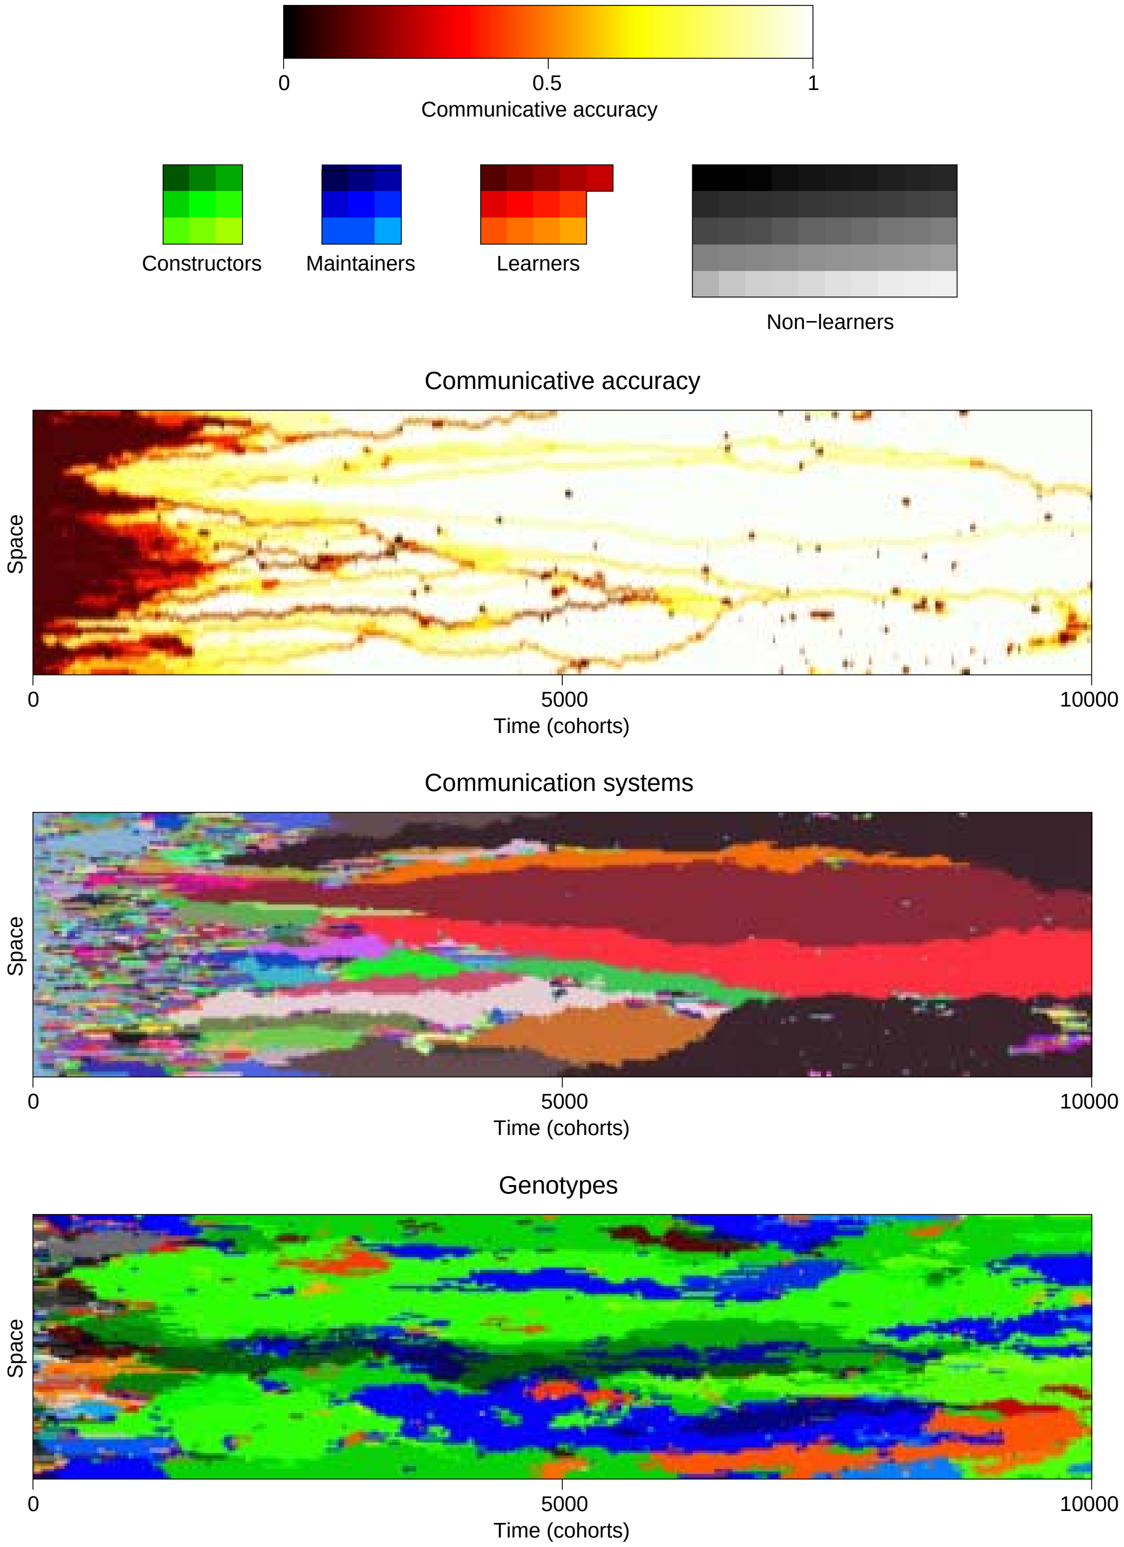

Figure 4.22: Combined spatialisation ($d_l = d_b = d_c = 1$). Both cultural and genetic diversity is maintained. Note that the early regions of increased communicative accuracy coincide with regions dominated by constructor genotypes. Later on, dialects are less associated with a particular genetic makeup. Note also the small subpopulation of learner agents, emerging at around 6000 cohorts and disappearing at 10000 cohorts. This build-up of [+learner, −maintainer] individuals leads to a drop in communicative accuracy, visible in the top diagram, and consequently the elimination of the learner genotypes.

e) or reducing the cultural population size (by introducing cultural spatialisation) speeds up the cultural convergence process, resulting in more frequent convergence on good genes and optimal communication systems. In successful runs, a kind of cultural niche construction (Odling-Smee 1988; Laland *et al.* 2000) takes place — drift provides constructor genotypes in sufficient numbers for sufficient time to construct a cultural niche, which constructor agents are then selected to occupy. However, the logical point remains that in a non-communicating population there can be no communicative advantage in being biased to acquire an optimal communication system.

4.5.8 Discussion

In Chapter 3 I argued that humans are biased in favour of acquiring one-to-one mappings between meanings and vocabulary items, and that this learning bias could result in the emergence of effective communication in human populations through purely cultural processes. The fairly natural assumption would be that this learning bias in human infants has evolved through natural selection in favour of communication. However, the results outlined in this Chapter cast doubt on this assumption — in the simulation model, such biases are unlikely to evolve, in spite of strong selection pressure in favour of successful communication.

What can these results tell us about the evolution of vocabulary acquisition biases in humans? There are two possible interpretations. I will take as the starting-point for both interpretations the assumption that pre-linguistic human populations resembled current-day primate populations in terms of communicative behaviour — they possessed some fairly limited set of stimulus-bound, innate signalling behaviours, possibly augmented by a few idiosyncratic, ontogenetically-ritualised signals shared between pairs of individuals. It should of course be remembered that non-human primates have undergone a significant period of evolution since we last shared a common ancestor, and do not constitute some kind of living fossil. However, modern humans still possess a small range of universal, innate signals (facial expressions for disgust, anger and so on) and it is reasonable to suggest that the basis of this common characteristic existed in the last common ancestor of chimpanzees and humans, and therefore in the first non-common ancestor on our side of the family tree. I will further assume that any emergent, culturally-transmitted symbolic communication system did not replace this innate, non-linguistic signalling system — the persistence of human facial expressions and the like shows that it did not.

Under the first, positive interpretation, the simulation results outlined in this Chapter can be seen as highlighting the conditions under which the learning bias in humans must have evolved. We can then draw parallels between the simulation parameters which lead most reliably to the evolution of appropriate learning biases and the hypothesised ecological setting of human evolution. Alternatively, we could emphasise the negative aspect of the simulation results, and mount an argument that the human vocabulary acquisition biases, apparently so well-designed for communication, are better understood as an exapted trait. These two contrasting interpretations are considered in the following two Sections.

4.5.8.1 The positive interpretation

The weakest positive interpretation of the simulation results is to conclude that the evolution of the human vocabulary acquisition bias was something of an evolutionary fluke — the simulation results indicate that such learning biases are unlikely to evolve but, given a large dose of luck, can do so. Under this interpretation, the uniqueness of human language as a learned symbolic communication system (at least among the primates) can be seen as a consequence of the fortuitous emergence of a significant number of individuals possessing a language-specific learning bias with the correct properties. After the cultural construction process got underway, these individuals reproduced with disproportionate success, due to the payoff of communication, and the genes encoding the bias became fixed in the population.

This kind of account is rather unsatisfactory, relying as it does on a double saltation. Firstly, the appropriate learning bias must emerge de novo in a population whose communication system was previously innate. Secondly, this learning bias must be maintained in a significant proportion of the population for a significant period of time, despite a lack of fitness advantage for those individuals possessing it. Of course, assuming that the appropriate bias is costless relative to inappropriate biases, genetic drift may preserve it, or even increase its numbers. However, given that the genes encoding the new bias will initially occur in small numbers, drift is more likely to result in their disappearance than their proliferation.

As a second possible interpretation, the simulation results could be taken as indicating that the emergent learning bias in human populations immediately provided a fitness advantage to those individuals who possessed it — in this case, the time-lag problem would not occur. It has been suggested (for example, in Chomsky (2002)) that an important function of language is individual-internal, allowing the formulation of plans and so on using “inner speech”. If this was the case, then it could be argued that the initial individuals possessing the human vocabulary acquisition bias could use their acquired, sub-optimal but non-random, communication system internally, thereby immediately gaining a fitness advantage over individuals who could not communicate with themselves.

I have never found this line of argument particularly convincing with respect to full-blown language, and it is even less convincing when stretched to cover unstructured communication. Firstly, it is not clear to what extent inner speech is a widespread phenomenon, or even a real phenomenon. Individuals who claim to think in speech could be imposing some kind of post-hoc rationalisation on their thought processes. Even if some individuals do think in speech, it is not known if this is common, or a matter of personal cognitive style. Even if all individuals do think in speech sometimes, it is also not clear that this mode of thought is applied to a wide range of tasks — while we might be happy to agree that we "think in language" when performing tasks like rehearsing a talk or writing a thesis, it is less clear to what extent we think in language when performing less wordy tasks like painting a shed, going to the shops, or hunting a wildebeest. Ignoring the empirical evidence (or lack thereof) regarding the degree to which we actually do think in language, if this inner speech provided the main impetus for the evolution of the linguistic capacity then we might expect language to be of a rather different form — see Hurford's points in Chapter 1, Section 1.4.1, on the apparent adaptation of language for the externalisation of propositional structures. With respect to thinking in an unstructured communication system, it is even difficult to see how thinking, for example, wildebeest$'$, signalling "wildebeest", and re-arriving at wildebeest$'$ offers any improvement in clarity of thought.

An alternative, more appealing possibility is that some aspect of the emergent learning bias in pre-linguistic hominids or the social structure of pre-linguistic hominids resulted in an immediate, or at least very rapid, fitness advantage to individuals possessing the appropriate bias. The simulation results in Sections 4.5.5 and 4.5.6 show that increased speed of cultural convergence, either by increasing the number of exposures learners receive or by reducing cultural population size, leads to more reliable emergence of learning biases supporting optimal communication. If a plausible argument could be presented that one or both of these factors were at play in early stages of the emergence of human language-learning biases then we might be more optimistic that these learning biases evolved specifically for communication in the hominid line.

Human infants have an unusually long period of immaturity, a consequence of the reduced size of the pelvic opening in humans, which requires that babies be born at a relatively immature stage, before the head is too big to fit through the birth canal (Martin 1992). This extended period of immaturity and plasticity could be very roughly equated with large e in the simulation models, which speeds up cultural convergence and leads to more reliable evolution of the one-to-one bias. Alternatively, and perhaps more plausibly, we could point to the speed of cultural convergence in creolization situations, arguably a consequence of a strong human learning bias. Bickerton argues that creolization takes place in a single generation (Bickerton (1990), but see, for example, (McWhorter 1997) for a contrasting view). If we accept for a moment that the emergent learning bias in the hominid line led to significant convergence on a shared communication system within one generation, or even a few generations, then natural selection would be more likely to quickly identify and select for this bias. This does, however, require that the bias in its initial form was extremely strong, ruling out initial emergence of a weak bias which was gradually selected for in stronger and stronger forms.

Reduced cultural population size leads to rapid cultural convergence in the simulation model. We could argue that this is characteristic of the population structure of pre-linguistic hominids, but not other primates, facilitating the evolution of a unique learning bias in the hominid line. Arguments based on group size have been made before in relation to the evolution of language in humans (Dunbar 1996). However, two main problems exist with this approach. Firstly, Dunbar's account is centred on the assumption that group size increased down the hominid line. This produces exactly the wrong prediction in terms of our model — larger group size should lead to slower cultural convergence and therefore a lower likelihood of evolving the correct bias. It could be argued that, while overall group size increased in hominids, groups became increasingly internally structured, which would produce the correct prediction. However, this brings us to the second problem with group-size arguments. In a recent review of the archaeological evidence underpinning theories of the evolution of language, Buckley & Steele (2002) conclude that there is no good evidence that group size changed during the Pleistocene period (from 1.6 million to 10,000 years ago), and that there is little to no direct evidence at all from earlier times. Making inferences from the fairly sparse archaeological evidence on group structure, a more subtle phenomenon, is likely to be even more fraught.

As a final, (semi-)positive interpretation, it could be concluded from these simulation results that a bias for the acquisition of vocabulary did not evolve under natural selection for communication in hominid populations, but a bias for the acquisition of *structured* communication did. This possibility will be investigated in Chapter 6. We should, however, remain sceptical — we would expect the logical point that there is no advantage to learning if there is nothing interesting to learn to pertain regardless of the degree of structure of the object of learning.

4.5.8.2 The negative interpretation

The contrasting position would be to interpret the results in this Chapter as indicating that the human capacity for vocabulary acquisition, and perhaps by extension language acquisition, did not evolve under natural selection for communication. This would mean that this capacity was initially not language-specific, although subsequent modifications may have refined the scope of the application of the bias.

How then should this learning bias in humans be explained? A first possibility is that this piece of cognitive apparatus is an exapted trait — the learning bias evolved under reliable selection pressure for some other task, and was later reappropriated for the acquisition of communication systems. A slight modification of this theory would be that it evolved initially for some general function, all the while being incidentally applied to the acquisition of vocabulary. The cultural construction of effective communication systems got underway, at which point the bias was selected directly for the communicative payoff it offered.

The main problem with these accounts is that it is not clear what the original function of the learning bias was — what other learning tasks require a one-to-one bias? As discussed in Chapter 1, some theories of the evolution language have suggested that the capacity for tool use and language are intertwined. It could be argued that the explosion in specialisation in tool function which occurred towards the end of the Paleolithic period (approximately 40,000 years ago) may have reflected a new bias in favour of one-to-one mappings between functions (meanings?) and forms (signals?). However, this is wildly speculative. An alternative, more plausible possibility is that the one-to-one bias was initially a side-effect of the particular method of learning which was selected in the population, for general or specific unknown purposes — the one-to-one bias was a spandrel of the selected learning mechanism.

An interesting variation on this possibility is that the one-to-one bias is a consequence of the evolution of a sophisticated theory of mind, an understanding of other individuals as intentional agents. Tomasello (e.g. Tomasello (1997) and Tomasello (1999)) argues that the human capacity for this type of mental gymnastics is uniquely sophisticated among the primates. Tomasello argues that this cognitive capacity could have evolved under an array of selection pressures, including selection for communication, cultural learning in general, cooperation and tool use.

How would this understanding of others as intentional agents lead to a learning bias in favour of the acquisition of one-to-one mappings between meanings and signals? Clark (1990) and Bloom & Markson (1998) present the vocabulary acquisition biases of humans in essentially pragmatic terms — children understand the nature of communicative acts, which is basically to draw their attention to particular items, and infer from that understanding that ambiguity (introduced by synonymy or homonymy) is undesirable. I have tended to shy away from this interpretation here, preferring to discuss learning bias as a consequence of weight-changing procedures. However, the two interpretations are roughly equivalent for my purposes — a one-to-one bias should either be imposed on the learning act itself (as I have it), or on the process by which the learner deduces the object which is being referred to (as Bloom & Markson (1998) have it). My account has the advantage of extending, as will be discussed in Chapter 5, to the evolution of structured communication, whereas Bloom & Markson see their account as being restricted to lexical learning. Bloom & Markson's (1998) account has the advantage that two processes which are distinct under my model (identifying the meaning of an utterance and learning with a one-to-one bias) are shown to be aspects of a single cognitive process.

4.6 Summary of the Chapter

At the beginning of this Chapter I briefly outlined approaches to modelling genetic transmission, and discussed some of the issues involved in applying such models to the investigation of the evolution of innate signalling systems. My primary concern, however, was to introduce the dual transmission model, which unifies models of biological and cultural evolution. The dual transmission model forms the basis of the Evolutionary Iterated Learning Model.

I then outlined two EILMs, based on extensions of the feedforward network and associative network ILMs discussed in Chapter 3. These EILMs allow us to investigate the interactions between the biological evolution of learning bias and the cultural evolution of communication systems acquired using this biased learning apparatus.

In the feedforward network EILM, there was found to be little significant interaction — the learning bias associated with the imitator or obverter network architectures proved to be the decisive factor in determining the behaviour of populations in the EILM. In the associative network EILM, the picture was more complex. One-to-one learning biases can evolve under natural selection pressure for communicative success. However, the evolution of such biases was dependent on an initial period of genetic drift, due to the lack of an immediate benefit to individuals with the appropriate learning bias. This dependence on genetic drift can be reduced by speeding up cultural convergence within populations, either by increasing the amount of learning each individual does or by introducing spatial organisation and effectively dividing a single population up into several smaller subpopulations. However, the point remains that there is no immediate benefit to being biased to acquire an optimal communication system in a non-communicating population.

This result has implications for our understanding of the evolution of vocabulary-learning biases in human populations. We could take the simulation results as indicative of the conditions under which such a bias evolved in hominid populations. The alternative, possibly stronger view, is to conclude that these learning biases did not evolve directly for the purpose of communication — human vocabulary learning biases are a spandrel, or an exapted trait.