CHAPTER 6

The evolution of compositionality in populations

In the previous Chapter the cultural evolution of compositionality was investigated in the context of a single individual learning their communication system from a single cultural parent, and transmitting their system to another single individual. In such a context the notion of communication has no place, as solitary individuals have no-one to communicate with. It is also meaningless to study genetic transmission in such a context, as a population consisting of a single individual will be genetically homogeneous by definition.

In this Chapter I will investigate the transmission of structured communication systems in the context of *populations* of individuals. Do the findings outlined in the previous Chapter hold in the context of cultural transmission within populations? What consequences does this have for communication within such populations? And how does the dual transmission of communication systems by cultural transmission and weight-update rules by genetic transmission impact on the population's communicative behaviour?

A review of population-level ILMs and EILMs which deal with the evolution of structured communication is carried out in Section 6.1. This review suggests that the ILM and EILM approaches can be extended to a consideration of linguistic evolutions in populations. However, they also highlight the fact that studies of the evolution of learning bias have typically been restricted to parameter-setting models of learning, or variants thereof. In Section 6.2 I discuss methods of measuring communicative accuracy among populations of individuals using the type of languages discussed in the previous Chapter. In Section 6.3 I go on to describe an extension of the ILM from the previous Chapter to non-trivial population sizes. Finally, in Section 6.4, I expand this model to a full Evolutionary Iterated Learning simulation, and outline results pertaining to the evolution of learning biases which support communicatively optimal, compositional language.

6.1 Models of the evolution of linguistic structure in populations

Iterated Learning Models of the cultural evolution of structured communication in populations do exist, and are discussed in Section 6.1.1. The population approach also leads fairly naturally to the use of Evolutionary Iterated Learning Models, and examples of these models are discussed in Section 6.1.2.

6.1.1 Cultural evolution in populations

I have discussed three models which demonstrate that the repeated expression and induction of linguistic form within a population can lead to the emergence of structured systems of meaning-signal mappings. Two models by John Batali (Batali (1998), described in Section 2.3.3.4 of Chapter 2, and Batali (2002), discussed in Section 2.3.6.2 of Chapter 2) demonstrate that morphological and syntactic structure can emerge in populations. However, as discussed earlier, the Negotiation Model framework makes it difficult to isolate the relative importance of learning bias and transmission bottleneck in shaping linguistic behaviour in these populations. Hurford (2000), covered in Section 5.1 of Chapter 5, demonstrates that recursively compositional language can emerge in small populations (Hurford uses a population size of 5) in a gradual ILM population model. Similarly, Kirby (2000) demonstrates that compositional language can emerge in small populations through purely cultural processes. Kirby's (2000) model is an earlier version of the model described in Kirby (2002). A gradual turnover ILM is used, with a simple, non-embedding semantics and a stochastic grammar inducer.

Parameter-setting models of language acquisition have also been used in population-level ILMs. Niyogi & Berwick (1997), using a mathematical model, consider the spread of linguistic variants in populations through Iterated Learning. In Niyogi & Berwick's model, there are two competing linguistic variants, L_1 and L_2, which differ with respect to the setting of a particular parameter[1]. Learners sample the languages of the adult population and decide, using a parameter-setting procedure known as the Trigger Learning Algorithm (Gibson & Wexler 1994), on which variant to acquire. Niyogi & Berwick demonstrate that, if one language, say L_1, produces triggers which are consistent with both L_1 and L_2, while L_2 produces triggers which are only consistent with L_2, then L_2 will come to dominate the population.

[1]Niyogi & Berwick's model need not be interpreted as a parameter-setting model — we could take the two linguistic variants to represent linguistic systems which differ arbitrarily from one another. However, the single parameter interpretation is the most natural one, given their use of the Trigger Learning Algorithm.

Briscoe (2000a) (discussed in Section 2.3.5.2, Chapter 2) and Kirby (1999) (described in Section 2.3.3.2, Chapter 2) present models which are similar to that of Niyogi & Berwick, which demonstrate the convergence of populations on shared parameter settings. This convergence is driven by a frequency-dependent bias in Briscoe's model, whereas a direct bias in favour of parsability drives populations in Kirby's model to a parameter setting which is yields optimally parsable utterances.

6.1.2 Gene-culture coevolution in populations

Parameter-setting approaches have also been used to model the co-evolution of languages and parameter-setting language acquisition devices. Kirby & Hurford (1997) describe an extension to Turkel's (2002) model of the evolution of co-ordination. As in Turkel's model, each individual's genotype consists of a string of 1s and 0s (representing inviolable principles) and ?s (representing settable parameters). An individual's mature phenotype is a string of 1s and 0s, with all genotype ?s being set to either 1 or 0. Unlike in Turkel's model, Kirby & Hurford use a generational EILM, with the setting of parameters in an individual's mature phenotype being determined by cultural transmission. Immature individuals receive a number of trigger utterances from mature individuals in the previous generation. Triggers specify the setting of a single parameter (as either 1 or 0), and mature individuals produce triggers consistent with their own grammar. Learners then set the values of ?s in their phenotype according to observed triggers and a learning procedure based on the Trigger Learning Algorithm. Mature individuals breed according to communicative success. Communicative success between two individuals is dependent on the number of matching settings they share in their mature phenotypes, but also on the 'parsability' of the utterances they produce — an arbitrary subset of the set of possible mature phenotypes were considered to produce more parsable utterances.

Kirby & Hurford report that, under these conditions, maximally functional (parsable) grammars do *not* emerge — the simulated populations converge on mature phenotypes which do not produce maximally parsable utterances. Not only does the population converge culturally on suboptimal grammars, but they nativize those suboptimal grammars — the Baldwin effect leads to the emergence of disfunctional, inviolable principles which prevent learners in the population from acquiring an optimal system. Kirby & Hurford attribute this to the overriding pressure for learners to learn the language of their community, regardless of whether it is optimally parsable or not.

In a second set of experiments, Kirby & Hurford introduce selection for parsability on cultural transmission. With a certain small probability, learners preferentially retain parameter settings which yield higher parsability over settings which yield lower parsability. Learners are therefore directly biased in favour of acquiring optimal parameter settings. Under this revised setup, the simulated populations converge on optimal grammars, and nativize those optimal grammars via the Baldwin effect.

Briscoe (2000b) presents an extension of his Iterated Learning Model (Briscoe (2000a), discussed above), which also demonstrates the role of the Baldwin effect in the nativization of linguistic structure, with languages which minimise working memory load and therefore improve parsability being preferentially nativized. His model shows that Kirby & Hurford's (1997) second result still stands under more realistic assumptions about grammars, parsability and population dynamics — more parsable and learnable languages emerge culturally, and the LADs of learners evolve a default setting which matches the dominant language in the population.

A paper by Martin Nowak and colleagues is of particular relevance to this Chapter, and merits discussion in detail. Nowak *et al.* (2000) compare the fitnesses of individuals pursuing two possible learning strategies — holistic learners and compositional learners. In their model, events consist of an action and an object acted upon. Nowak *et al.* vary the number of possible objects and actions, and also the possible combinations of objects with actions — for example, events involving action 1 and object 1 may occur with a certain frequency, whereas events involving action 1 and object 2 may never occur. Nowak *et al.*'s model of events is therefore equivalent to (and formed the basis for) my model of meaning spaces and environments — their event space corresponds to a two-dimensional meaning space, with V for each dimension given by the number of possible objects and actions, and their *event rate matrix* (which specifies which action-object combinations may occur) is equivalent to an environment in my model.

Nowak *et al.* assume there are two possible types of learners — holistic learners, who attempt to learn a single word for each event, and compositional learners, who learn separate words for actions and objects and then combine those words to form descriptions of events. Nowak *et al.* calculate the equilibrium frequency of words in populations consisting solely of holistic or compositional learners, calculate the levels of communicative accuracy associated with these word frequencies and compare these values across populations.

In populations of holistic learners, the frequency of individuals who use a word W_{ij} to refer to an event E_{ij} is given by $x\left(W_{ij}\right)$. Nowak *et al.* assume a generational ILM where each individual receives b exposures to the linguistic behaviour of the previous generation. During each of these exposures, a holistic learner successfully learns a word with probability q. $\phi(E_{ij})$ gives the frequency of occurrence of event E_{ij}. After a single generation, the frequency of word W_{ij} to refer to event E_{ij} is given by:

$$x'(W_{ij}) = R(W_{ij}) \cdot x(W_{ij}) \cdot (1 - x(W_{ij})) - x(W_{ij})$$

where $x'(W_{ij})$ is the proportion of individuals who know the word after transmission, $x(W_{ij})$ is the frequency of individuals who know the word prior to transmission and $R(W_{ij})$ is the reproductive rate of the word, and is given by:

$$R(W_{ij}) = bq\phi(E_{ij})$$

In other words, the reproductive rate of word W_{ij} is the product of the number of exposures each individual receives (b), the probability of learning that word after a single exposure (q) and the probability that you hear someone talking about the event E_{ij} ($\phi(E_{ij})$). In the population equation, the change in the proportion of people who know the word depends on this reproductive rate and the cultural variance in the population ($x(W_{ij}) \cdot (1 - x(W_{ij}))$). The term $\ldots - x(W_{ij})$ simply keeps the population size constant.

Given this equation, Nowak *et al.* go on to calculate the equilibrium frequency of individuals who know word W_{ij} for event E_{ij} — the frequency which, possibly infinitely many, iterated learning events will lead to, assuming that the word has a reproductive rate of greater than 1. This quantity, $x^*(W_{ij})$, is given by:

$$x^*(W_{ij}) = 1 - \frac{1}{R(W_{ij})}$$

In other words, the equilibrium frequency of W_{ij} will depend on its reproductive rate. If the reproductive rate of a word is very high then virtually everyone in the population will know the word. As all events which actually occur are assumed to occur with equal frequency, the key factors in determining $R(W_{ij})$ are number of exposures b and learnability q.

The communicative accuracy, and therefore fitness, of a population of such holistic learners at equilibrium is simply the probability that any two individuals will know the word for an event, summed over all possible events:

$$F_{holistic} = \sum_{i,j} \gamma_{ij} \cdot \left[x^* \left(W_{ij}\right)\right]^2$$

where γ_{ij} is 1 if event E_{ij} is allowed by the event rate matrix (environment) and 0 otherwise. Note two things about this measurement. Firstly, it is implicit in this measure that there is only ever one word for each event in the population — the measure is not summed over all possible words for each event. Secondly, it is also implicit that each event is associated with a unique word — words for one event are never confused with words for other events.

Nowak *et al.* go through a similar processes for calculating word frequencies in populations of compositional learners. Compositional learners learn separate words for objects O_i and actions A_j of events E_{ij}. The word associated with object O_i is N_i (a noun), while the word associated with action A_j is V_j (a verb). The assumption is made that learners spend half their time learning nouns and half their time learning verbs. The reproductive rates of N_i and V_j are given by:

$$R\left(N_i\right) = \left(b/2\right) q_c \phi \left(E_i\right)$$

$$R\left(V_j\right) = \left(b/2\right) q_c \phi \left(E_j\right)$$

where q_c is the probability of learning a noun or verb after a single exposure (taken to be lower than q), $\phi\left(E_i\right)$ gives the frequency of events involving object O_i and $\phi\left(E_j\right)$ gives the frequency of events involving action A_j. These equations are clearly variants of the holistic learner equations. The equilibrium frequency of individuals who know both N_i and V_j is given by:

$$x^*\left(N_i V_j\right) = \frac{\left(1 - 1/R\left(N_i\right)\right) \cdot \left(1 - 1/R\left(V_j\right)\right)}{1 - 1/\left(R\left(N_i\right) + R\left(V_j\right)\right)}$$

This leads to a communicative accuracy at equilibrium of:

$$F_{compositional} = \sum_{i,j} \gamma_{ij} \cdot \left[x^*\left(N_i V_j\right)\right]^2$$

In other words, as for the holistic learners, communicative accuracy depends on the probability of two individuals being able to make a signal for each event, multiplied by the probability of that event occurring.

Based on these equations, Nowak *et al.* go on to identify the conditions under which compositional learners will be preferred to holistic learners — the circumstances under which $F_{compositional} > F_{holistic}$. To do this they make several more assumptions. Firstly, they assume that there are n objects and n actions — there are as many objects as actions. This yields n^2 possible events. Suppose some fraction p of these events occur (for example, if $p = 0.5$ then only half of all possible combinations of objects and actions actually occur in the environment), and further assume that these events are distributed randomly through the space of possible events. Under these conditions, compositional learners will be preferred when:

$$n > \frac{3q}{pq_s}$$

In other words, there is a critical value for the number of objects and actions that must be exceeded before compositional learners are favoured. Nowak *et al.* give the following illustrative example. If nouns and verbs are twice as hard to memorise as holistic words ($q_c = q/2$) and if one third of all possible events actually occur in the environment ($p = 1/3$), then n must be greater than 18 before compositional learners are favoured.

This is an interesting result, and Nowak *et al.* successfully demonstrate that mathematical techniques can be fruitfully applied to the investigation of the evolution of learning apparatus underlying language. However, several criticisms of their model can be made. Firstly, as noted above, they assume that the relationship between events and words, or nouns and objects, or actions and verbs, is perfectly one-to-one. As we have seen throughout this thesis, arriving at this situation is far from straightforward. Secondly, they also rule out any gene-culture coevolution — they calculate the communicative accuracy of genetically homogeneous populations at equilibrium, then compare across populations. The situation is likely to be more complex in heterogeneous populations, and more complex still in heterogeneous populations undergoing selection for communicative success. Finally, they assume that events are randomly scattered in the event matrix — as we saw in Chapter 5, a non-random sampling from the space of possible meanings can have consequences for cultural evolution in populations, and therefore might have implications for gene-culture coevolution.

6.2 Languages, communication and communicative agents

The model of structured languages is identical to that used in the previous Chapter. A language L consists of a production function $p(m)$, mapping from meanings m to signals s, and a reception function $r(s)$, mapping from signals s to meanings m. Each $m \in \mathcal{M}$ is a vector drawn from an F-dimensional space, where each dimension has V possible values, and each signal $s \in \mathcal{S}$ is a string of characters of length 1 to l_{max}, where the characters are drawn from the alphabet Σ.

We can calculate the communicative accuracy of two individuals in exactly the same way as that outlined in Chapter 3, Section 3.2. If $p(m)$ is converted to a probabilistic function $p(s_j|m_i)$, which gives the probability of producing signal s_j given meaning m_i, and $r(s)$ is similarly viewed as a probabilistic function $r(m_i|s_j)$ then the communicative accuracy of a producer P with production function $p(s|m)$ signalling to a receiver R with reception function $r(m|s)$, averaged over all meanings, is:

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

In a population possessing an optimal communication system $ca(P, R) = 1$ for any choice of P and R.

Note that, given the distance function between meanings given in Chapter 5, based around Hamming distance, it would be possible to have a communicative accuracy measurement which awarded partial credit for getting a proportion of the meaning across to the receiver. This is the approach taken in Batali (1998) and Batali (2002). In this case $ca(P, R)$ would be defined as:

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

where $sim(m_i, m_k)$ gives the degree of similarity between two meanings, and is defined as:

$$sim(m_i, m_k) = \frac{(F - HD(m_i, m_k))}{F}$$

For the moment I will persevere with the all-or-nothing measurement. However, in Section 6.4, I will return to the partial payoff measure when considering the dual-transmission of compositional systems.

The model of a communicative agent is identical to that used in the preceding Chapter — each individual is modelled by an associative network capable of manipulating mappings between structured meanings and structured signals, and each individual acquires its system based on observation and the application of a weight-update rule W, specified by the 4-tuple $(\alpha\ \beta\ \gamma\ \delta)$.

6.3 Cultural evolution in populations

The simulations described in Chapter 5 demonstrate that, in populations consisting of a single individual, the cultural transmission of meaning-signal mappings leads to the emergence of compositional language. This is dependent on learners being biased in an appropriate fashion, the presence of a bottleneck on cultural transmission, and a degree of structure in the environment. Do these results scale up when we consider populations which consist of more than one individual at any one time?

This question can be addressed using a gradual population turnover ILM. The model of languages, communication and communicative agents is as given above in Section 6.2. The initialisation and iteration processes are given below.

Initialisation Create a population of N agents[2], each using the weight-update rule W and having an initial set of connection weights $\mathcal{W}$, where each $w \in \mathcal{W}$ has a weight of 0.

Iteration

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

[2] $N = 20$ for all ILMs outlined in this Section. In previous Chapters $N = 100$ was typically used. However, the more complex models of communication and communicative agents increases the computational cost of each cohort. In the gradual population turnover model computational complexity is constant with respect to population size, as each cohort involves replacing a single individual. However, larger populations take longer to converge on a shared system. $N = 20$ reduces this factor, while still allowing meaningful population-level dynamics.

4. The new agent receives e exposures to the population's observable behaviour and updates their connection weights according to the observed meaning-signal pairs and their weight-update rule W. See below for more detail.
5. The new agent joins the population. Return to 1.

Each pass through the iteration process will be termed a *cohort*, and as with other ILMs there is no genetic diversity within the population and no selection based on communicative ability.

Step 4 of the iteration process offers some complications. In the ILM outlined in Chapter 3, each of the e exposures consists of exposure to the complete set of observable behaviour generated by a single, randomly selected individual. In the model outlined in Chapter 5, each of the e exposures consisted of an exposure to a single meaning-signal pair produced by the individual's single cultural parent, and the exposures were either selected exhaustively from the environment $\mathcal{E}$ (in the no-bottleneck condition) or randomly (in the bottleneck condition).

In this model, neither of these methods of transmission is entirely suitable. If each of the e exposures consisted of exposure to the complete set of observable behaviour generated by a single, randomly selected individual then we immediately rule out a bottleneck on cultural transmission. If each of the e exposures consists of an exposure to a single meaning-signal pair produced by a single cultural parent, then convergence within the population will occur only by chance — true, non-random convergence requires that individuals sample the behaviour of several individuals. It is therefore necessary to introduce a new parameter τ, which is the number of cultural parents an individual has. Two versions of step 4 of the iteration process will be defined, one for the no-bottleneck condition and one for the bottleneck condition.

4 (No-bottleneck) The new agent selects τ cultural parents[3] at random from the population. The new agent receives $e = |\mathcal{E}|$ exposures to the communicative behaviour produced by those τ parents. During each of these e exposures the new agent observes the meaning-signal pairs produced by each parent for a single meaning $m \in \mathcal{E}$ and updates their connection weights according to the observed meaning-signal pairs and their weight-update rule W. Each $m \in \mathcal{E}$ is selected in turn, therefore the learner observes the full set of observable behaviour produced by each of the τ parents.

[3] $\tau = 3$ for all simulation runs reported here. This means that each individual will observe the behaviour of three individuals, which was the case for the associative network ILM discussed in Chapter 3.

4 (Bottleneck) The new agent selects τ cultural parents at random from the population. The new agent receives e exposures to the communicative behaviour produced by those τ parents. During each of these e exposures the new agent observes the meaning-signal pairs produced by each parent for a single, randomly selected, meaning and updates their connection weights according to the observed meaning-signal pairs and their weight-update rule W. The agent will therefore observe approximately $|\mathcal{E}| \cdot c(\mathcal{E}, e)$ distinct meanings, paired with the corresponding signals produced by each of the τ parents.

The no-bottleneck version leads, as will be discussed in Section 6.3.1, to the emergence of shared stable communication systems. However, the bottleneck version as given above does not. This appears to be due to the high level of variability in the behaviour observed by learners during the early stages of a simulation run. The bottleneck version of step 4 is therefore revised as follows. Each agent selects θ individuals at random from the population, where θ is randomly selected from the range $[1, \tau]$. The agent then selects τ cultural parents at random with replacement from among these θ individuals — in other words, learners are exposed to the same size of data set regardless of the number of distinct cultural parents they have, but the data set can contain the behaviour of between 1 and τ individuals. Each individual therefore receives τ exposures to each meaning, as in the bottleneck version of step 4 given above. However, these exposures will be to the behaviour of at most τ distinct individuals. Alternatively, given that $\tau = 3$ for all runs reported here, they will have 2 distinct cultural parents and observe one of them twice, or have a single cultural parent and observe that individual's behaviour three times[4].

I will consider an ILM where every agent uses the weight-update rule $W = (1\ \ -1\ \ -1\ \ 0)$. As shown in Chapter 5, this is one of the two [+constructor, +ic-preserver] rules. For the results described in Sections 6.3.1 and 6.3.2, $F = 3$, $V = 5$, $l_{max} = 3$ and $\sigma = \{a, b, c, d, e, f, g, h, i, j\}$. The initial agents have connection weights of 0, and therefore use the maximum entropy system where every meaning analysis-signal analysis pair occurs with equal probability. This is the same experimental setup as for the ILM described in Section 5.3, the only difference being the scaling up to larger populations.

[4]As part of my current research project, I am working on an extension to Kirby's (2002) model of the evolution of recursive syntax. One part of the project involves scaling this model up from populations consisting of a single individual to larger populations. Interestingly, a similar problem is encountered with Kirby's model — the set of cultural parents for each individual must be fairly tightly constrained, otherwise stable systems of meaning-signal mappings never emerge.

6.3.1 *Linguistic evolution in the absence of a bottleneck*

Runs of the ILM described above were carried out, using the no-bottleneck variant of step 4 — each individual observes the complete set of behaviour of $\tau = 3$ members of the population. 100 runs[5] of the ILM were carried out for each of the sparse environments shown in Figures 5.5 and 5.6 in Chapter 5. 50 runs of the ILM were carried out for each of the medium density environments shown in Figures 5.5 and 5.6 in Chapter 5. As in Chapter 5, the e-compositionality of the emergent languages (averaged over all members of the population) is the key measure of linguistic structure. The population's communicative accuracy is also measured, to establish whether the emergent languages are functional and shared by all members of the population. Communicative accuracy is estimated by evaluating every individual's average communicative accuracy as both producer and receiver with two randomly selected partners according to the all-or-nothing measure $ca(P, R)$ given in Section 6.2, averaging over all individuals in the population. Runs were allowed to proceed to a stable state, where the population exhibits no linguistic diversity.

In all simulations runs in each environment the populations converge upon an optimal shared communication system which yields $ca(P, R) = 1$ for any choice of P and R. This is as expected, given the one-to-one learning bias associated with the weight-update rule used by learners in these populations. Figures 6.1 and 6.2 plot the compositionality of the initial and final, stable systems for the sparse and medium-density environments. The results for the medium-density environment are similar to those shown in Figure 5.9 in Chapter 5 (for the same environments with an ILM involving isolated individuals).

The results for the sparse environments are rather different from the results from the single-individual ILM. In the isolated individual ILM (see Figure 5.8), the majority of runs converged on non-compositional systems. Partially compositional systems did occur, with their frequency being greatest when the environment was unstructured. Highly compositional systems were very infrequent, and occurred only when the environment was structured.

The results for the population ILM shown in Figure 6.1 show a much stronger tendency towards compositionality. The majority of the final systems are not holistic. Partially compositional systems occur with comparatively high frequency in both unstructured

[5] 1000 runs were carried out for the no-bottleneck condition of the single-individual ILM. A smaller number of runs were carried out in the population ILM due to two factors: 1) the increased computational memory requirements introduced by having 20, rather than 1, associative network in the population and 2) the increased number of cohorts required for a population to reach a stable state.

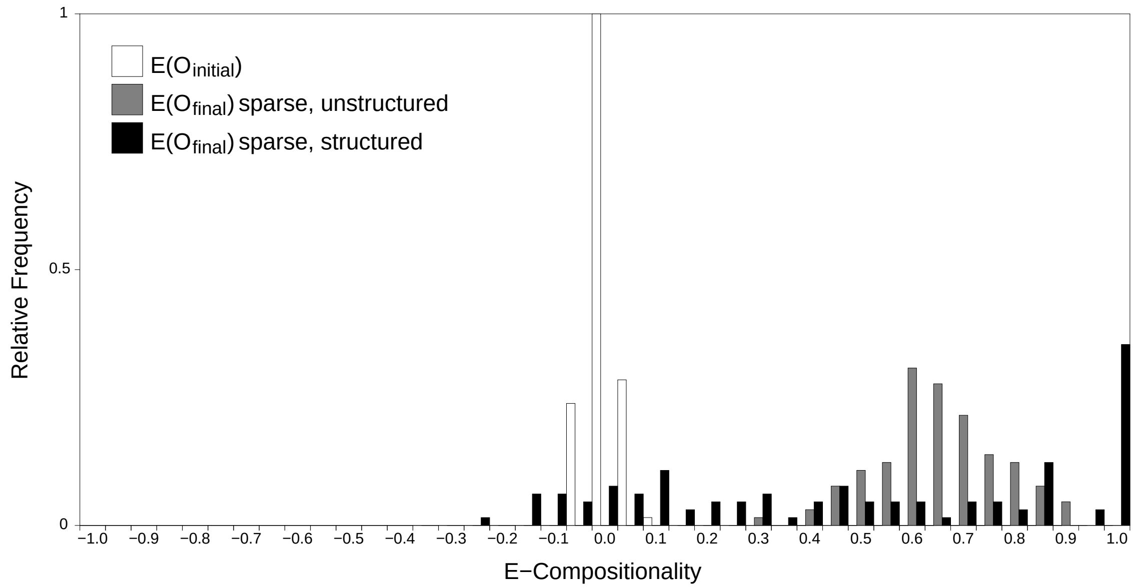

Figure 6.1: E-compositionality of initial and final, stable systems in sparse environments, when there is no bottleneck on transmission. The initial systems have low e-compositionality. The final systems are of partial or high e-compositionality. Highly e-compositional systems occur most frequently when the environment is structured.

and structured environments. Perfectly compositional systems emerge with fairly high frequency, but only when the environment is structured.

Why does the expansion to non-trivial population size lead to the more frequent emergence of compositionality, but only in the sparse environment? Recall from Chapter 5 that [+constructor, +ic-preserver] agents are biased in favour of acquiring i-compositional systems, and are further biased in favour of acquiring one-to-one mappings between feature values and signal substrings, which leads to a bias in favour of e-compositional language. In the single-agent population case, this bias can lead to the emergence of highly compositional language even in the absence of a bottleneck on cultural transmission, but only if the initial, random language already exhibits slight compositional tendencies. In the single-agent population case, learners are essentially stuck with the system of their single cultural parent. Their learning bias has a reduced impact, due to the absence of competing variants to select between (recall from B&R's model given in Chapter 2 that the rate of increase of the variant favoured by directly-biased transmission is dependent on the degree of cultural variation present in the population).

In the population ILM, each individual has several cultural parents and therefore biased acquisition potentially has a greater impact. An individual attempting to acquire two systems will be more influenced by the system which conforms more fully to their bias. In the population ILM, this means that highly compositional systems will be preferred to

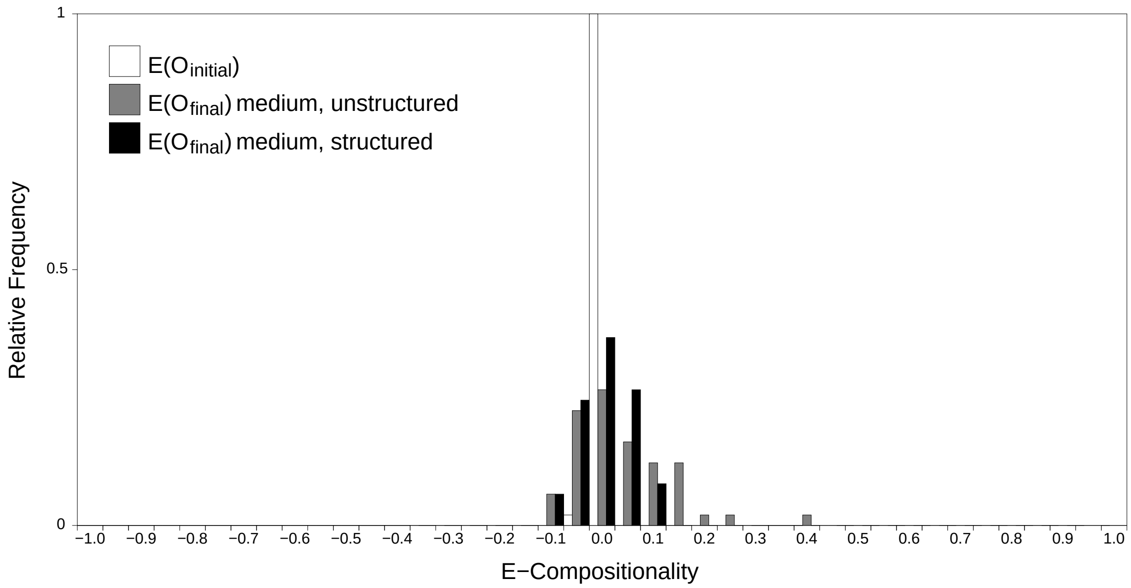

Figure 6.2: E-compositionality of initial and final, stable systems in medium density environments, when there is no bottleneck on transmission. The initial systems and the vast majority of the final systems have low e-compositionality. Partially e-compositional final systems occur with very low frequency, and only when the environment is unstructured.

less compositional systems. This results, in sparse environments, in the frequent emergence of highly compositional systems. The difference between unstructured and structured sparse environments is due, as discussed in Chapter 5, to the greater potential spread of compositional mappings in the structured environment, due to the number of feature values shared between meanings.

The fact that each learner has several cultural parents, and therefore several possible communication systems to choose from, increases the force of the learner's bias and results in the emergence of systems which conform to that bias. However, this only happens in sparse environments — in medium density environments, the final stable systems tend overwhelmingly to be holistic, with partially compositional systems occurring very infrequently and only when the environment is unstructured. Why?

Recall from Chapter 5 that the compositionality of the final system in the single-individual ILM is sensitive to the compositionality of the initial, random system. Where this initial mapping exhibits compositional tendencies, yielding $E\left(\mathcal{O}_{initial}\right)$ above the mean, there is an increased likelihood of the system moving, over iterated learning events, towards more compositional languages. The compositional tendencies of the initial system spread to other parts of the system over time, resulting in an increase in compositionality. For the more densely-filled environments, partially or highly compositional systems emerge infrequently due to the fact that the initial systems tend to be clustered more tightly around the non-compositional mean. When the environment contains few meanings the initial system may, by chance, exhibit some compositional tendencies. However, when the environment contains a large number of meanings such tendencies are likely to be drowned out by the majority non-compositional mapping.

In the population ILM, with medium density environments, the lack of compositional tendencies in the early, random mappings of the population prevents highly compositional systems from ever emerging. Even though each learner observes several individuals, their set of cultural parents is essentially homogeneous with respect to compositionality — each parent uses a non-compositional system (although not necessarily the same one). This lack of cultural variation effectively nullifies the learner preference for compositionality. In contrast, in sparse environments the initial random systems are more widely distributed, and more likely to exhibit some compositional tendencies which the learner bias can exploit.

Partially compositional systems do emerge with low frequency in medium density, unstructured environments. This is due to the fact that, in such environments, fewer meanings share feature values, therefore the initial random system is more likely to exhibit slight compositional tendencies — the initial systems in unstructured environments has to be less 'lucky' in the assignment of characters to feature values. This can provide some cultural variation among an individual's cultural parents, allowing the learner bias to have some effect.

6.3.2 Linguistic evolution in the presence of a bottleneck

The simulation results outlined in the previous Section show that, in the absence of a bottleneck on cultural transmission, highly compositional languages can emerge in populations. Their emergence is dependent on the density and structure of the environment, and there is a degree of sensitivity to the compositionality of the original, random systems of meaning-signal mappings. It is now time to investigate how a transmission bottleneck impacts on the compositionality of emergent systems in populations.

To this end, runs of the ILM described above were carried out, using the bottleneck variant of step 4 — each individual observes e meaning-signal pairs, randomly selected from the set of behaviour produced by $\theta \leq \tau = 3$ different members of the population. 10 runs of the ILM were carried out for each of the sparse environments shown in Figures 5.5 and 5.6 in Chapter 5, with a bottleneck of $c(\mathcal{E}, e) = 0.8$ (e = 19) and 10 runs were carried out for each of the medium density environments shown in Figures 5.5 and 5.6

(a)

Density	$c(\mathcal{E}, e)$	Proportion compositional	Average $E(O_{final})$	Average *ca* (final)
sparse	0.8	0	0.59	0.51
medium	0.4	0	0.28	0.05
medium	0.5	0.2	0.47	0.26
medium	0.6	1.0	0.99	1.0

(b)

Density	$c(\mathcal{E}, e)$	Proportion compositional	Average $E(O_{final})$	Average *ca* (final)
sparse	0.8	1	1.0	1.0
medium	0.4	1	0.99	0.98
medium	0.5	1	0.99	0.99
medium	0.6	1	0.99	1.0

Table 6.1: Summary of results for the population ILM. (a) gives the proportion of runs converging on a highly compositional system, the average e-compositionality of the final systems and the average communicative accuracy yielded by the final systems for *unstructured* environments. Highly compositional, communicatively-optimal languages only reliably emerge in the medium density when the bottleneck is wide. (b) gives the same measurements for runs of the ILM in *structured* environments. Highly compositional, communicatively-optimal languages always emerge when the environment is structured.

with bottlenecks of $c(\mathcal{E}, e) = 0.4$, 0.5 and 0.6 (e = 16, 21 and 28 respectively)[6]. Runs were allowed to proceed for 5000 cohorts.

Table 6.1 summarises the results of these simulation runs, in terms of the proportion of runs converging on a highly compositional system ($E(\mathcal{O}_{final}) > 0.95$), and the average final levels of e-compositionality and communicative accuracy. Figure 6.3 show compositionality and communicative accuracy against time in five representative runs of the ILM.

As shown in the Table, environment structure has a significant impact on the compositionality of the emergent systems. In the unstructured environments, highly compositional, communicatively-optimal systems only reliably emerge in the medium density environment with a relatively wide bottleneck ($c(\mathcal{E}, e) = 0.6$). In contrast, in the structured environments, highly compositional, optimal systems *always* emerge.

[6] This is obviously a significantly smaller number of runs than was carried out for the single-individual ILM, and is due to the increased computational cost of population-level ILMs, as discussed above. Each run of the population ILM with a medium density environment takes approximately 36 hours on a 2.5GHz Pentium 4 processor.

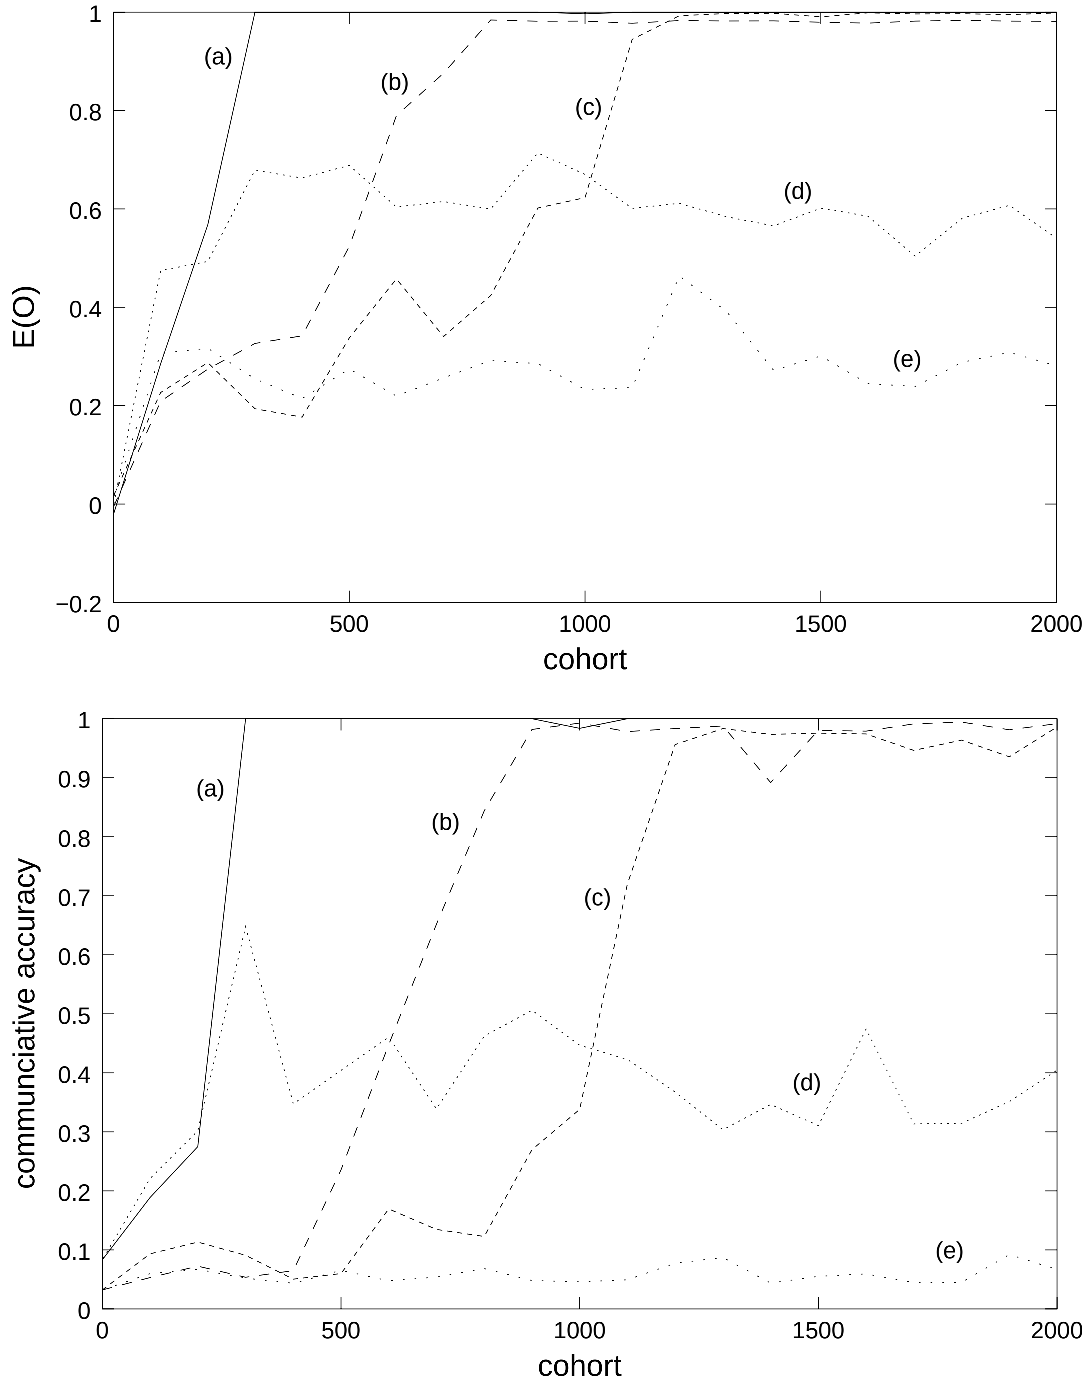

Figure 6.3: Plots of e-compositionality (top) and communicative accuracy (bottom) against time in runs of the population ILM. (a) shows the progress of a simulation run in the sparse, structured environment. (b) is an example of a convergent run in a medium density, structured environment. (c) is a convergent run from a medium density, unstructured environment where $c = 0.6$. (d) is a non-convergent run from the sparse, unstructured environment. (e) is a non-convergent run from a medium density, unstructured environment with a tight bottleneck on transmission ($c = 0.4$).

Learners in the population ILM observe and learn from the linguistic behaviour of between one and three cultural parents. If these cultural parents have strongly conflicting languages, or if their languages are non-compositional, then the learner will tend to arrive at a non-compositional or partially compositional system, depending on the number of meanings in the environment (as discussed above, the e-compositionality of random systems is sensitive to the number of meanings expressed in that system). If, on the other hand, there is broad agreement in the linguistic behaviour of a learner's cultural parents, then the learner will converge on a system which is similar to that of its cultural parents, and which exhibits a similar level of compositionality.

Stability and regularity at the individual level therefore form the basis for convergence in the population. In structured environments, this individual-level stability emerges fairly straightforwardly. Compositional languages have a strong advantage in such environments, due to their generalizability. When the environment is structured, individual members of the population will converge fairly quickly on systems which are at least partially compositional. These systems will then spread fairly rapidly through the population, until the population converges on a communicatively optimal, compositional language. This is reflected in the reliable emergence of compositional language in structured environments, regardless of the degree of environment density or the severity of transmission bottleneck. These results show that the ILM results from the previous Chapter can scale up to the case where the population at any one time consists of more than one individual.

However, in unstructured environments stability at the individual level is rather more difficult to achieve, due to the lack of shared feature values between meanings. This problem can be overridden by a relatively wide bottleneck. For example, when learners see 60% of the language of the previous generation, there is very little difference between structured and unstructured environments, as can be seen in the final distribution of languages in Figure 5.14 in Chapter 5. The wide bottleneck allows individuals to reach highly compositional, consistent systems. As shown in the Table, in the population ILM these systems spread through the population — when the bottleneck is wide ($c = 0.6$), highly compositional languages always emerge.

For the lower levels of bottleneck in the medium environments, and in the sparse environment, the lack of structure in the environment is more problematic. As can be seen from Figures 5.12 and 5.13 in Chapter 5, in the single individual ILMs, for tight bottlenecks, unstructured environments lead to languages which are partially compositional, and therefore only partially stable.

What consequences does this have in the multi-agent population model? Each learner observes and learns from the communicative behaviour of several other individuals. If these individuals are using a partially compositional system then there will be some randomness in their linguistic behaviour — while a learner's cultural parents may share some part of the meaning-signal mapping, some of their behaviour will be random and therefore unlikely to be shared. This means that the learner will receive contradictory learning input — each of their cultural parents will produce different signals for certain meanings. As discussed above, this means that the learner's bias and the number of meanings in the environment become important. In the sparse environment, partially compositional systems can still emerge — based on a set of conflicting observations, the learners tend to arrive at a partially compositional system. The system never becomes perfectly compositional due to the bottleneck, which reintroduces instability.

However, in the medium density environments, such compositional systems never get off the ground — as can be seen from the Table, highly compositional systems emerge infrequently in unstructured environments when the bottleneck is relatively tight ($c = 0.4$ or 0.5). In these circumstances, learners will be faced with a large set of contradictory input. As a consequence, they will tend to acquire a non-compositional system — as discussed above, given a large number of meanings, the majority non-compositional mapping tends to drown out any weak compositional tendencies. As a consequence, partially compositional systems do not emerge in unstructured environments — the system of meaning-signal mappings remains random, with each individual's system tending to be rather different from the systems of other members of the population.

6.3.3 Summary

Communicatively optimal, compositional languages can emerge in populations through purely cultural processes. When there is no bottleneck on cultural transmission, the fact that learners make observations of several cultural parents increases the impact of their learning bias, effectively allowing them to pick the system of meaning-signal mappings which most closely matches their bias. This leads to the emergence of compositional languages with reasonably high frequency, although only when the environment is sparse and structured — in the absence of a bottleneck, the results are sensitive both to the compositionality of the early systems and to the potential for spread of compositionality.

In the presence of a bottleneck, highly compositional systems emerge with high frequency when the environment is structured. However, when the environment is unstructured such systems only emerge when the bottleneck is relatively wide. Convergence on a compositional language first requires a degree of stability at the individual level. This is straightforwardly achieved when the environment is structured, due to the high potential for generalisation. In unstructured environments, this stability can be achieved when the bottleneck on transmission is not too tight. However, when the bottleneck is tight individual members of the population never arrive at stable systems, and as a consequence the population never converges on a shared language.

6.4 The evolution of learning biases for compositional language

We have established that compositional language can evolve through cultural processes in a population, provided that learners have the appropriate learning bias (of the [+constructor, +ic-preserver] classification). The final question is to investigate whether this learning bias can evolve through natural selection for communicative success. The simulation results described in Chapter 4 suggest that one-to-one biases for vocabulary acquisition are unlikely to evolve specifically for their communicative function, due to the time delay between the emergence of such a bias and a communicative payoff for individuals possessing it. This should make us skeptical as to whether such a bias can evolve for the acquisition of a (potentially) structured system of meaning-signal mappings.

The model of languages and communication is as described in Section 6.2. As discussed in that Section, there are two possible methods of evaluating communicative accuracy between two individuals — one which counts a communicative episode as a success only if speaker and hearer arrive at exactly the same meaning, and one which gives partial credit for speaker and hearer arriving at partially overlapping meanings. I will investigate both alternatives here.

6.4.1 Genotypes, phenotypes and reproduction

The model of a phenotype communicative agent is as described in Section 6.2 — an associative network capable of representing structured meaning-signal mappings, with an initial set of connection weights $\mathcal{W}$ and a weight-update rule W.

As in the EILM for the simple associative network outlined in Chapter 4, Section 4.5, I will assume that an individual's weight-update rule W is genetically-encoded. 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. 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.

To recap, there are 81 possible genotypes, which encode the 81 possible weight-update rules discussed in Chapter 5, Section 5.4. These 81 weight-update rules can be split into four classifications:

- 63 are classified as [$-$maintainer], and are therefore unable to acquire an e-compositional language.
- 11 are classified as [$+$maintainer, $\pm$constructor, $-$ic-preserver], and are able to acquire an e-compositional language, but represent it in an internally-holistic fashion.
- 5 are classified as [$+$maintainer, $-$constructor, $+$ic-preserver], and are able to acquire an e-compositional language, but unable to maintain such a language in the presence of a bottleneck.
- 2 are classified as [$+$constructor, $+$ic-preserver], and are able to acquire, maintain and construct an e-compositional language.

Individuals inherit their genes from their parents. As in earlier EILMs, organisms are haploid but sexual recombination (involving crossover, in an identical fashion to that outlined for the previous EILMs) is used. Newly-formed genotypes are also subject to mutation.[7]

6.4.2 *The EILM*

A gradual EILM is used — at each cohort, a single individual is selectively removed from the population, the remaining members of the population breed according to communicative success to produce a new individual, and that new individual acquires its communication system based on observations of the population's behaviour.

Initialisation Create a population of N agents[8]. 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.

[7]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.

[8]$N = 50$ for all simulations outlined in this section.

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 the environment $\mathcal{E}$.
3. Create a new agent. The new agent inherits their genotype from their parents, who are selected from the population according to the reproduction procedure outlined below.
4. The new agent selects θ individuals at random from the population, where θ is randomly selected from the range $[1, \tau]$. The agent then selects τ cultural parents at random from among these θ individuals[9]. The new agent receives e exposures to the communicative behaviour produced by those cultural parents. During each of these e exposures[10] the new agent observes the meaning-signal pairs produced by each parent for a single, randomly selected meaning and updates their connection weights according to the observed meaning-signal pairs and their weight-update rule W. The agent will therefore observe approximately $|\mathcal{E}| \cdot c(\mathcal{E}, e)$ distinct meanings, paired with their corresponding signals produced by each of the τ parents.
5. The new agent joins the population. Return to 1.

As with the EILM outlined in Chapter 4, Section 4.5, tournament selection is used to determine reproduction and death. During each tournament T individuals [11] are selected from the population at random and evaluated. Each individual is scored according to their average communicative accuracy (according to one of the two measures) when acting as both producer and receiver with two randomly selected partners. 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.

Note from the iteration procedure that each agent observes a subset of the language of its cultural parents — there is a bottleneck on cultural transmission.

[9] $\tau = 3$ for the runs outlined here.

[10] $e = 24$ for all EILMs outlined in this section, which yields a bottleneck of $c(\mathcal{E}, e) = 0.6$ with respect to the environment described below — each learner observes approximately 60% of the language of its cultural parents.

[11] As in the simple associative network EILM, $T = 3$.

6.4.3 The environment

As discussed above with reference to the population ILM, the associative network model is computationally expensive, both in terms of memory and CPU cycles, particularly when used in the context of a population. This is largely due to the large size of the associative network, and the large number of analysis pairs which have to be evaluated during production and reception.

These problems can be alleviated by reducing the number of feature values (V) and the size of the character inventory ($|\Sigma|$). To this end, for all EILMs outlined in this Section $F = 3$, $V = 3$, $l_{max} = 3$, $\Sigma = \{a, b, c, d, e, f\}$.

This selection of F and V means that the environments used in Chapters 5 and Section 6.3 of this Chapter can no longer be used, due to the changed space of possible meanings $\mathcal{M}$. Instead, an environment is used where $\mathcal{E} = \mathcal{M}$ — every possible meaning in the meaning space is present in the environment. This allows us to simplify away from the structured-unstructured distinction with respect to environments.

The change in environment, number of exposures ($e = 24$ is used in the EILM) and also the change in the population size ($N = 20$ in the population ILM, whereas $N = 50$ in the population EILM) compared to Section 6.3 makes it necessary to rerun the ILM with the new environment and population size. Ten runs of the ILM were carried out, using a [+constructor, +ic-preserver] weight-update rule. All runs converged on a communicatively optimal, highly compositional ($E(\mathcal{O}) \geq 0.95$) language, with the mean time to convergence being 3680 cohorts, although half of the runs converged on a stable system within 2000 cohorts. These runs are plotted in Figure 6.4. This demonstrates that, as before, given the appropriate learning bias, communicatively optimal, compositional language can emerge through cultural processes given this experimental setup.

6.4.4 A negative result

Ten runs of the EILM were carried out, using the all-or-nothing evaluation of communicative accuracy given in Section 6.2 — a communicative episode was only considered a success if speaker and hearer arrived at exactly the same meaning. Runs were allowed to proceed for 10000 cohorts. None of these runs converged on a communicatively-optimal or compositional communication system — all runs remained stuck with a random, e-holistic communication system, which yields chance levels of communicative accuracy. All populations became fixated on genotypes which encoded [$-$maintainer] weight-update rules.

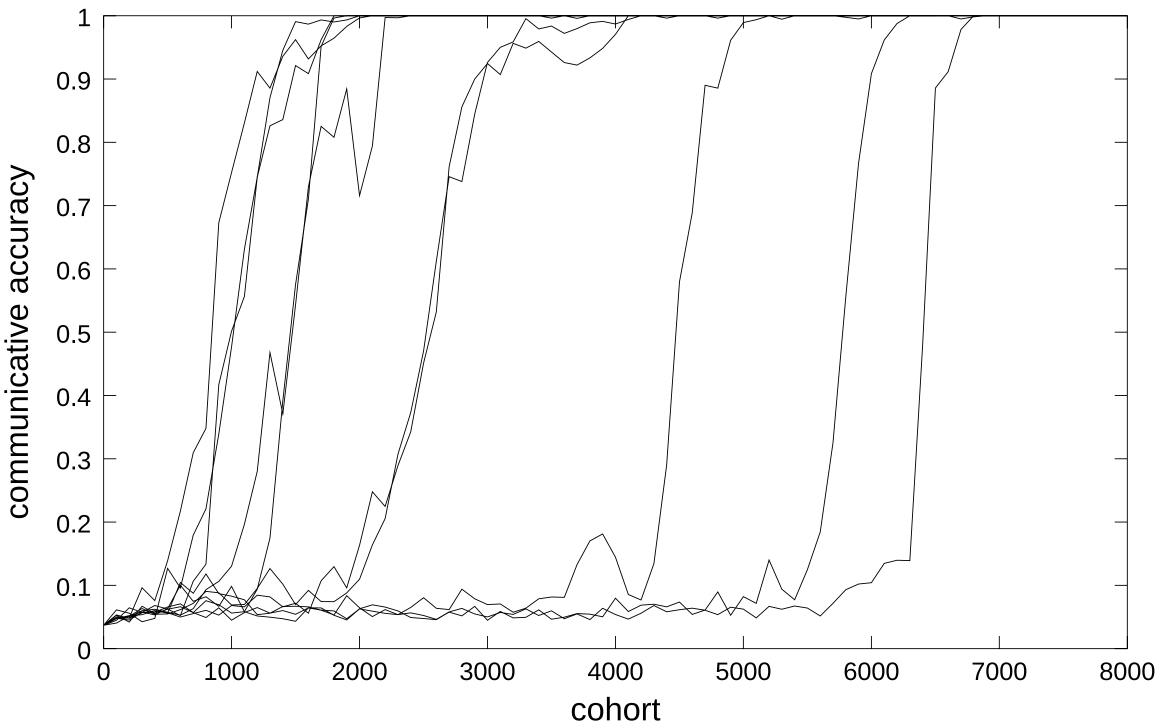

Figure 6.4: Communicative accuracy against time in the ILM with the parameter setting which will be used for the EILM. Communicative accuracy was evaluated according to the all-or-nothing measure. All runs converge on an optimal system, although time to convergence varies considerably from run to run.

These results are unsurprising. We saw in Chapter 4 that appropriate learning biases are unlikely to evolve, given the time-lag between the emergence of such biases and a payoff to individuals possessing them. The evolutionary task for the populations here is much harder — only two of the 81 genotypes are any use (as opposed to nine of 81 in the associative network EILM in Chapter 4) and cultural convergence, even given the correct learning bias, is potentially somewhat slow.

6.4.5 A positive result: the evolution of learning biases for compositional language

A further ten runs of the EILM were carried out, using the partial credit evaluation of communicative accuracy given in Section 6.2 — individuals receive a payoff from communication which is proportional to the similarity between the meaning the speaker was attempting to convey and the meaning the hearer arrives at.

Figures 6.5 and 6.6 show the progress of a simulation run where the population constructs a communicatively optimal, compositional language. This is a typical example of a successful run. Five of the ten EILM runs were successful in this respect — learning

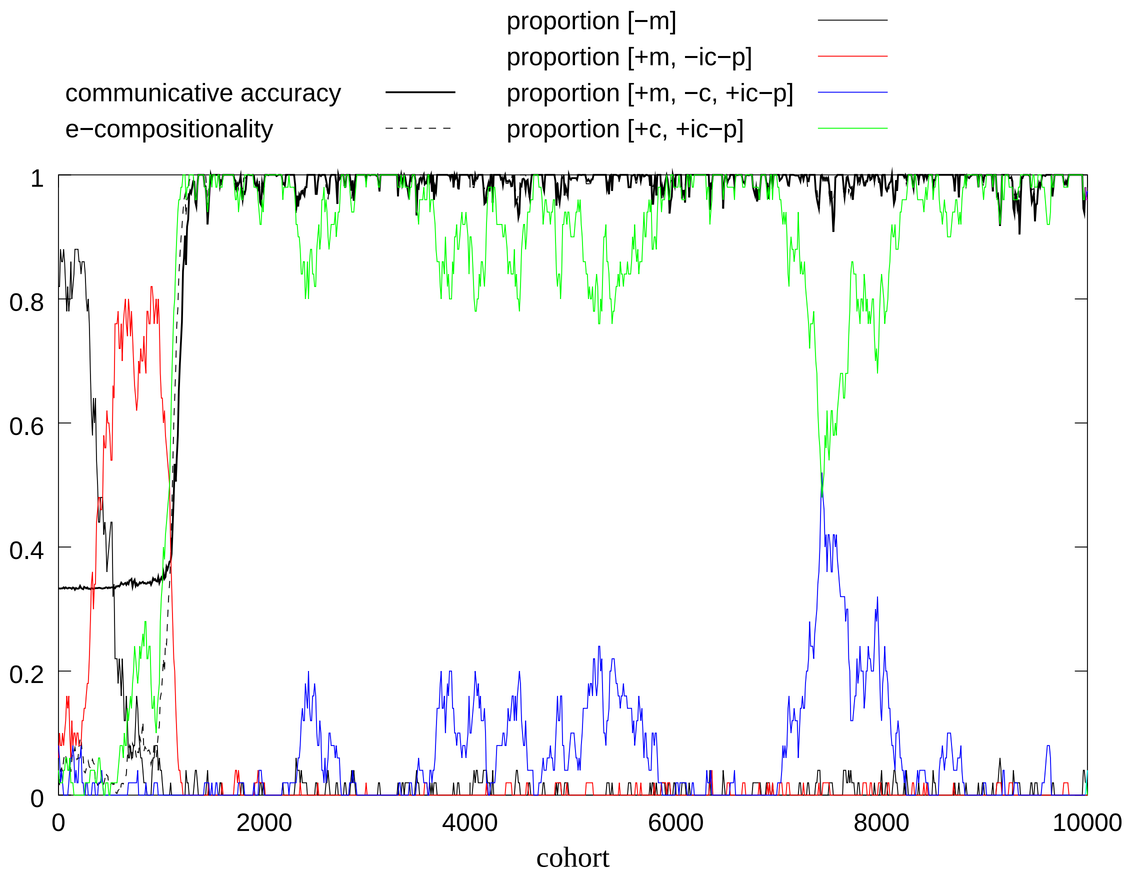

Figure 6.5: The evolution of learning bias leading to communicatively optimal, compositional language. Proportions of various groups of genotypes are given ([±m] stands for [±maintainer], [±c] stands for [±constructor], [±ic-p] stands for [±ic-preserver]). The population's communicative accuracy and compositionality reach maximal values with 2000 cohorts. The population comes to be dominated by [+constructor, +ic-preserver] weight-update rules, although [+maintainer, −constructor, +ic-preserver] weight-update rules do drift in and out after 2000 cohorts.

biases supporting the cultural evolution of a compositional language emerge 50% of the time. Figure 6.7 shows the relationships between the population's average communicative accuracy, e-compositionality and the proportion of individuals in the population with weight-update rules encoding [+constructor, +ic-preserver] weight-update rules. There is a clear relationship — as the number of [+constructor, +ic-preserver] individuals in the population increases, so too does compositionality and, latterly, communicative accuracy.

In Chapter 4, we saw that the evolution of one-to-one biases for vocabulary acquisition consisted of three stages — an initial stage of drift, a stage of selection for the appropriate learning biases, then a further stage of drift. This same three-stage process is evident

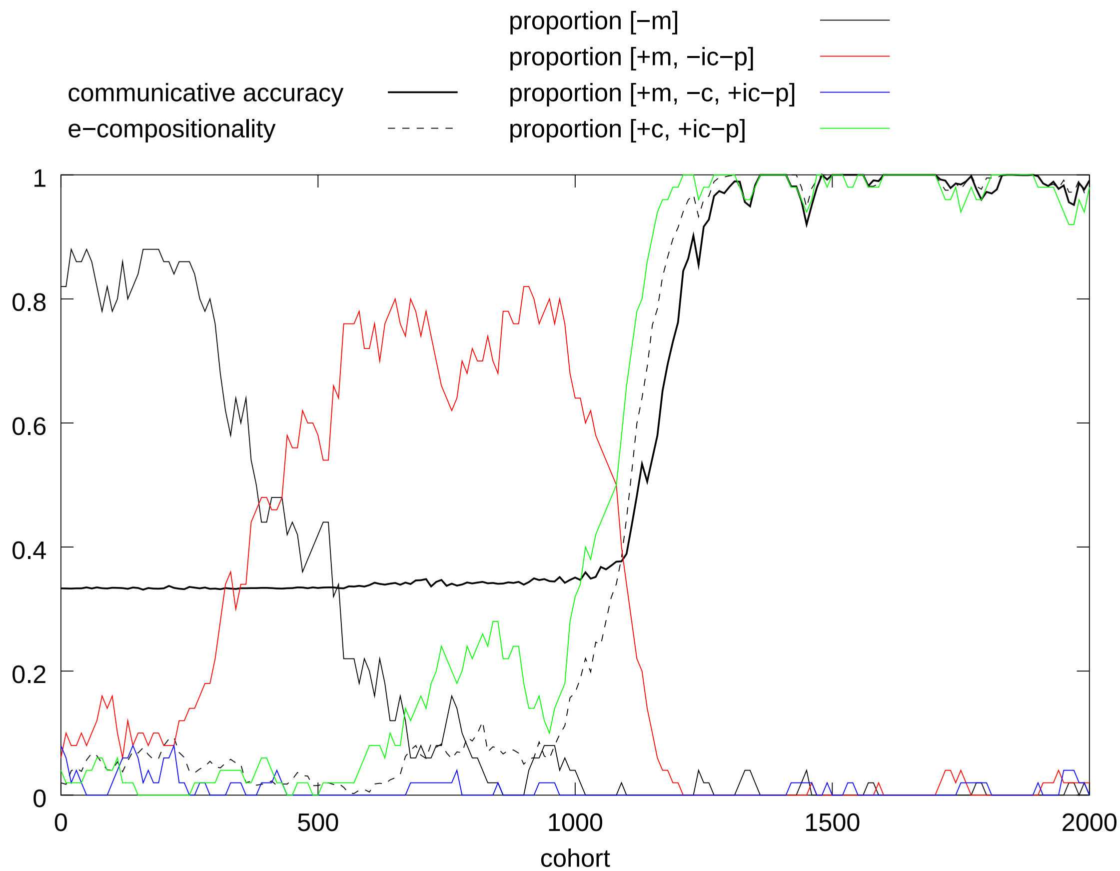

Figure 6.6: The first 2000 cohorts of the simulation run in Figure 6.5. At about 900 cohorts the population is dominated by [+maintainer, ±constructor, −ic-preserver] individuals. There are also a small number of [+constructor, +ic-preserver] individuals present. The numbers of the [+constructor, +ic-preserver] individuals increases sharply from 900 cohorts, and as a consequence the communicative accuracy and e-compositionality of the population's language increases to maximum values.

in runs where learning biases supporting communicatively optimal, compositional language emerge. Figures 6.8–6.11 plot the relative communicative accuracies (*rca*s) for four classes of genotypes in this run. As can be seen from these Figures, the first 2000 cohorts of the simulation run consists of a stage of drift, followed by a stage of selection. The initial drift stage lasts from 0 to 900 cohorts. During this time the *rca* of all four classes of genotypes remains around 1, indicating that no genotype is associated with above-average levels of communicative accuracy. Access to breeding in the population is random for this time. Genetic drift results in an increase in the numbers of [+maintainer, −ic-preserver] individuals (incapable of constructing an optimal system through a bottleneck), and later an increase in the number of [+constructor, +ic-preserver] individuals (capable of constructing such a system).

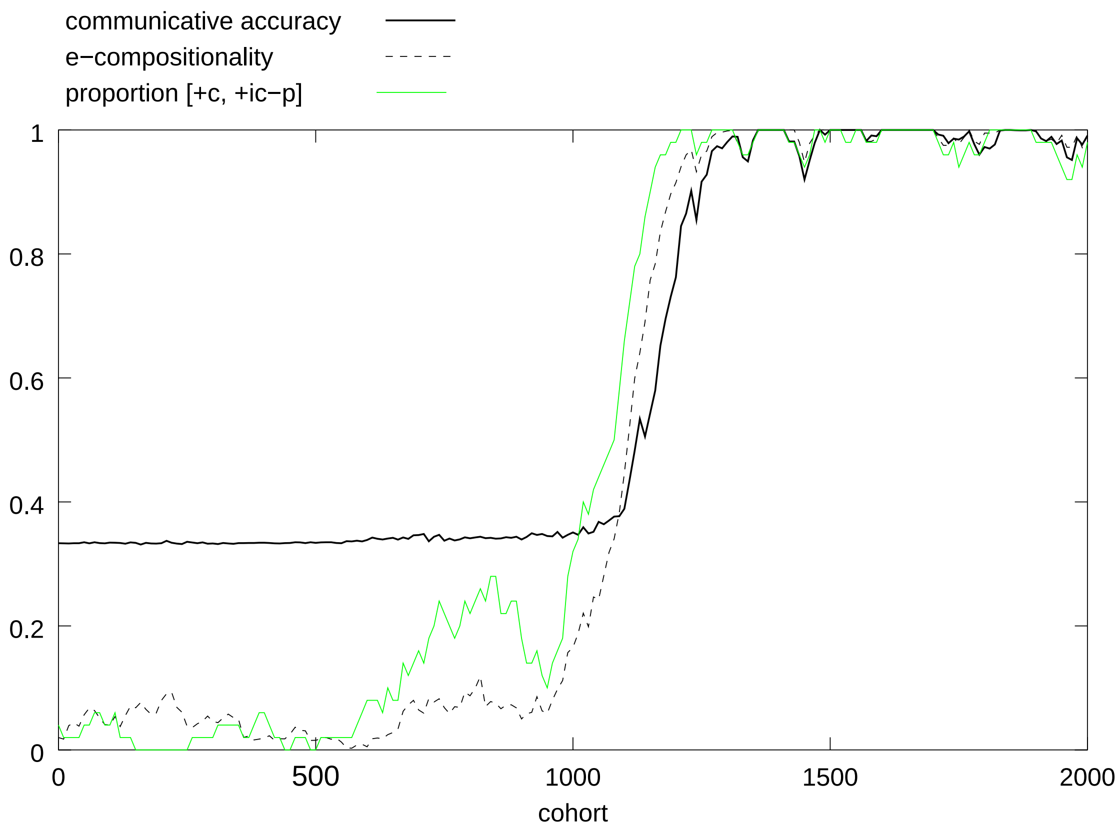

Figure 6.7: The relationship between communicative accuracy, e-compositionality and proportion of [+constructor, +ic-preserver] individuals in the successful run.

The mini peak of [+constructor, +ic-preserver] individuals results in the beginnings of a communicatively useful, partially compositional system of meaning-signal mappings. Consequently, individuals with such genotypes, who are capable of acquiring and contributing to the construction of such a system, receive a communicative payoff — *rca* for these genotypes rises above 1, they receive disproportionate access to breeding roles and their numbers increase sharply in the population. *rca* for other genotypes (particularly [+maintainer, −ic-preserver] genotypes, which form a significant proportion of the population up until 900 cohorts) drops below 1, and their numbers decrease sharply. Individuals with genotypes which make them capable of acquiring and constructing an optimal, compositional language are selected for, to the detriment of other genotypes.

This selection proceeds until the population consists entirely of [+constructor, +ic-preserver] individuals. Shortly after, compositionality and communicative accuracy reach maximum levels — the population converges on a communicatively optimal, compositional language. A second period of drift then ensues. During this period, as can be seen in Figure 6.5, [+maintainer, −constructor, +ic-preserver] individuals are introduced

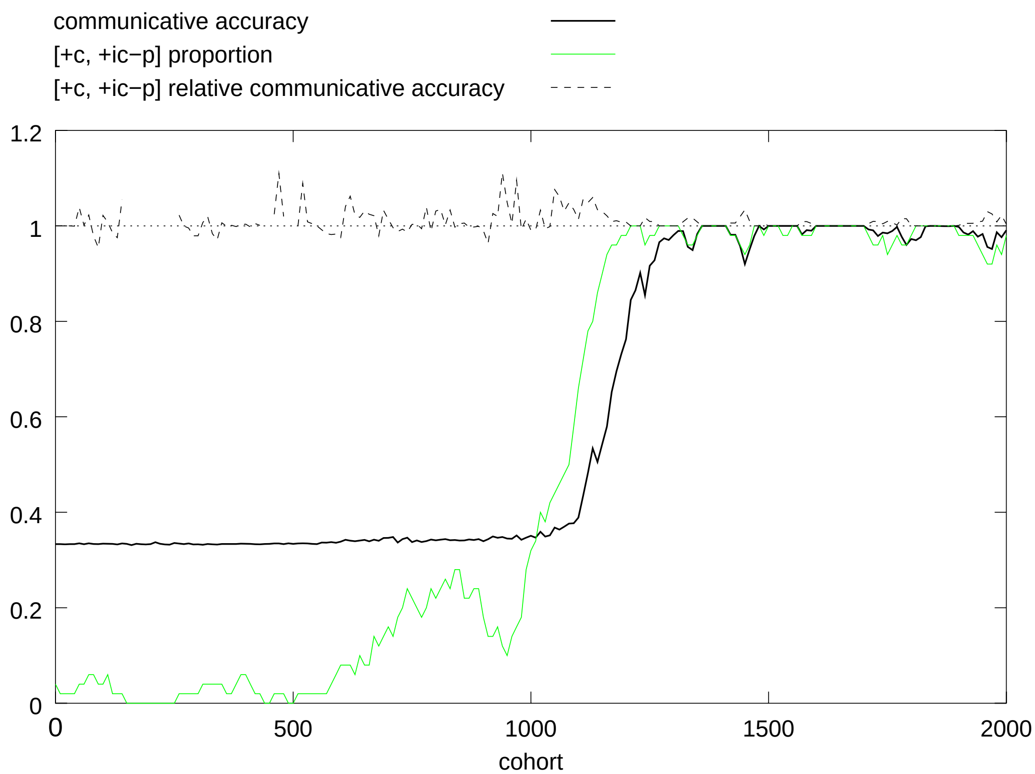

Figure 6.8: The relative communicative accuracy of individuals with [+constructor, +ic-preserver] weight-update rules. This value fluctuates around 1, and is clearly above one for the period from 900 to 1200 cohorts, at which point the numbers of such individuals increase sharply.

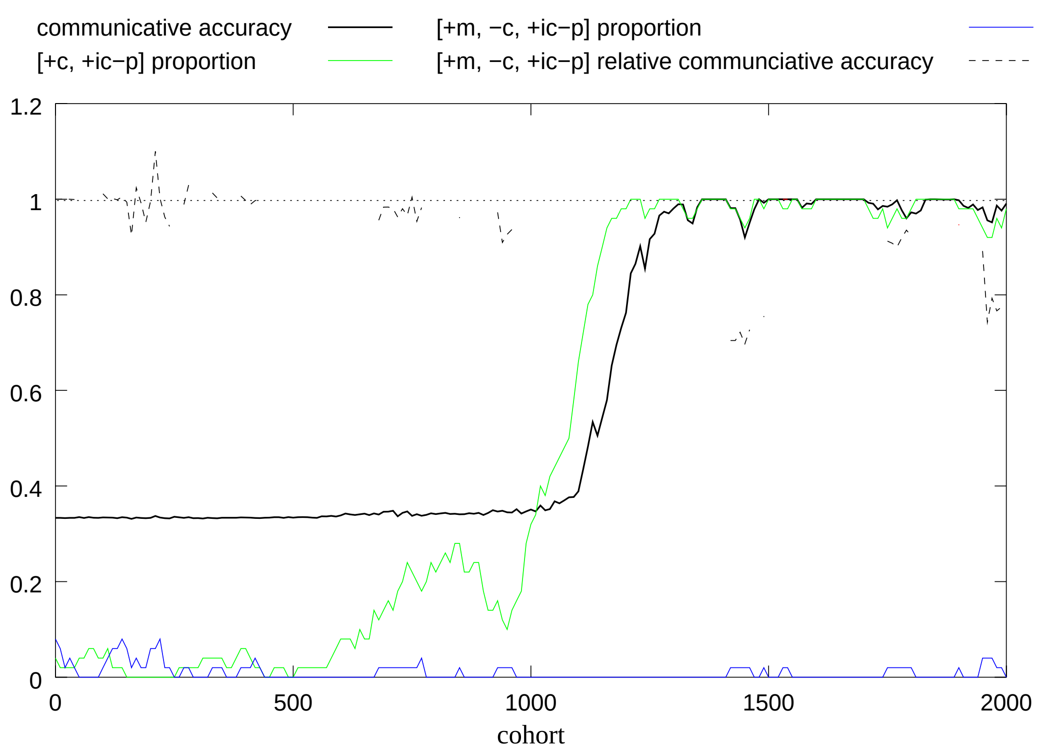

Figure 6.9: The relative communicative accuracy of individuals with [+maintainer, −constructor, +ic-preserver] weight-update rules. These individuals are not present in significant numbers in the early stages of the run.

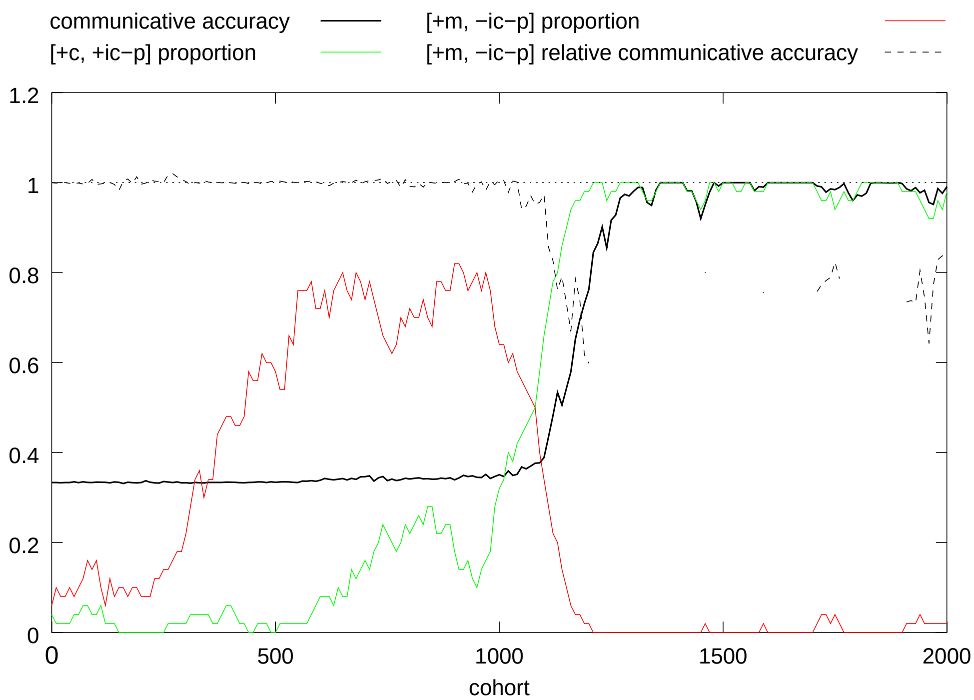

Figure 6.10: The relative communicative accuracy of individuals with [+maintainer, ±constructor, −ic-preserver] weight-update rules. The numbers of these individuals increases from 200 to 600 cohorts. However, the fact that their *rca* remains at 1 suggests that this increase is due to drift. Their numbers drop sharply from 900 cohorts, at which point the number of [+constructor, +ic-preserver] individuals increases sharply. The *rca* of [+maintainer, ±constructor, −ic-preserver] drops below 1 around this point.

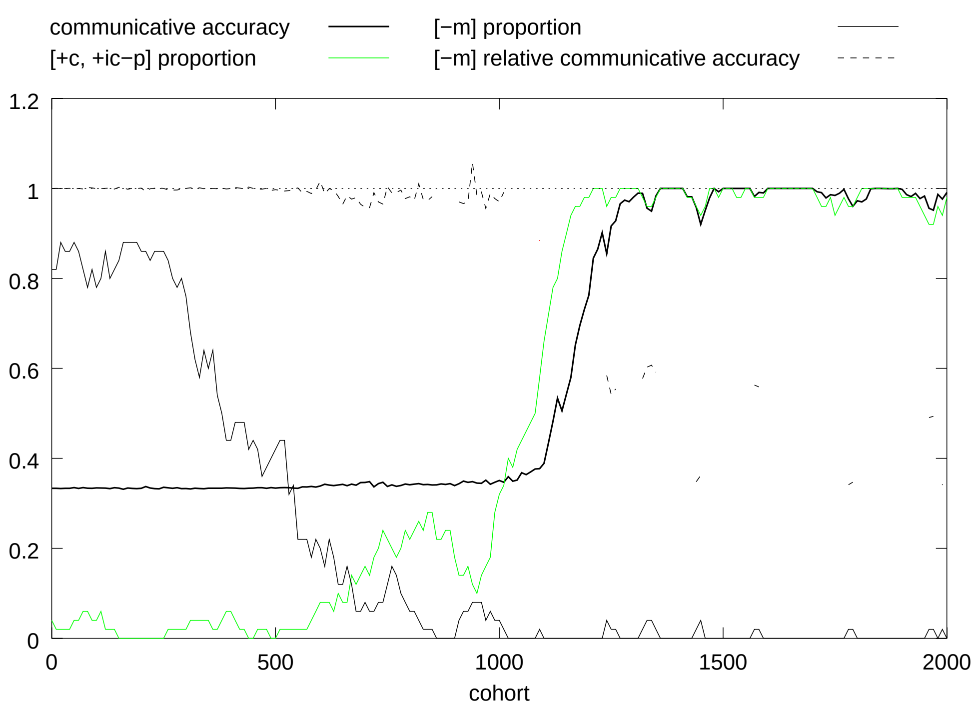

Figure 6.11: The relative communicative accuracy of individuals with [−maintainer] weight-update rules. Their numbers decline from 200 to 700 cohorts. The fact that their *rca* remains around 1 at this point suggests that this is due to drift.

into the population, and their numbers fluctuate randomly, although generally remaining fairly low. We saw in Chapter 5 that individuals using [+maintainer, −constructor, +ic-preserver] cannot maintain an optimal compositional system in a single-agent ILM. However, in a mixed population with [+constructor, +ic-preserver] individuals, given a reasonably wide bottleneck on transmission, such individuals can maintain the optimal system, provided their numbers do not get too great. Consequently, drift allows them to enter the population. At around 7500 cohorts they make up 50% of the population. At this point, they begin to lose the optimal system, the population's overall communicative accuracy drops somewhat, and the [+maintainer, −constructor, +ic-preserver] individuals are selected against until their numbers drop to lower levels.

The successful runs of the EILM therefore exhibit the same three-stage process of drift-selection-drift that we saw in Chapter 4, although the second period of drift is somewhat more constrained. The emergence of the optimal learning bias is therefore dependent on an initial period of drift. In the EILM in Chapter 4 this resulted in a very low number of runs converging on optimal system. In this EILM 50% of runs converge on optimal systems. Why?

The partial-credit measurement of communicative accuracy plays an important role — if the all-or-nothing measurement is used, optimal systems never emerge. The partial-credit measurement reduces the time it takes individuals with [+constructor, +ic-preserver] weight-update rules to get some communicative payoff, and therefore reduces the period of vulnerability to drift. Individuals who learn using [+constructor, +ic-preserver] weight-update rules will not necessarily arrive at the same overall system of meaning-signal mappings, but they may, based on commonalities in the observable behaviour they learn from, arrive at a shared system of mappings from one or two feature values to signal substrings. In the all-or-nothing measurement of communicative accuracy, there is no reward for this — the whole meaning must be correct. However, under the partial-credit scheme, these individuals will receive some small communicative payoff, and therefore be more likely to breed and add more [+constructor, +ic-preserver] individuals to the population. The partial-credit measurement of communicative accuracy smoothes the fitness landscape, making the evolution of appropriate learning biases more straightforward.

A second reason that appropriate learning biases evolve so rapidly is due to the speed of cultural convergence in the EILM — the population moves from random levels of communicative accuracy to optimal levels in around 300 cohorts. As discussed in Chapter 4, speed of cultural convergence plays an important role in the EILM, with a reduction in the time to cultural convergence leading to less sensitivity to drift and consequently more frequent genetic convergence.

The fact that the population in the EILM converges so quickly appears to be at odds with the results from the ILM discussed in Section 6.4.3 above — in a pure ILM, convergence takes on the order of thousands, rather than hundreds, of generations. There are two factors that can explain this disparity. Firstly, in the EILM there is (weak) natural selection of cultural variants, which weeds out systems of meaning-signal mappings which offer below-average communicative accuracy. Secondly, and more importantly, in the ILM the initial system is completely random. In the EILM, at the point where the construction of the optimal system takes off, the population's communication system is not entirely random — 1000 cohorts of 'preparatory' cultural evolution have occurred in the population. This preparatory stage ensures that new individuals entering the population will observe and learn from linguistic behaviour which has common elements. These common elements will be tend to be picked out by [+constructor, +ic-preserver] individuals, who will consequently share some subset of the system of meaning-signal mappings with other [+constructor, +ic-preserver] individuals, thereby reducing the time to cultural convergence.

6.4.6 Summary

Learning biases which lead to the evolution of communicatively optimal, compositional language can evolve through natural selection acting on genetic transmission. However, this only occurs when the partial-credit measurement of communicative accuracy is used. This measurement ensures that individuals with appropriate learning biases receive a fitness payoff fairly rapidly.

Even with the partial-credit measurement scheme, learning biases for compositionality only emerge 50% of the time. This is due to a dependence on genetic drift — as seen in Chapter 4, successful runs exhibit a drift-selection-drift pattern, where the initial period of drift is required to provide appropriate genotypes in sufficient numbers for cultural evolution to get underway. In the structured model of communication, there is less dependence on this initial period of drift, due to the partial-credit evaluation function and the rapid cultural convergence observed in the population. However, the fact remains that this initial period of drift is necessary — there is no immediate advantage in being biased to acquire a communicatively optimal, compositional language in a population which has no established communication system.

6.5 Discussion

What can the results from the Evolutionary Iterated Learning Model tell us about the evolution of language acquisition biases in humans? Much of the discussion relating to the evolution of vocabulary acquisition biases given in Chapter 4 remains pertinent. We can either draw a positive conclusion, and argue that these results show that human-like learning biases *can* evolve through natural selection, or we can take a negative position and argue that these results show that learning biases in humans must have arisen by (initially) non-adaptive mechanisms.

The negative conclusion remains the strongest one — despite the comparatively high level of success in the EILM using the partial-credit communicative accuracy evaluation, the point remains that the evolution of one-to-one biases requires an initial, fortuitous period of genetic drift. The natural conclusion to draw from this is that some mechanisms other than natural selection for communication must have provided appropriate genotypes in sufficient numbers to allow the cultural construction process to get underway. As discussed in Chapter 4, perhaps this one-to-one property was an incidental feature of some learning apparatus which evolved for some other purpose (i.e. the one-to-one bias is a spandrel). Alternatively, the appropriate learning bias may have evolved for some other function, then been pressed into service for communication.

These results do suggest an interesting alteration to the positive interpretation, however. Comparison of the EILM results in this chapter and those in Chapter 4 show that one-to-one learning biases are more likely to evolve when the communication system is potentially structured, given the partial credit fitness function. This suggests that an appropriate learning bias is more likely to evolve for the acquisition of a structured language than an unstructured language — the possibility of regularities in subparts of the meaning-signal mapping smoothes the fitness landscape and simplifies the evolutionary problem. In other words, if we accept the positive interpretation, the models in this Chapter and Chapter 4 show that evolution is more likely to go the whole hog and evolve a learning bias for language, rather than evolving a learning bias for unstructured vocabulary then later elaborating this bias.

6.6 Summary of the Chapter

In the first part of this Chapter I demonstrated that compositional language can emerge in populations of linguistic individuals through cultural processes, given a bottleneck on cultural transmission, a structured environment and the appropriate learning bias. I then went on to show that this learning bias can evolve under natural selection for communication. However, the evolution of this bias is dependent on an initial period of genetic drift — there is no immediate advantage to individuals possessing the appropriate bias in a population with no established communication system. The results and discussion from Chapter 4 therefore pertain to the evolution of learning biases for structured communication — such biases may best be explained as a spandrel or exapted trait.