THE ROYAL SOCIETY

doi 10.1098/rspb.2000.1150

The mimetic transition: a simulation study of the evolution of learning by imitation

Paul G. Higgs

School of Biological Sciences, University of Manchester, Oxford Road, Manchester M13 9PT, UK (*paul.higgs@man.ac.uk*)

Culturally transmitted ideas or memes must have had a large effect on the survival and fecundity of early humans. Those with better techniques of obtaining food and making tools, clothing and shelters would have had a substantial advantage. It has been proposed that memes can explain why our species has an unusually large brain and high cognitive ability: the brain evolved because of selection for the ability to imitate. This article presents an evolutionary model of a population in which culturally transmitted memes can have both positive and negative effects on the fitness of individuals. It is found that genes for increased imitative ability are selectively favoured. The model predicts that imitative ability increases slowly until a mimetic transition occurs where memes become able to spread like an epidemic. At this point there is a dramatic increase in the imitative ability, the number of memes known per individual and the mean fitness of the population. Selection for increased imitative ability is able to overcome substantial selection against increased brain size in some cases.

Keywords: memes; imitation; cultural evolution; human evolution

1. INTRODUCTION

A meme is an element of culture which can be passed on by imitation (Dawkins 1976, 1993; Dennett 1995; Blackmore 1999; Fog 1999). Memes can be trivial (e.g. catchy tunes) or can have huge influence (e.g. religious faiths). Memes can be abstract (e.g. the principle of democracy) or down to earth (e.g. stone axes). Certain memes, such as new methods of crop cultivation and new techniques for making tools and weapons, would have had a significant effect on the ability of individuals in primitive societies to survive and bring up children. Thus, the acquisition of memes affects the fitness of individuals. Blackmore (1999) argued that memes provide an explanation for the unusually large brain size and intelligence of our species. Individuals with genes which confer a better ability to learn by imitation are more likely to acquire memes which increase their fitness and, hence, are more likely to pass on their genes. Thus, imitative ability increases due to natural selection. Large brain size and high cognitive abilities are therefore seen as necessary for successful imitation.

Much effort in the field of animal behaviour has gone into determining to what extent various animal species can imitate and into distinguishing true imitation from simpler phenomena such as stimulus enhancement (see the review by Byrne & Russon (1998)). Although particular examples of imitation have been identified in species such as the great apes, it seems clear that humans have an ability to imitate which is more accurate and more general than any other species. Donald (1991, 1993) called ape-like culture 'episodic' and supposed that, around five million years (Myr) ago, *Australopithecus* had an episodic culture. The first stone tools appeared with *Homo habilis* around 2.5 Myr ago. Approximately 1.5 Myr ago *Homo erectus* arose with a substantially larger brain and a more advanced 'mimetic' culture. The word mimetic emphasizes that cultural traits spread between individuals by imitation. Thus, somewhere between 2.5 and 1.5 Myr ago our ancestors passed through a mimetic transition where imitative spread of culture got under way.

Several authors have produced mathematical treatments of cultural evolution (Cavalli-Sforza & Feldman 1981; Lumsden & Wilson 1981; Boyd & Richerson 1985, 1995; Durham 1991; Laland 1992). It is generally assumed that the human population possesses a high imitative ability and, thus, that memes have the ability to spread. The focus of the theories is therefore to ask which memes will spread and consider the competition between memes with different properties. In contrast, the model presented here considers a situation in which the population has a very low imitative ability and memes tend to die out. The model shows that imitative ability is selectively favoured in these conditions and, thus, provides theoretical support for the 'big brain' argument of Blackmore (1999). It is shown that imitative ability increases gradually under the action of selection until a mimetic transition point is reached where memes have the ability to spread like an epidemic. The model predicts a sudden change at this transition point where the number of memes known per individual increases by orders of magnitude and the imitative ability of the population also rises dramatically.

One reason why the evolution of a large brain size is puzzling from an evolutionary point of view is that brain tissue has a high metabolic energy requirement in comparison to most other organs of the body. This should provide a downward selective pressure on brain size in the absence of other selective effects. In addition, in humans the brain size has become large enough to lead to a significant number of problems at childbirth. These two factors imply that there are substantial costs associated with having a large brain (at least for present-day humans) and, hence, that there must also be relatively large selective benefits from the brain in order to counter these costs. The first part of this paper ignores these costs and demonstrates that there is a selective advantage to imitative learning in the absence of costs. In the second part of the paper, costs are introduced and it is shown that selection for imitative ability is able to counter costs of a substantial size.

Received 21 February 2000 *Accepted* 23 March 2000

2. A MODEL FOR THE EVOLUTION OF IMITATIVE ABILITY

The model studied here considers a population of individuals in which memes are constantly being invented and passed between individuals by imitation. There is a fixed population size of N individuals and generations are treated as non-overlapping. Each individual in the population has a biological fitness w and a cultural fitness v. The biological fitness determines the probability of reproduction and, hence, of passing genes to the next generation. The cultural fitness determines the probability of being imitated and passing memes to the next generation. It is a measure of status and the ability to influence others. This is equivalent to the Boyd & Richerson (1985) definition of cultural fitness. It is assumed that both the biological and cultural fitnesses of an individual are determined by the set of memes that he or she knows. Each meme (m) has a biological fitness effect $w_m = 1 + s_m$ and a cultural fitness effect $v_m = 1 + c_m$. These quantities represent the effects of the meme on the fitness of the individual. For each meme the values of s_m and c_m are determined randomly from normal distributions with mean zero and standard deviations σ_s and σ_c. Extreme cases with w_m (or v_m) less than zero are treated as w_m (or v_m) = 0. The biological and cultural fitnesses of an individual are $w = \Pi_m w_m$ and $v = \Pi_m v_m$, where the products are over the set of memes known by the individual. A 'naive' individual knowing no memes has $w = v = 1$. For simplicity it is assumed that the fitness effects of different memes are multiplicative, i.e. a meme always has the same relative effect, independent of which other memes are known.

The ability to learn by imitation is determined using the simplest possible, one-locus, diploid genetic system. Variant alleles are present at this locus, each of which specifies a learning ability l. An individual has a learning ability which is the average l-value of the two genes possessed. Offspring inherit one or other of the l genes at random from each parent. Mutation occurs with probability u per gene per generation. When a mutation occurs the new allele has a learning value $l_{new} = l_{old} + \delta l$, where δl is chosen from a normal distribution with mean $\overline{\delta l}$ and standard deviation σ_l. Alleles with negative l-values are not permitted and are set to $l = 0$.

Whilst genes are inherited from biological parents only, memes are inherited from both biological parents and cultural parents. Cultural parents are defined as non-related individuals in the parental generation from whom the children copy memes. Each new individual created in the population has two biological parents randomly chosen from the previous generation with a probability proportional to their biological fitness (i.e. roulette wheel selection) and K cultural parents randomly chosen from the previous generation with a probability proportional to their cultural fitness. The $K + 2$ adult models for any one child must all be different. However, an individual in the parental generation may be chosen any number of times as a parent or cultural parent of different children.

An individual with learning ability l has a probability $L(l) = 1 - \exp(-l)$ of successfully learning a meme from any one of its $K + 2$ models. This probability is obtained by the following argument. Suppose that the individual has n occasions of seeing the meme demonstrated by the adult model and that there is a small probability al of successful imitation on each occasion. The probability of acquiring the meme is the probability of successful imitation on at least one of the n occasions: $L(l) = 1 - (1 - al)^n$. Assuming $n >> 1$, $al << 1$ and $na = 1$ gives $L(l) = 1 - \exp(-l)$. Setting na to 1 merely sets a scale to the l-values; hence, it can be done with no loss of generality. We already have parameters $\overline{\delta l}$ and σ_l governing the size of mutations in l; therefore, we do not require the extra parameters n and a. Note that $L(0) = 0$ (i.e. an individual with no learning ability has zero probability of acquiring a meme) and $L(l)$ tends to 1 when l is very large (i.e. an individual with extremely high learning ability can always acquire any meme).

In order to determine the set of memes known by an individual in a new generation, each of the memes of the first parent is considered in turn and acquired with probability $L(l)$. The memes of the second parent and the cultural parents are then considered in turn and are acquired with probability $L(l)$ if they have not already been learned. An individual cannot learn the same meme twice. However, he or she may have more than one attempt at learning one meme if more than one of the adult models knows the meme. After attempting to learn memes by imitation, an individual also has a probability p_{inv} of inventing a new meme. New memes have values of s and c determined randomly and independently of previous memes.

3. RESULTS: THE MIMETIC TRANSITION

Run 1 of the imitator model considers the case where the cultural and biological fitness effects of memes are equal ($c = s$ for all memes). This situation arises if we interpret biological fitness as determining the probability of survival to adult reproductive age and if cultural parents are chosen randomly from among viable adults. This is the simplest situation because there can be no conflict between genes and memes. The simulation begins with a population of moderately innovative individuals ($p_{inv} = 0.1$) with very little ability to learn by imitation ($l = 0.01$). The values of the other parameters are given in the notes to table 1. Figure 1 shows the evolution of the mean learning ability $\bar{l}$, the mean biological fitness $\bar{w}$, and the mean number of memes known by an individual $\bar{m}$. Initially, $\bar{m} = p_{inv}$ because the only memes known by an individual are those he or she has invented personally. Newly invented memes are equally likely to have positive and negative fitness effects and, hence, $\bar{w} = 1$. In the initial part of the simulation there is a general steady increase in $\bar{l}$ indicating that selection is favouring individuals with a higher learning ability. For most of the simulation $\bar{m}$ is only slightly larger than p_{inv}, i.e. most individuals fail to imitate memes. The fate of most new memes (even those with an advantageous effect on fitness) is to die out in a few generations. Whilst l is low, memes cannot spread through the population because the average number of offspring (either biological or cultural) who inherit a meme from a single adult is less than one.

A sudden change happens after *ca.* 3200 generations, as shown in figure 1. At this point the mean learning ability reaches a high enough value for memes to be able to

Table 1. *The transition times for runs 1–8*

(The simulations begin with all learning alleles having an identical value of $l = 0.01$ and consider populations with $p_{\rm inv} = 0.1$ and $K = 2$. In run 1, $\mathcal{N} = 500$, $u = 0.001$, $\overline{\delta l} = 0$, $\sigma_l = 0.05$ and $\sigma_s = 0.2$. In subsequent runs most parameters are as in run 1 and parameters which differ from run 1 are given in the notes column. The transition times are given as means plus or minus the standard error in the mean (estimated from 100 independent runs for each set of parameters).)

run	transition time $\overline{T}$	notes
1	5270 ± 320	standard parameter set
2	1850 ± 100	σ_s increased to 0.4
3	19 600 ± 1000	u decreased to 0.0002
4	8590 ± 490	$\mathcal{N} = 100$
5	3050 ± 100	$\mathcal{N} = 2000$
6	7770 ± 460	c and s independent, $\sigma_s = 0.2$ and $\sigma_c = 0.2$
7	21 300 ± 1500	$c = s$ and $\sigma_s = 0.2$
8	17 000 ± 1000	$\overline{\delta l} = -0.25$

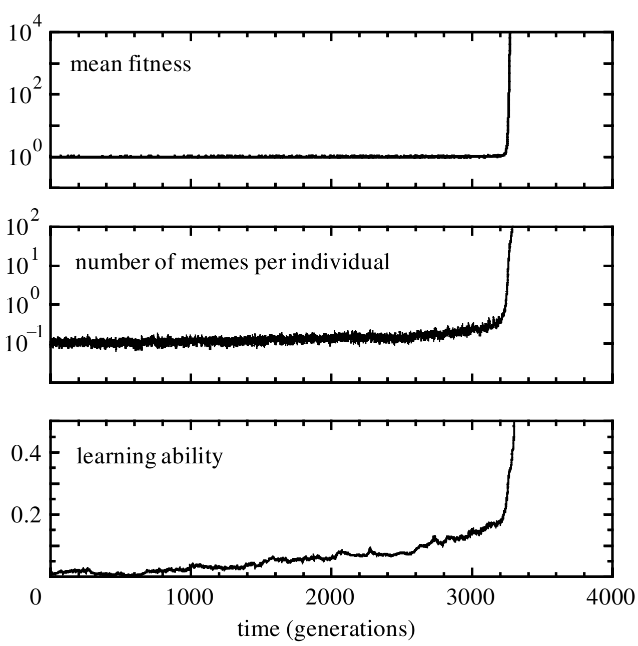

Figure 1. The results of a typical simulation of the imitator model with the parameters as in run 1.

spread like an epidemic. From this point onwards most new memes will rise to a high frequency. The number of memes known per individual suddenly shoots up and would continue to rise indefinitely if the simulation were not halted. We define the mimetic transition as the point at which $\overline{m}$ first rises above unity. In independent simulations with the same parameters the time taken to the mimetic transition varied considerably, but the transition always occurred. The mean time taken $\overline{T}$ is given in table 1. The minimum value of l necessary for an epidemic-like spread of memes is calculated in Appendix A.

Table 1 also lists the mean transition times for various other sets of parameters for comparison with run 1. The range of fitness effects of memes in run 2 is $\sigma_s = 0.4$,

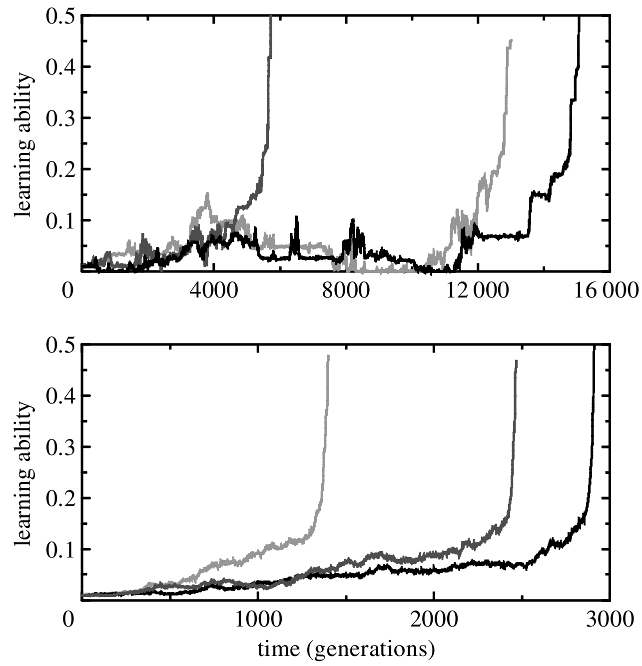

Figure 2. The mean learning ability as a function of time for three typical simulations with $\mathcal{N} = 100$ (top) and $\mathcal{N} = 2000$ (bottom). These correspond to runs 4 and 5 in table 1.

which is twice as much as in run 1. This speeds up the transition as would be expected, since the selective advantage to higher l-values is now larger. The transition time also depends on the mutation rate u. Run 3 differs from run 1 in that u is reduced by a factor of five and $\overline{T}$ is substantially increased. The transition time also depends on the population size $\mathcal{N}$. In run 1, $\mathcal{N}$ was 500. In runs 4 and 5, $\mathcal{N}$ is set to 100 and 2000, respectively. The time to the transition decreases as the population size increases. This is because high-l alleles are advantageous. Standard population genetics theory shows that the rate of fixation of advantageous alleles increases with $\mathcal{N}$, whereas for deleterious alleles the rate of fixation decreases with $\mathcal{N}$. The faster the fixation of high-l alleles in the population, the sooner $\overline{l}$ reaches the threshold necessary for meme spread.

Figure 2 shows the mean learning ability in three independent runs with $\mathcal{N} = 100$ and $\mathcal{N} = 2000$. These curves show that there is considerable variability in the time taken to the transition when the simulation is repeated with the same parameters. This is because there are several stochastic elements to the model. The times at which new learning alleles arise by mutation are random, as are the l-values of the new alleles. The selective effect of each newly invented meme is also random. In small populations random sampling of parents at each generation leads to substantial genetic drift in the gene frequencies. This is seen in the curves for $\mathcal{N} = 100$, where $\overline{l}$ drifts up and down before the transition takes effect. For $\mathcal{N} = 2000$ genetic drift is less significant and a steady upward trend is seen in learning ability. As $\overline{l}$ increases in each simulation run, the number of memes in the population that are available to be imitated also increases. If an individual has a slightly higher l than average, the number of extra memes that the individual learns compared to average will be proportional to the number of memes in the

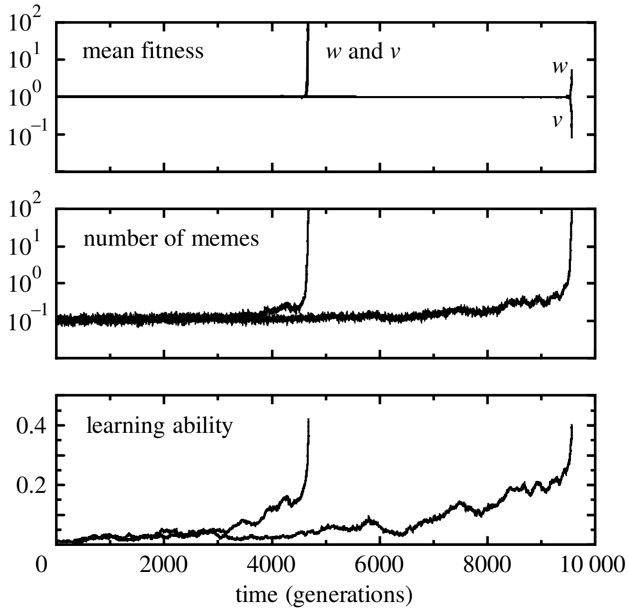

Figure 3. The results of one typical simulation with c and s independent (transition occurring close to time 4600) and one typical simulation with $c = -s$ (transition occurring close to time 9600). The standard parameter set is used except that $N = 2000$.

population. This means that the strength of selection for increasing l also increases as the mean l increases. At the end of the runs, the number of memes is large and this leads to a strong directional selection on l which dominates genetic drift, even for $N = 100$.

It is worth considering in more detail exactly how selection operates to increase l. If an individual with higher than average l happens to acquire a positive meme (with $s > 0$), that individual will have more offspring than average. Hence, high-l genes will increase in frequency. The offspring of this individual also have a higher than average chance of acquiring the positive meme from their parent and, hence, of having a higher than average fitness in the next generation. Of course, high-l individuals are also more likely to acquire negative memes (with $s < 0$) as well as positive memes. High-l individuals have a wide spread of fitnesses, whilst low-l individuals all have fitnesses close to one. What matters is that the highest fitness individuals will be among those which have the highest l; hence, there is selection to increase l.

This argument assumes that there is a range of memes of both positive and negative fitness effects and that individuals imitate memes with equal probability, regardless of their effects. In reality, individuals must have some ability for discriminating between positive and negative memes and are more likely to imitate memes which they perceive to be beneficial. Simulations were also carried out in which individuals were allowed to discriminate positive from negative memes and had a lower probability of imitating negative memes than positive ones. The result was as expected (therefore details will not be given): the greater the discriminatory ability, the less likely it is that high-l genes become associated with memes of negative s, hence the greater the selective advantage to imitation and the shorter the time to the mimetic transition. The occurrence of the mimetic transition in the simple simulations with no discrimination is therefore quite a strong result. Even if memes are equally likely to have a positive or negative effect, and even if individuals imitate blindly, there is still a selective advantage to imitation.

The criterion for meme spread given in Appendix A shows that memes of positive effect spread more easily that memes of negative effect (even with blind imitation). Hence, the majority of memes in the population after the transition will be positive ones. Therefore, the mean fitness $\overline{w}$ increases exponentially with the number of memes known (note the logarithmic scale in figure 1). It should be remembered that these fitnesses are measured relative to a naive individual with no memes. Individuals after the transition are at a tremendous advantage relative to those before. However, the competition between individuals is much stronger after the transition. It is assumed that the population size is controlled by limited resources. The number of offspring of an individual is thus proportional to his or her fitness relative to the mean of the current generation and not relative to the naive individuals of previous generations. The new high-fitness memes are necessary for surviving after the transition, whereas before the transition it was not necessary to have memes because no-one else had them either.

4. FACTORS ACTING AGAINST THE MIMETIC TRANSITION

So far the biological and cultural fitness effects of the memes have been set equal. However, there is no reason why this should necessarily be the case. A meme with a positive effect c on the cultural fitness is (by definition) one which makes the individual more likely to be imitated. This could be by conferring status, power or communicative ability on the individual or simply by making him or her more fashionable. These characteristics are not necessarily correlated with the biological fitness. Since the transmission of memes is not necessarily reliant on the successful transmission of the genes of their host, there is a potential coevolutionary conflict between the memes and the genes. Dawkins (1993) called memes 'viruses of the mind' for this reason. Appendix A shows that memes with a positive cultural fitness effect could spread relatively easily, even if they had a negative biological fitness effect.

Simulations were therefore performed in which both c and s were determined randomly and independently of each other (run 6 in table 1). The principal result was that the mimetic transition still occurred, but it took significantly longer than for run 1 since the cultural and biological fitnesses were now sometimes in conflict with each other. In order to maximize this conflict the simulations in run 7 were performed, in which s was determined as usual and c was set equal to $-s$, i.e. all memes with a positive biological fitness effect had a negative cultural fitness effect. The mean transition time in run 7 was substantially longer than the other cases, but once again the transition still occurred.

Figure 3 shows typical simulations with independent c and s and with $c = -$s. The curves for learning ability and the number of memes per individual appear as in the previous cases. The top graph shows both the mean

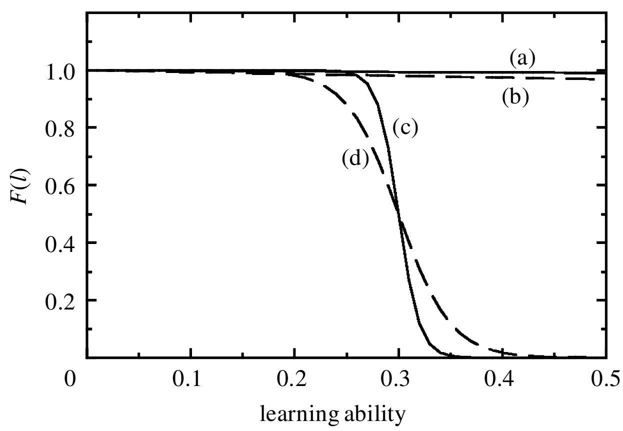

Figure 4. The fitness functions $F(l)$ for various types of cost functions. Curve (a), linear costs with $l_0 = 50$. Curve (*b*), linear costs with $l_0 = 15$. Curve (c), exponential costs with $l_0 = 0.3$ and $\beta = 100$. Curve (d), exponential costs with $l_0 = 0.3$ and $\beta = 40$.

biological and cultural fitnesses. When c and s are independent both $\overline{w}$ and $\overline{v}$ shoot up at the transition. These quantities remain almost equal to one another and, therefore, only one curve can be distinguished on the graph. This means that the memes which spread effectively after the transition are those which have both s and c positive. In contrast, for the example with $c = -s$, $\overline{w}$ increases at the transition but $\overline{v}$ decreases. This shows that the biological fitness effect of the memes is driving the mimetic transition despite the counteraction of the cultural fitness effect. Memes with positive s are transmitted vertically to biological offspring even though the oblique transmission of these memes occurs at a reduced rate. The result is to reduce but not eliminate the selective advantage to increased imitative ability. This is reminiscent of the 'leash principle' (Lumsden & Wilson 1981; Dennett 1995; Blackmore 1999) which argues that cultural evolution is kept on the leash by biological evolution. A more apt metaphor for the situation here is that the memes are dragging the genes along by the leash. The memes cannot proceed if the genes do not follow because it is the high-l genes which allow the memes to spread. However, it is selection on the memes which drives the genetic evolution towards a high learning ability.

Another factor acting to prevent the mimetic transition is mutational bias. In runs 1–7, the mutations were such that l was as likely to increase as to decrease. We might expect that there are more ways that learning ability could be decreased by mutations than ways in which it could be increased. Thus, there should be biased mutation towards a low l. In run 6 the mean change in l per mutation was set to $\overline{\delta l} = -0.025$. With this value, only *ca.* 30% of mutations increase l, so if there were no selection for an increased learning ability the mean l of the population would remain very close to zero (i.e. use it or lose it). The mimetic transition still occurred in run 8, although with a substantially longer waiting time than run 1. This demonstrates that there is a positive selection towards increased l which is sufficient to overcome substantial downward mutational bias.

Table 2. *The effects of the costs of a large brain size on the mean time to the mimetic transition*

(The simulations have the same parameter values as in run 1, with the addition of a cost function with the parameters given above. One hundred replicate runs were performed with each cost function. The simulations were halted after 100 000 generations if no transition occurred. Where a mean and standard error are quoted this implies that a transition occurred in every one of the 100 replicates. In the case of linear costs with $l_0 = 15$, a transition occurred before 100 000 generations in 76 out of 100 replicates. In the cases denoted as 'no transition' there was no transition in any of the 100 replicates.)

linear cost		exponential cost		
l_0	transition time $\overline{T}$	l_0	β	transition time $\overline{T}$
∞	5270 ± 320	0.50	100	5140 ± 320
500	5630 ± 390	0.30	100	5150 ± 330
200	6350 ± 420	0.26	100	6810 ± 340
50	7820 ± 560	0.24	100	10 310 ± 540
30	13 670 ± 970	0.22	100	no transition
15	*ca.* 100 000	0.30	50	6950 ± 370
10	no transition	0.30	40	19 260 ± 1000
		0.30	30	no transition

The main factor acting against the mimetic transition has been left until last. This is the selective cost against increased brain size because of increased energy expenditure and possibly because a large head size causes increased risk at childbirth. To model this effect it is supposed that the relative cost of the brain for an individual with learning ability l is $1 + f(l)$, where the level of energy expenditure for an individual with no learning ability is 1 and the additional cost due to increased brain size is $f(l)$. The biological fitnesses of individuals are assumed to vary inversely with this cost, i.e. we set

$$w = F(l) \prod_m w_m, \tag{1}$$

where $F(l) = 1/(1 + f(l))$.

Initially, suppose the additional cost is linear in l, i.e. $f(l) = l/l_0$, where l_0 is a parameter determining the magnitude of the cost. The results are shown in table 2. An infinite l_0 corresponds to no cost. As l_0 is reduced the costs are increased and the time to the transition is increased. For l_0 below around 15 no transition occurred. Thus, small linear costs can be overcome by the selective advantage of the memes, whereas larger costs prevent the transition from occurring.

It may be more reasonable to suppose that the costs are nonlinear with l because the brain size may not be not directly proportional to l and because the costs may not be directly proportional to the brain size. In particular, we might expect that the costs are very small for a small l and then increase rapidly if l becomes too large. For example, difficulties at childbirth would only become apparent if the head were larger than a threshold size. It could also be argued that brains are required for many things other than imitation and that a low imitative ability could be obtained with an existing brain at no additional metabolic cost. The additional costs of imitation would only set in

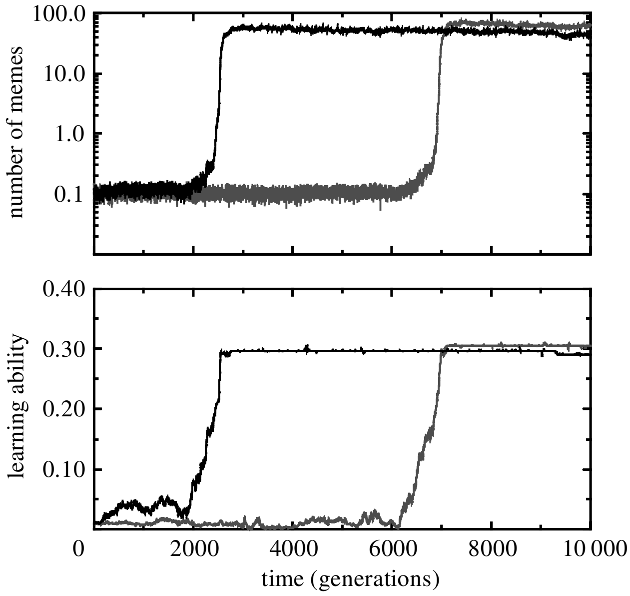

Figure 5. The results of two typical simulations with exponential costs ($l_0 = 0.3$ and $\beta = 100$). These costs do not prevent the transition, but they prevent the growth of l beyond a certain point.

once the brain size became significantly larger than that which was required for other purposes. Simulations were therefore performed with an exponential cost function $f(l) = \exp(\beta(l - l_0))$. For both exponential and linear functions $f(l) = 1$ when $l = l_0$. The effect of these costs on the fitness are determined by the function $F(l)$ which multiplies w and v. Examples of this function are shown in figure 4.

The results in table 2 show that, when $\beta = 100$ and $l_0 > 0.3$, there is no significant difference in the mean transition time from the case with no cost. A decreasing l_0 corresponds to shifting curve (c) in figure 4 to the left. This causes a significant increase in the transition time and eventually prevents the transition altogether. Similarly, if l_0 remains at 0.3 but β is decreased, the costs begin to take effect at smaller l (as in curve (d) in figure 4). This again causes the mean transition time to increase and eventually prevents the transition.

The results with the nonlinear cost function are much more encouraging than those with the linear cost function. Curve (b) in figure 4 is the maximum strength of the linear costs for which the transition was observed ($l_0 = 15$), i.e. even rather small linear costs are sufficient to prevent the transition. On the other hand, when the costs are nonlinear, selection in favour of imitation can overcome substantial costs. The results of two runs of the simulation with $\beta = 100$ and $l_0 = 0.3$ are shown in figure 5. The initial parts of the runs are as before, but growth in the learning ability is stopped at a point after the transition where the costs become sufficient to prevent further increase. This results in a corresponding halt to the increase in the number of memes known. The value of l reached in this example is *ca.* 0.3, which means that $F(l)$ is *ca.* 0.5. Thus, there is a twofold disadvantage to the large brain in this example and selection in favour of imitation is sufficient to counter this. Simulations with a power law cost function $f(l) = (l/l_0)^k$ for exponents k greater than unity were also carried out and produced results similar to those with the exponential cost function.

The strength of the selection against increasing l is proportional to the slope of the cost function $\mathrm{d}F/\mathrm{d}l$. This slope is approximately constant for linear costs, i.e. the costs take full effect even when l is small. When l is low and memes are rare the selective advantage to increased l is small and, hence, selection for imitation can only beat the costs if the slope is small. For linear costs, if the costs beat selection for imitation initially then l is driven down to an optimum of zero. If imitation beats the costs initially, then it will gain an even greater advantage later in the simulation. Linear costs are therefore not sufficient for preventing the rapid rise of l once the transition has occurred. However, for nonlinear costs, if the slope is initially small then imitation will win initially and the transition will occur. A rapid rise in the costs after the transition can then check the growth in brain size, as happens in figure 5. This situation may well apply to the human species, where the costs of the large brain have become significant, but clearly the advantages are sufficient for outweighing these costs. The sensitivity of the transition to the shape of the cost function may also go some way to explaining why the mimetic transition has not happened in other species.

5. CONCLUSIONS

This model gives theoretical support to the big brain argument given qualitatively by Blackmore (1999, pp. 74–81). It shows that imitation is selectively favoured under a range of circumstances and that selection for imitation can overcome downward mutational biases and the costs of a large brain size provided that these costs are not too large when the imitative ability is very low. The model makes the novel prediction that a mimetic transition will occur in which a sudden rise in imitative ability and the number of memes takes place. It seems likely that such a phenomenon took place during the evolution of our own species, although we recognize that there are alternative theories for the origin of culture which we cannot attempt to review here (see, for example, Gabora 1998).

The model assumes that memes can have both positive and negative effects on fitness and that individuals imitate memes blindly with no account being taken of their selective effect. This has been done so that the theory is conservative: even when memes are learned blindly there is still a selective advantage to imitation. In reality, individuals should have some ability to discriminate good from bad memes and to learn the good ones preferentially. This just makes the transition happen more easily and, therefore, does not affect the main conclusions of this paper. An increased ability to discriminate memes is presumably a selective advantage and, therefore, discriminative ability is likely to increase alongside imitative ability. In addition, the ability to create new memes will also be a selective advantage and this aspect of cognitive ability is also likely to increase during evolution. If p_{inv} were allowed to increase during a simulation, this would also make the transition happen more easily. However, if p_{inv} were allowed to evolve whilst l were fixed at a low value, there would be no transition. Creative individuals would evolve which were adept at individual problem solving, but they would have no ability to pass on their discoveries or to benefit from the discoveries of others. Following this argument, creative ability and discriminatory ability are therefore both secondary to imitative ability.

This model treats all memes as different from one another, i.e. it does not consider that some memes are alternative versions of the same thing. After the mimetic transition there will clearly be many alternative competing memes for any aspect of human culture. Theories of cultural evolution after the transition therefore need to account for competition between alternative memes and for the evolution of memes that arise by the modification of old ones. Some of these effects are considered in other theories and simulations (Cavalli-Sforza & Feldman 1981; Lumsden & Wilson 1981; Boyd & Richerson 1985; Gabora 1995). However, these aspects are not important before culture gets under way. When the average number of memes known is less than unity, any individual is unlikely to know two memes for the same thing. When most memes die out after a few generations there is no time for them to be modified to alternatives. Thus, without unnecessary complexities, this model clearly shows that imitative ability is at an advantage when that ability is low and that this creates the situation in which meme spread and cultural evolution can get under way. More complex models will be needed to consider meme evolution after the transition.

APPENDIX A

If the variation in learning ability between individuals is small, then the condition for meme spread can be calculated as a function of the average learning ability of the population. Consider a meme with biological and cultural effects s and c. Let the fraction of individuals who know this meme at time t be $x(t)$ before selection. The frequency of the meme in the biological parents after selection is

$$x_B(t) = x(t)\frac{(1+s)}{\overline{w}} = \frac{x(t)(1+s)}{1+sx(t)}. \tag{A1}$$

The frequency of the meme in the cultural parents after cultural selection is

$$x_C(t) = x(t)\frac{(1+c)}{\overline{v}} = \frac{x(t)(1+c)}{1+cx(t)}. \tag{A2}$$

Similar equations to these were used by Cavalli-Sforza & Feldman (1981). The probability that an individual learns the meme from any one of its biological (or cultural) parents is $x_{\rm B}(t)L(l)$ (or $x_{\rm C}(t)L(l)$). This is just the probability that the parent knows the meme multiplied by the learning probability. The frequency $x(t+1)$ of the meme at the next generation before selection is given by the probability that an individual learns the meme from at least one parent:

$$x(t+1) = 1-\left(1-\frac{x(t)(1+s)L(l)}{1+sx(t)}\right)^2\left(1-\frac{x(t)(1+c)L(l)}{1+cx(t)}\right)^K. \tag{A3}$$

A meme will increase in frequency when rare if $x(t+1) > x(t)$ when $x(t) << 1$. This gives the criterion

$$(2(1+s) + K(1+c))L(l) > 1. \tag{A4}$$

When condition (A4) is true the meme reaches a non-zero stationary frequency which is the solution to equation (A3) with $x(t+1) = x(t)$. When condition (A4) is not true the only solution to equation (A3) is $x = 0$, i.e. the meme dies out. Condition (A4) shows that memes spread more easily when K is higher, that memes with positive s and c spread more easily than those with negative fitness effects and that memes with negative effects can still spread if $L(l)$ is large enough. In run 1, a neutral meme with $s = c = 0$ can spread if $l > 0.288$. A positive meme with $s = c = 0.4$ (two standard deviations above the mean) can spread if $l > 0.197$. This seems consistent with the simulation in the figure, in which $\overline{l} = 0.22$ when $\overline{m}$ reaches 1.

REFERENCES

Blackmore, S. 1999 *The meme machine.* Oxford University Press.

Boyd, R. & Richerson, P. J. 1985 *Culture and the evolutionary process.* University of Chicago Press.

Boyd, R. & Richerson, P. J. 1995 Why does culture increase human adaptability? *Ethol. Sociobiol.* **16**, 125–143.

Byrne, R. W. & Russon, A. E. 1998 Learning by imitation: a hierarchical approach. *Behav. Brain Sci.* **21**, 667–721.

Cavalli-Sforza, L. L. & Feldman, M. W. 1981 *Cultural transmission and evolution: a quantitative approach.* Princeton University Press.

Dawkins, R. 1976 *The selfish gene.* Oxford University Press.

Dawkins, R. 1993 Viruses of the mind. In *Dennett and his critics: demystifying the mind* (ed. B. Dahlbom), pp. 13–27. Oxford, UK: Blackwell.

Dennett, D. 1995 *Darwin's dangerous idea.* London: Allen Lane.

Donald, M. 1991 *Origins of the modern mind.* Harvard University Press.

Donald, M. 1993 Precis of *Origins of the modern mind* and associated commentary. *Behav. Brain Sci.* **16**, 737–791.

Durham, W. H. 1991 *Coevolution. Genes, culture, and human diversity.* Stanford, CA: Stanford University Press.

Fog, A. 1999 *Cultural selection.* Dordrecht, The Netherlands: Kluwer Academic Publishers.

Gabora, L. 1995 Meme and variations: a computer simulation of cultural evolution. In *1993 Lectures in complex systems*, vol. VI (ed. L. Nadel & D. L. Stein). California: Addison Wesley.

Gabora, L. 1998 Autocatalytic closure in a cognitive system: a tentative scenario for the origin of culture. *Psycoloquy* **9**, 67.

Laland, K. N. 1992 A theoretical investigation of the role of social transmission in evolution. *Ethol. Sociobiol.* **13**, 87–113.

Lumsden, C. J. & Wilson, E. O. 1981 *Genes, mind, and culture.* Harvard University Press.