How Culture Can Guide Evolution: An Inquiry into Gene/Meme Enhancement and Opposition

To appear in the *Journal of Adaptive Behavior.*

Michael L. Best

Media Laboratory
Massachusetts Institute of Technology
E15-312
20 Ames Street
Cambridge, MA 02139
USA
mikeb@media.mit.edu
www.media.mit.edu/~mikeb

+1 (617) 253-4662
FAX +1 (617) 258-6264

How Culture Can Guide Evolution

Abstract

We study the relationship between genetic evolution, learning, and culture. We start with the simulation environment of Hinton and Nowlan in which individual learning was shown to guide genetic evolution towards a difficult adaptive goal. We then consider, in lieu of individual learning, culture in the form of social learning by imitation. Our results demonstrate that when genes and culture cooperate, or enhance one another, culture too is able to guide genetic evolution towards an adaptive goal. Further, we show that social learning is superior to individual learning insofar as it with genetic evolution converges more quickly to the goal. However, the social learning algorithm results in slower genetic assimilation of adaptive alleles than with individual learning. It is as if, we argue, the adaptive values are stored in the culture rather than in the genes. Finally, we consider what happens when culture and genes pursue diametrically opposed goals. Here we show that culture, in the form of social learning, is no real match when opposed to genetic evolution with individual learning. In fact, only the most herculean of social learning algorithms is able to keep a neutralizing toe-hold against the slow plodding force of genetic evolution. Finally, our results suggest that in both cases, opposition and enhancement, transmission forces such as the ratio of teacher to learner are central to the success of social learning.

Keywords: Cultural evolution, gene/culture co-evolution, social learning, genetic algorithm.

1.0 Introduction

Happy is she who is able to employ a multiplicity of methods to acquire adaptive behaviors. To a first approximation, living organisms on Earth have but a single method, namely evolution of genetic material. Lucky are those organisms whose adaptive mechanisms are augmented by lifetime individual learning. Lucky and few are those entities who add various modes of cultural learning to the adaptive mix.

For the social animals, genes, individual learning, and culture mix together creating a rich and complex adaptive agency. It is the character of this mix, the ways in which these three adaptive mechanisms interact and co-evolve, which is a central question in the development of biological theories of culture. In fact, these very questions have been the ground-zero for heated debates about the evolution of culture, in particular human culture (Caplan 1978).

So far the best work exploring the relationship between genes, memes (Dawkins 1976), and learning has been empirical studies of animal models (e.g. Huffman's (1984) consideration of adaptively neutral stone handling in Japanese Macaques or Galef's (1994) studies of the social transmission of food preference among Norwegian rats), primarily theoretical models (e.g. Lumsden &Wilson 1981; Cavalli-Sforza & Feldman 1981; Boyd & Richerson 1985) and work employing empirical studies of humans as support of particular theoretical models (e.g. Cosmides & Tooby (1989) and Tooby & Cosmides' (1989) study of social exchange, or Durham's (1991) coevolutionary model supported by empirical studies of incest taboos, lactose intolerance, and so forth). While simulation models of social agents provide a powerful method to explore the gene/meme mix, so far these techniques go underutilized (in Section 3.0 we overview work of this sort).

In this paper we explore a simple simulation environment of evolving social agents. In our simulation, the interactions of lifetime individual or social learning with trans-generational genetic evolution is studied. By *social learning* we mean a learning algorithm where the interaction is between conspecifics rather than the environment. Our simulation takes as its starting point the method of G.E. Hinton and J. Nowlan (1987). In their "modern classic," Hinton and Nowlan augmented the *genetic algorithm* (GA) (Holland 1975) affording each agent the capacity for individual learning. That is, each agent has additional adaptive plasticity as a result of non-social interaction with the environment. In this paper we add to Hinton and Nowlan's duo of adaptive mechanisms, culture in the form of a *simple* social learning algorithm, namely *imitation*. In our view it is the transparent simplicity of Hinton and Nowlan's original work which is critical to its impact; we have strived for such simplicity even while adding mechanism.

We study both gene/meme *enhancement* and *opposition* (Durham 1991). That is, we study cases where genetic evolution and social learning search the same adaptive landscape and cases where their adaptive surfaces are diametrically opposed. In results reported here we find that in the case of gene/meme enhancement social learning can, in the absence of individual learning, guide genetic evolution towards a goal. Thus, we argue that culture can guide genetic evolution in much the same way that Hinton and Nowlan demonstrated for individual learning. Moreover, within our simulation social learning is superior to individual learning -- the cultural algorithm produces more rapid convergence to the adaptive goal than individual learning. But in the case of gene/meme opposition we find that social learning is hardly a match for genetic evolution. Only the most herculean of social learning algorithms allows culture to keep a neutralizing toe-hold against the plodding march of genetic evolution.

1.1 When genes and memes interact

We are studying the relationships between genetic evolution and learning, but by this we do *not* mean how a youngster's grade school algebra class is relevant to their Darwinian success. (That would be overly simplistic sociobiology.) In humans high-level learning, including that done in an algebra class, generally does not meaningfully interact with genetic evolution or correspond to particular genetically coded traits. But in other animals (and for cases in humans) lifetime learning, and in particular that which has a social component, often does have some interactions with genetic traits and evolution. Some examples of this includes learning within the problem domains of spatial orientation and memory, foraging, food preferences, mate selection, predator and risk avoidance, and so forth. It is these behaviors that do indeed often have strong relationships with genetic evolution and that motivate our simulation.

Within these problem domains learning, for some animals, is a rather simple cognitive behavior. Indeed, by "learning" we simply mean some form of lifetime phenotypic plasticity due to non-trivial associations between stimuli (social and non-social) and behavior (Dukas 1998). As will be seen in Section 2.0, learning need not imply a process that is particularly clever, studious, or thoughtful.

1.2 Overview of the paper

In the next section we will review the Hinton and Nowlan GA model and their original results. In Section 3.0 we describe some other related work in simulation of social agents. In the following section we describe the social learning algorithm and report results of cultural enhancement in the absence of individual learning. In Section 5.0 we develop analytically the affordances of social learning in comparison to individual learning. In the next section we consider social learning in opposition to genetic evolution. We end with our conclusions.

2.0 Hinton and Nowlan's Simulation

Hinton and Nowlan used a simple GA simulation to demonstrate the effects of individual learning on genetic evolution (Hinton & Nowlan 1987). In particular their simulation gave support to the Baldwin effect. This general phenomena, described first by J. Mark Baldwin (1896a) and C. Lloyd Morgan (1896), argues that learned traits or "adaptations secured by conscious agency.... imitation, gregarious influences, maternal instruction, the lessons of pleasure and pain, and of experience generally" (Baldwin 1896a), while not heritable (*contra* Lamark (1914)) could still have effect on genetic evolution. In Hinton and Nowlan's reading, this meant that individual lifetime learning (the "lessons of pleasure and pain") could alter the fitness landscape searched by genetic evolution.

In Hinton and Nowlan's simulation each agent is represented by a string of twenty characters, this is the agent's genotype. Each position, or *locus*, along the string can take on one of three values, namely, "0", "1", or "?". The values, "0" and "1", represent genotypic codings, *alleles*, which are described as connecting relay switches for a neural net with 20 potential connections. In the simulation a population of 1000 agents were initially seeded randomly, each locus having a 25% chance of being assigned a “0”, 25% chance a “1”, and 50% chance a “?”.

The “?” value is what allows for lifetime behavioral plasticity amongst the agents. Each agent enjoys during its lifetime a series of learning rounds. During each round the agent runs a simple individual learning algorithm which for each “?” in the genotype assigns either a “0” or “1” as the expressed value. Note that the values assigned to the question mark positions during a learning round are never written back into the genotype nor passed directly on to progeny.

The individual learning algorithm itself is quite simple. For each learning round each agent queries a random number generator once for each question mark. With 50% probability a “0” is assigned to the locus and with 50% probability a “1” is assigned. Thus the agent performs during the learning rounds a random walk of a subspace of the boolean hypercubic landscape (as described by the loci marked by question marks). This search clearly is random with respect to the resultant fitness; it is not biased towards adaptive behaviors but has equal probability to assign a zero or one. For each round of genetic evolution every agent engages in 1000 learning rounds during which each of its “?’s” are subject to individual learning.

The reproductive rounds proceed using, more or less, the traditional genetic algorithm. The fitness of each agent is evaluated according to an external function and agents are selected (with replacement) with probability proportional to that fitness. The selected agents are paired at random and the crossover operator is applied to their genotypes. The mutation operator is not used.

In their simulation the external fitness landscape is unusual; indeed it describes a rather cruel space. The function is low for all dichotomies of the genotype except for the string of all ones for which it is high. Thus it can be described as an impulse function. This fitness landscape is highly correlated and provides a difficult space indeed for genetic evolution to search -- a needle-in-a-haystack search. During the selection stage of the GA an agent that is all “1’s” will be assigned a high fitness and an agent with *any* “0’s” will receive a low value.

But, given the general sketch above, how exactly do we evaluate each agents fitness during the selection phase of the GA? The agent's fitness is a function of its successes during the learning rounds as well as its initial genetic makeup. If during the set of learning rounds the agent happened upon the winning dichotomy, namely the string of all ones, then the agent notices that it has found the co-adaptive set of behaviors and will stop its learning. After all 1000 learning rounds the fitness for an agent is given by $\left(1 + \frac{19n}{1000}\right)$ where n is the number of learning rounds required to find the all one string. Thus should the genotype contain any zeros, or should the agent never discover the all one dichotomy during learning, the fitness will be twenty times worse than that of an agent who has the all-one string at the first learning round ($1 + \frac{19(1)}{1000} = 1.019$ versus $1 + \frac{19(1000)}{1000} = 20$).

The principal measurement of the state of the simulation is the proportion of loci occupied by "0's", "1's", or "?'s". Should genetic evolution make any progress one would expect the percentage of "0's" to drop towards zero and the percentage of "1's" to rise. Since the fitness function defined above penalizes agents who have to spend time learning the all-one configuration one expects there to be selective pressure towards the genetic assimilation of "1's" and the removal of "?'s" from the genotype. We have re-implemented the Hinton and Nowlan algorithm and summarize their main result in Figure 2.1. Here we see the percentage of "0's" quickly drop to zero and the percentage of "?'s" slowly decline. Note that our results show a slightly more pronounced, though still quite gradual, drop off in the number of question marks than do the original results reported by Hinton and Nowlan; we suspect this is due to subtle implementation differences.

The important finding of Hinton and Nowlan's original work is revealed when one compares the results of the simulation in the presence of individual learning against the exact algorithm run without any learning. Without learning there is no significant variation in the frequency of "0's", "1's", and "?'s" from their original random assignment, even after runs of 1,000's of generations. Thus the conclusion is that in the absence of individual learning genetic evolution is lost; it will never find and hold onto the needle-in-the-haystack. But when coupled with the adaptive learning algorithm, genetic evolution is able to find and keep hold of the all-one configuration, bringing along the entire population of agents. The learning mechanism not only increases the *likelihood* of discovering the needle due to searching. It also allows agents which are *close* to the needle to receive some fitness (and therefore reproductive) benefit since they are likely to find the needle during their learning trials. This in effect smooths-out the spiked landscape creating a virtual slope through the fitness landscape; and it is this slope which the genetic algorithm is then able to hillclimb. Thus individual learning "guides" evolution to the adaptive goal.

3.0 Other Related Work

We now review previous work which has studied and enhanced Hinton and Nowlan's model. We also describe some related work using simulation to explore culture, social learning, and communication.

The work most closely related to ours is that of R. Belew (1990). Belew proposed a very useful metaphor for the relationship of genetic evolution, individual learning, and culture. He argues that these three modalities each act as an articulation point of a robot arm. At the end of this arm are fingers which grope about some adaptive landscape. Thus these three mechanisms act as components to a single adaptive system. Belew notes that this metaphor "helps to capture the way in which the plastic system ('joints') of [genetic] evolution, culture and learning mediate between the more rigid constraints ('bones') of extant genetic material, species, and artifacts." (Belew 1990: 45).

While the bulk of Belew's paper is devoted to an analysis of the original Hinton and Nowlan model he does spend some time discussing a culturally-enriched version of the model. He proposes that "winning" agents (those that find the all-one string) are able to vertically transmit, parent to offspring, a cultural advantage (CA). This cultural advantage comes in the form of a learning bias; for instance, those offspring enjoying an CA might during a learning round have a 60% chance of selecting a "1" rather than a 50% chance. Belew finds that the vertical transmission of cultural advantage causes the population to converge more quickly to the adaptive goal. Further, he notes that the population maintains a larger number of question marks than under the original simulation. Apparently, with the addition of culture there is less pressure to adopt the optimal genotype. (Our results confirm these findings for horizontal, or peer-to-peer, transmission and social learning.)

Belew also studied oblique transmission, calling it the "broadcast" mode. Here a winner from the previous generation is able to transmit a cultural advantage to a randomly selected subset of the current generation. Belew discovered that the broadcast mode was effective relative to simple vertical transmission when the size of the attendant population is large.

While Belew's results are impressive we feel he does not address the most compelling questions, namely Can culture, without the influence of individual learning, guide evolution? In other words, can cultural learning guide genetic evolution in the same way that Hinton and Nowlan show individual learning does. And exactly what are the additional affordances of culture relative to individual learning? Can culture oppose genetic evolution and learning? Below we attempt to answer these questions.

Further, we find Belew's notion of a cultural advantage a bit difficult. Belew seems to propose that a parent (or other model) is able to transmit to offspring a "learning strategy" that biases the agents individual learning towards the winning results. This seems quite an elaborate cultural mechanism since strategies, in contrast to simply memes or behaviors, are hard to distill (even when you use one) and even more difficult to transmit. We find it more likely that each model simply transmits an actual set of behaviors, which are imitated with more or less success. This is the model of cultural transmission we propose.

A number of other researchers have studied and enhanced the Hinton and Nowlan model. This includes work by G. Mayley (1996) who explored the cost of learning with varying landscapes; instead of considering solely the impulse function Mayley employed Kauffman's (1993) NK-model to produce landscapes of varied ruggedness. Inman Harvey (1993) studied the persistence of the large frequency of question marks within the evolving population of agents. Through a probabilistic diffusion model he argues that the persistent question marks are due to genetic drift in the system.

A number of researchers have used simulation environments to study the evolution of language and the role of communication. The bulk of this growing area of research has been studying how cooperative communication systems might *develop*, and how they might *sustain* themselves, all within an evolutionary context. In particular, recent work has studied the evolution of shared meaning and lexicon formation (e.g. Werner & Dyer 1992; MacLennan 1992; Ackley & Littman 1994; Steels 1996; Noble & Cliff 1996; Saunders & Pollack 1996). Further, some researchers have examined cultural transmission in general within simulation. For instance Gabora (1995) has built a simulation system and explored a collection of memetic transmission mechanisms.

The above work examines through simulation the evolution *of* communication and social agency. We, in contrast, assume the pre-existence of some communication mechanism and are studying the co-evolutionary role of this social mode, genetic evolution, and individual learning. In this way we take an approach similar to Hinton and Nowlan who assumed the presence of individual learning and explored its relationship to genetic evolution.

4.0 Gene/Meme Enhancement with Social Learning

When genetic and cultural (memetic) forces select for identical traits they are potentially in an enhancing relationship. Specifically, when genetic and cultural selection act to favor the same trait and the resultant combined selective force is stronger than would be the case just due to genetic influences we have an instance of gene/meme enhancement (Durham 1991). Within the Hinton and Nowlan model this means that culture should somehow select for the adaptive spike (the all-one configuration) and do so in some way that is above-and-beyond the selective force already present thanks to genetic evolution.

We wish to answer the following two questions regarding gene/meme enhancement:

- First, can a cultural algorithm operating in the absence of individual learning "guide" evolution towards the adaptive goal just as does individual learning? In other words, can we recreate Hinton and Nowlan's important result using a cultural rather than individual learning mechanism?
- Second, what exactly are the affordances of the cultural algorithm when compared to the individual learning algorithm?

We will tackle the first question in this section and address the latter in Section 5.0.

The immediate matter at hand is to devise some cultural algorithm which is simple yet powerful. We have already seen Belew's cultural advantage algorithm, which made use of vertical and oblique transmission modes. We employ instead a *social learning* algorithm which is *imitation* based. Rather than conferring a cultural advantage, agents are selected as social models and their behaviors are imitated.

In our current study we have only examined *horizontal transmission*, that is transmission between any two (usually unrelated) agents within a single generation (Cavalli-Sforza & Feldman 1981). The relative use and importance of these various modes of transmission is, in our view, not well understood. For instance, compare Cavalli-Sforza, Feldman, Chen, & Dornbusch's (1982) findings amongst students of Stanford (vertical transmission wins the day) versus Huffman's (1984) observations of Macaques of Japan (horizontal transmission rules). But at least amongst nonhuman social animals, it has been argued that most social learning is highly horizontal (Galef 1988; Laland, Richerson, & Boyd 1996). The results reported here only make use of horizontal transmission; work currently in preparation studies in simulation the relative advantages of these various transmission modes.

It has been argued that imitation is at the very root of human-level culture (e.g. Piaget 1962; Dawkins 1976; Galef 1992; Meltzoff 1996; Heyes 1996). Therefore, an imitation centered social learning algorithm is an appropriate first mechanism to probe the fundamental relationships between genes, individual learning, and culture. (Note, however, that we do not here presume to enter into the debate around the patterns and prevalence of different forms of social learning, e.g. true imitation, local enhancement, stimulus enhancement, and so forth (see Galef (1988) for a review). We find that imitation is the most useful way to *think* of our social learning algorithm. However, these other modes of learning are also modeled by our social learning algorithm.)

We will consider the social learning algorithm used *in lieu* of the individual learning algorithm. Between each round of genetic evolution we now have 1000 rounds of social learning rather than 1000 rounds of individual learning. Again, our main question is Can culture "guide" genetic evolution?

Our social learning algorithm is quite similar to the genetic algorithm. This is due to our belief (Pocklington & Best 1997; Best 1998; and argued by many others), that cultural evolution is fundamentally the same as genetic evolution (they are two different examples of one common abstract mechanism). The social learning algorithm proceeds as follows: Each agent is evaluated by an external cultural fitness function. Social models are chosen based on *fitness-proportionate selection* (Mitchell 1996). A learner is selected randomly for each model. For each question-mark in the learner's genotype, that is for each locus which enjoys lifetime plasticity, the model's value is imitated with some probability and subject to some mutation rate. Unlike the genetic algorithm, the social learning algorithm does not employ sex and crossover. Of course values acquired during the learning rounds are not passed backed into the genome nor inherited by offspring.

In order to study gene/meme *enhancement* the cultural fitness landscape should be identical to the genetic landscape; to wit, it is small everywhere but for the all-one configuration for which it is big. Fitness-proportionate selection means that the expected number of times that an agent is selected as a model is proportional to its relative fitness. We have chosen a fitness span between the winning all-one configuration and any of the losing configurations of 500. In other words, a winner will be 500 times more fit then a loser. This is considerably larger than the span of 20 used by Hinton and Nowlan, though the end results are not particularly sensitive to this. Unlike the normal genetic algorithm, however, we do not employ a selection mechanism which replaces the entire population with each round. In our simulation, the expected number of times an agent is selected as a model during a particular learning round is given by,

$$E(selected) = \begin{cases} \alpha, & \text{winner} \\ \dfrac{\alpha}{500}, & \text{loser.} \end{cases} \tag{4.1}$$

Thus, for α=1, *E(win)* = 1 and *E(lose)* = .002. With this α and a population of 1000 agents we expect only two non-winning agents (1000*.002 = 2) to act as models and therefore, given no winners, only two agents will do any learning at all (since each time a model is selected, with replacement, it only gets one learner). Clearly the amount of learning due to the cultural algorithm, given this value for α, is quite small.

Thus, the value α parametrizes the *transmission force* (Cavalli-Sforza & Feldman 1981) of the social learning algorithm. It controls the expected total number of models (and therefore learners). An outcome of this is variation in the model to learner *ratio*. For instance with α=100, each winner is expected to model for 100 different learners giving a model to learner ratio of 1/100. A loser, in contrast, is expected to be selected with a 1 in 5 chance (100/500) to act as a model for a single learner. Thus, one-fifth of the agents without the all-one configuration will have a 1/1 model to learner ratio and the other non-winners do not act as models at all.

For our simulation we used a mutation rate of zero and each social learner is able to imitate perfectly the model at all of its question-mark loci; in other words, for each question mark in the learner's genotype the learner is able to copy the value from the model perfectly. This mutation mechanism (or lack thereof) was chosen for its similarity to the genetic evolution stage, which also does not make use of mutation. But, more importantly, we wanted to eliminate all randomness in the social learning stage beyond that due to the selection and pairing mechanism. This is important since any randomness (for instance a mutation rate or a probabilistic transmission rate) would "appear" as something akin to the original individual learning -- that is, some sort of random shuffling of the expressed values. We want a social learning algorithm which, given some learner and its model, deterministically copies "0's" and "1's" from the model to the learner and therefore keeps itself quite distinct from the stochastic individual learning method. Under social learning there is no random assignments of "0's" or "1's" but just the transmission of "0's" and "1's" from model to learner.

We experimented with a variety of values for the transmission force, α. These experiments are summarized in Figures 4.1 and 4.2 which show results from single runs of the simulation. For very low values of α, as shown in Figure 4.1, the social learning algorithm is not sufficient to allow genetic evolution to gain a handle on the adaptive goal. This is true even when the simulation is run for considerably more than 50 generations (much as Hinton and Nowlan discovered when they turned off individual learning). Recall that for these small values of α we expect only a few agents to engage in any learning at all during each round. Compare that to the original individual learning algorithm in which, for each learning round, each and every agent engages in learning.

In Figure 4.2 we show three simulation rounds for $\alpha = 10$, 20, and 100. For the value $\alpha = 10$ we see a phase transition, with a rapid drop-off in frequency of "0" alleles, around the 45th generation. Given a transmission force of $\alpha = 10$ we expect only (1000 * .02) = 20 agents to engage in learning during each round in the absence of winners. But as soon as we have winners the number of social learners will quickly escalate since each winner is expected to model for ten agents. Each winner by definition has the all-one configuration so they are only giving out the "right" answer. Thus they are likely to create more winners through their effective modeling, creating a snow-ball effect.

Already for $\alpha = 20$ the phase transition point is very close to the point for the individual learning algorithm, as seen in Figure 2.1. And in Figure 4.3 we directly compare the results from the social learning algorithm with a transmission force $\alpha = 100$, and the individual learning algorithm. Clearly, the social learning mechanism outperforms individual learning insofar as it converges on the adaptive goal more quickly.

Note that, similar to the results of Belew (1990), the social learning algorithm does not eliminate question-marks at the rate that individual learning does. Why is that? The adaptive cost of learning, paid at the point of genetic evolution, is the same for both social and individual learning. In both cases we use the same fitness function for the genetic evolution stage which charges for each learning round required to discover the winning configuration (as described in Section 2.0). Under individual learning the number of rounds required to find the winning configuration is a probabilistic function in the number of question marks for that genotype. And so the forces at hand are genetic drift (c.f. (Harvey 1993)) versus this selection pressure towards assimilation. These same forces are in place under social learning as well. However, the cost to learn the winning dichotomy is considerably less under social learning given the presence of a fair number of winners. Since the winning agents are expected to be selected as a social model for a large number of agents it is likely that non-winning agents will quickly learn from a winning model the correct value for all of their plastic loci. Social diffusion rates are expected to follow the logistic (Cavalli-Sforza & Feldman 1981), and thus within a small number of learning rounds agents with even a large number of question marks will learn the winning configuration. The assimilation pressure is then not sufficient to squash out the question-marks since the correct configuration is so quickly learned. Another way to think of this is that the agents store the correct configuration within the culture rather than in their genes, and quickly adopt the adaptive behavior via social learning.

We have seen that the social learning algorithm is more effective than the individual learning algorithm though it produces genotypes with more question marks. More importantly, though, we have shown that a cultural algorithm, namely social learning via imitation, is sufficient to “guide” genetic evolution towards a needle-in-the-haystack adaptive goal. This power of social learning is a function of the transmission force, α. Only with a suitably large transmission force is social learning able to guide genetic evolution. Recall that these high values for α correspond to both overall increased levels of learning and increased model to learner ratios.

4.1 Social Learning used in Conjunction with Individual Learning

The above analysis considered the social learning algorithm used in lieu of the individual learning algorithm. This was done in order to study whether a cultural algorithm can guide evolution. However, we suspect there are no animals which rely exclusively on social learning and make no use of individual learning, though the opposite is clearly the case for non-social learners.

We did perform a number of simulations of the social learning algorithm used in conjunction with individual learning. Our results, however, where not significantly different than those discussed in Section 4.0. In other words, when the value of α was quite low then the social learning scheme was not able to improve on the results observed given individual learning working in isolation. And, when the value of α was high, the results where similar to those reported in Section 4.0 -- speedier convergence than under individual learning operating alone.

5.0 Gene/Meme Enhancement -- Individual versus Social Learning

We will now analytically compare the social learning algorithm with the individual learning algorithm. The results from Section 4.0 demonstrate that the social learning algorithm does not require individual learning in order to make progress towards the adaptive goal (and similarly the original Hinton and Nowlan result demonstrates the reverse, that individual learning can guide evolution without social mechanisms). Further, our simulation demonstrated that the social learning algorithm converges faster to the adaptive goal than does the individual learning algorithm.

The question of interest, then, is Exactly how superior is the social learning algorithm compared to the individual learning algorithm? Rather than explore this via simulation we will try to make progress through an analytical treatment.

Following the terminology of Belew (1990), a *winner* is an agent with all-ones either within the genotype or learned. A *potential winner* is an agent with no zeros in its genotype and perhaps some question marks. Given an initial population we wish to determine what is the probability that all potential winners will become winners under the individual versus the social learning algorithm. We then ask How would we have to bias the individual learning algorithm toward assigning “1’s” so that these two algorithms are equivalent?

Allow W to represent a winner, PW a potential winner, and q the number of question marks in the agent’s genotype. Then the probability that a potential winner becomes a winner given q question marks and the individual learning algorithm is given by,

$$Pr_{IL}[W|q] = \left(\frac{1}{2}\right)^q. \tag{5.1}$$

And following Harvey (1993) the probability, given a potential winner (no zeros), to have q question marks is the expansion of $\left(\frac{2}{3}+\frac{1}{3}\right)^{20-1}$, or,

$$Pr[q \text{ question marks}] = \binom{20}{q}\left(\frac{2}{3}\right)^q\left(\frac{1}{3}\right)^{20-q}. \tag{5.2}$$

This falls out of the initial probabilities: a 50% chance for “?” and 25% chance for “1”. Thus, if we sum across all possible values for q, then the probability that the individual learning algorithm will produce a winner given a potential winner in one round is,

$$Pr_{IL}[W] = \sum_{q=0}^{20}\binom{20}{1}\left(\frac{2}{3}\right)^q\left(\frac{1}{3}\right)^{20-1}\left(\frac{1}{2}\right)^q. \tag{5.3}$$

To determine a similar probability under the social learning algorithm is a much more difficult proposition and is served by a few simplifying assumptions. Namely, we wish to fix the number of potential winners. The probability of a potential winner in the initial population, that is an agent with no zeros, is given by $\left(1-\frac{1}{4}\right)^{20}$. Therefore, the expected number of potential winners in a population of 1000 randomly seeded agents is,

$$E[PW] = 1000\left(1-\frac{1}{4}\right)^{20} \approx 3.$$

We will consider an initial population with exactly three potential winners. The other simplifying assumption we make is that for any single learning round no more than one potential winner will become a winner. This is a conservative simplification; it will serve to underestimate the strength of social learning.

We wish to compute the probability that a potential winner becomes a winner given the social learning algorithm. But this is a function of how many winners currently exist -- since each winner has an opportunity to model for one or many potential winner. Thus the probability that exactly one of the remaining potential winners becomes a winner through learning from a winner in round *n+1,* given k winners in learning round n, $1 \le k < 3$ is,

$$Pr_{SL,\,winner}[W_{n+1}^{k+1} | W_n^k] = 1 - \left(1 - \frac{\alpha}{1000}\right)^k - Pr_{SL}[\text{more than 1 winner in } n+1] \; . \quad (5.4)$$

This is because there is an α-in-1000 chance of learning from a winner model given the selection method of Equation 4.1. Here $Pr_{SL}[\text{more than 1 winner in } n+1]$ is simply the probability that more than one potential winner becomes a winner in a single round which we rule out, simplifying the analysis below. Similarly, the probability that the social learning algorithm produces a winner through modeling by a non-winner is given by,

$$Pr_{SL,\,nonwinner}[W_{n+1}^{k+1} | W_n^k] = \left(1 - \left(1 - \left(\frac{1}{1000}\right)\left(\frac{\alpha}{500}\right)\right)^{999-k}\right) \sum_{q=0}^{20} \binom{20}{1} \left(\frac{2}{3}\right)^q \left(\frac{1}{3}\right)^{20-1} \left(\frac{1}{2}\right)^q . \quad (5.5)$$

This is the joint probability that a non-winner acts as a model (from Equation 4.1), the probability that the potential winner is selected as the learner (1-in-1000), and the probability that the loci modeled result in an all-one configuration (from Equation 5.3). This latter probability we compute as though each of the q question-marks are assigned a “0” or “1” with equal probability. This makes sense since this is the initial distribution from which the values for the model were taken. Thus, the probability that a potential winner becomes a winner given no current winners is given by Equation 5.5 (which accounts for non-winner models). And from Equation 5.4 and 5.5 we can derive the probability of a potential winner becoming a winner under social learning as,

$$Pr_{SL}[W_{n+1}^{k+1} | W_n^k] = 1 - (1 - Pr_{SL,\,winner}[W_{n+1}^{k+1} | W_n^k])(1 - Pr_{SL,nonwinner}[W_{n+1}^{k+1} | W_n^k]) . \quad (5.6)$$

What is the probability that all three potential winners will be winners at learning round n using social learning? This probability is computed by summing the probabilities of all the independent pathways in which the social learning algorithm could have produced a winner. This is given by,

$$Pr_{SL}[W_n^3] = \sum_{\substack{i,j,k \\ i<j<k\le n}} \left(\begin{array}{l} \left(\prod_{a=0}^{i-1} 1 - Pr_{SL}[W_a^1 | W_{a-1}^0]\right) Pr_{SL}[W_i^1 | W_{i-1}^0] \; * \\ \left(\prod_{b=i+1}^{j-1} 1 - Pr_{SL}[W_b^2 | W_{b-1}^1]\right) Pr_{SL}[W_j^2 | W_{j-1}^1] \; * \\ \left(\prod_{c=j+1}^{k-1} 1 - Pr_{SL}[W_c^3 | W_{c-1}^2]\right) Pr_{SL}[W_k^3 | W_{k-1}^2] \end{array} \right). \qquad (5.7)$$

Computing the probability that three potential winners become winners in an initial population with $n = 1000$ and $\alpha = 100$, we find that $Pr_{SL}[W_{1000}^3] \approx 0.035$. In other words, the probability that three potential winners will become winners in 1000 rounds of social learning is roughly 0.035 for the moderate value of $\alpha = 100$.

We wish to compare this to the probability of the same outcome given the individual learning algorithm. This is easily computed since every agent's individual learning activities are independent. From Equation 5.3 it is easy to see that the probability in 1000 rounds for three potential winners to become winners is given by,

$$\begin{aligned} (Pr_{IL}[W_{1000}^3] &= (1 - ((1 - Pr_{IL}[W])^{1000}))^3) \\ &\approx 0.018. \end{aligned} \qquad (5.8)$$

Thus the probability doubles (.018 versus .035) when moving from individual to social learning.

Finally, we wish to find what individual learning bias would produce the same probability of converting three potential winners to winners in 1000 rounds as is given by the social learning algorithm. That is, how can we re-work Equation 5.8 to equal the probability for social learning, 0.035. This is the same as, within Belew's framework, looking for a cultural advantage which provides the same probability of success as our imitation based social learning algorithm. Working with Equation 5.3 and Equation 5.8 we are looking for an individual learning bias, β, that will satisfy the equality,

$$\left(1 - \left(1 - \sum_{q=0}^{20} \binom{20}{1} \left(\frac{2}{3}\right)^q \left(\frac{1}{3}\right)^{20-1} (\beta)^q \right)^{1000}\right)^3 = 0.035 \qquad (5.9)$$

Recall the learning bias given by Hinton and Nowlan, $\beta = .5$. In other words, there is an equal probability of "0's" and "1's". Solving Equation 5.9, we find $\beta \approx 0.514$ satisfies the equality. Thus if the individual learning algorithm selects "1's" with 51.4% probability and "0's" with 48.6% probability then the three potential winners would become winners with the same probability as under social learning.

6.0 Gene/Meme Opposition

Finally, we consider the case when culture and genes are in opposition. This occurs, according to Durham (1991), when cultural selection favors memes which are not the most advantageous in terms of genetic inclusive fitness. Under our simulation model this is easily implemented. We utilize the same simulation environment as described above except that we will simultaneously employ all three adaptive mechanisms -- genetic evolution, individual learning, and social learning. While genetic evolution continues to search for the all-one adaptive goal, social learning now has as its goal the all-zero configuration. Thus, social-learning has a hard needle-in-the-haystack function to search for; and it is diametrically opposed to the fitness function guiding genetic evolution. We continue to use the same fitness proportionate cultural selection mechanism as suggested by Equation 4.1. However, unlike the simulations of Section 4.0 we do not remove the individual learning mechanism; this is to ensure that genetic evolution always has a fighting chance of finding the solution even when very little social learning is occurring. The order we apply the learning mechanisms is clearly critical. For each learning round agents first employ the individual algorithm and then follow it with the social learning algorithm. Since the individual learning algorithm is so completely disruptive, with every "?" randomly reset, the reverse ordering of events would completely neutralize social learning. But in the absence of many winners it is not the case that social learning undoes totally the effects of individual learning.

Our results are demonstrated in Figures 6.1 and 6.2. In Figure 6.1 we see social learning in opposition to genetic evolution for three transmission forces, $\alpha = 200$, 400, and 600. By increasing α the model to learner ratio increases. This serves to disrupt genetic evolution requiring more generations before the transition point to the all-one configurations can occur. But, clearly, sooner or later genetic evolution wins out over the opposing goals of culture. In Figure 6.2 we directly compare a small and large value for α. For the small transmission force, $\alpha = 100$, no real difference is noted when weak social learning opposes genetic evolution; the results look like the original simulation shown in Figure 2.1. However with $\alpha = 900$, genetic evolution is not able to get a foothold and there is no motion away from the original frequencies of “0’s”, “1’s”, and “?’s”. This result is true for runs of 100’s of generations. One could argue that social learning is not able to win insofar as it is not able to move the frequency of alleles towards its goal of an all-zero configuration. However, a strong opposing culture is at least able to neutralize the effects of genetic evolution. It is something of a draw with neither culture nor genes able to push the frequency of alleles towards their opposing goals.

These results demonstrate the force of genetic evolution along with individual learning. With social learning pitted directly against it, genetic evolution is hardly affected except when opposed by very strong social learning. Only with the herculean social learning mechanism, where any social winner essentially broadcasts its all-zero configuration to the entire population, is social learning able to keep a toe-hold against the plodding march of genetic evolution.

The effect seems fairly easy to explain. When social learning is in this near-broadcast mode, and it happens upon the all-zero configuration early enough, than the individual learning stage never is given a chance to find the all-one configuration. Anytime individual learning does find an all-one configuration, and thus gives genetic evolution a chance, social learning quickly undoes all of this work by offering all-zero agents as models. But if social learning is not strong, than even a few all-one configurations can slip through and that is enough for genetic evolution to quickly gain full control.

It is worth noting just how unpleasant a situation the cultural algorithm is placed into. We have put social learning in direct opposition to both genetic evolution and individual learning, given it a hard adaptive problem, and have not offered culture any particular help beyond transmission forces. Under these circumstances we find that culture is able to keep a toe-hold by making use of a very strong transmission force. Some researchers have argued that cultural transmission is able to oppose genetic evolution with reasonable ease given other tools. Boyd and Richerson (1985) show that biased transmission, that is the preference for certain cultural variants, in the presence of relatively weak genetic selection, will support gene/meme opposition. Takahasi (1998) has described a model in which the coevoltuion of cultural variants along with preferences can place genes and memes in opposition to one another. And Durham (1991) lists a collection of potential reasons for genes and culture to work against each other including simply an agent's inability to acquire accurate and timely information when evaluating variant behaviors. Given this range of conditions that may support it, we should not be surprised at the frequency in which genes and memes are shown to be in opposition.

7.0 Conclusions

We have explored the relationship between genes, learning, and culture. By adding a social learning mechanism to the model of Hinton and Nowlan we have explored how genes, memes, and learning interact. Our results indicate that culture in the form of social learning, when enhancing genetic evolution, can guide evolution in the same way as individual learning. Moreover, our social learning algorithm is superior to the individual learning mechanism insofar as it converges on the adaptive goal more quickly. We discovered that roughly a 1.4% bias in the individual learning algorithm towards the adaptive goal is sufficient to make the social and individual learning mechanisms equivalent. However, the results under social learning are less optimal; the proportion of question marks within the population remains higher than under individual learning. It is as if the adaptive values are stored in the culture instead of in the genes. We also considered what happens when culture and genes pursue diametrically opposed goals. We showed that culture, in the form of social learning, is no real match when opposed to genetic evolution with individual learning. In fact, only the most herculean of social learning algorithms is able to keep a neutralizing toe-hold against the slow plodding force of genetic evolution.

The role of the transmission force, α, is shown to be central to the effectiveness of social learning within our simulation. A variety of forces have been proposed as influential to the differential transmission of cultural elements (Durham 1990). These include "cultural selection", where memes are directly selected for due to their cultural advantage; "natural selection", where the meme's effect on reproductive advantage translates back to replication advantage for the meme; and "transmission forces", with factors such as teacher-to-learner ratios. Our results demonstrate that the transmission forces parameterized by α do indeed play an important role in social learning.

We note that Baldwin himself considered social learning via imitation as central to his theory of an adaptive system operating at the ontogenetic level. Baldwin proposed "social heredity" as an additional adaptive system alongside genetic evolution: "two great kinds of influence, each in a sense hereditary; there is *natural heredity*... and there is '*social heredity*'" (Baldwin 1896b, italics his). And central to social heredity was imitation and other observational learning: "it is evident that when young creatures have these imitative, intelligent, or quasi-social tendencies to any extent, they are able to pick up *for themselves*, by imitation, instruction, experience generally, the functions which their parents and other creatures perform in their presence," (Baldwin 1896a, italics his). Baldwin did not, however, make a real distinction between the adaptive affordances of individual learning and the adaptive qualities of cultural-social systems. We believe this to be a mistake. Curiously, the current interpretations of the Baldwin effect seem to concentrate on individual learning (lessons from "pleasure and pain") and there is little consideration of the "social" elements of "social heredity". This work describes a reorientation to the central consideration of Baldwin.

Acknowledgments

Thanks to the anonymous reviewers for their helpful comments. We gratefully acknowledge the assistance and advise of Pattie Maes, Nelson Minar, Richard Pocklington, and Bernd Schoner. Thanks go to the sponsors of the MIT Media Lab for their financial support.

References

Ackley, D.H. & Littman, M.L. (1994). Altruism in the evolution of communication. In R.A. Books & P. Maes (Eds.), *Artificial Life IV* (pp. 40-48). Cambridge, MA: MIT Press.

Baldwin, J.M. (1896a). A new factor in evolution. *American Naturalist,* Vol 30, 441-451.

Baldwin, J.M. (1896b). Heredity and instinct (I). *Science,* March 20, 1896.

Best, M.L. (1998). An ecology of text: using text retrieval to study alife on the net. *Journal of Artificial Life,* Vol 3, No 4, 261-287.

Belew, R.K.(1990). Evolution, learning, and culture: Computational metaphors for adaptive algorithms. *Complex Systems,* Vol 4, 11-49.

Boyd, R. & Richerson, P.H. (1985). *Culture and the evolutionary process.* Chicago: University of Chicago Press.

Caplan, A.L. (Ed.) (1974). *The sociobiology debate: Readings on ethical and scientific issues.* New York: Harper & Row Publishers.

Cavalli-Sforza, L. & Feldman, M. (1981). *Cultural transmission and evolution: A quantitative approach.* Princeton, NJ: Princeton University Press.

Cavalli-Sforza, L., Feldman, M.W., Chen, K.H., Dornbusch, S.M. (1982). Theory and observation in cultural transmission. *Science,* Vol 218, 19-27.

Cosmides, L. & Tooby, J. (1989). Evolutionary psychology and the generation of culture, part II: Case study: A computational theory of social exchange. *Ethology and Sociobiology*, Vol 10, 51-97.

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

Dukas, R. (1998). Evolutionary ecology of learning. In R. Dukas (Ed.), *Cognitive ecology: The evolutionary ecology of information processing and decision making* (pp. 129-174). Chicago: The University of Chicago Press.

Durham, W.H. (1990). Advances in evolutionary culture theory. *Annual Review of Anthropology.* Vol 19, 187-210.

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

Gabora, L.M. (1995). Meme and variations: A computational model of cultural evolution. In L. Nadel & D. Stein (Eds.), *1993 Lectures in Complex Systems* (pp. 471-485). Redwood City, CA: Addison-Wesley.

Galef, B.G., Jr. (1988). Imitation in animals: History, definition, and interpretation of data from the psychological laboratory. In T.R. Zentall & B.G. Galef, Jr. (Eds.), *Social learning: Psychological and biological perspectives* (pp. 3-28). New Jersey: Lawrence-Erlbaum Associates.

Galef, B.G., Jr. (1992). The question of animal culture. *Human Nature,* Vol 3, 157-178.

Galef, B.G., Jr. (1994). Olfactory communications about foods among rats: A review of recent findings. In B.G. Galef, Jr., M. Mainardi, & P. Valsecchi (Eds.), *Behavioral aspects of feeding: Basic and applied research in mammals* (pp. 83-102). Chur, Switzerland: Harwood Academic Publishers.

Harvey, I. (1993). The puzzle of the persistent question marks: A case study of genetic drift. In S. Forrest (Ed.), *Proceedings of the fifth international conference on genetic algorithms* (pp. 15-22). San Mateo, CA: Morgan Kaufmann.

Heyes, C.M. (1996). Genuine imitation. In C.M. Heyes & B.G. Galef, Jr. (Eds.), *Social learning in animals: The roots of culture.* San Diego, CA: Academic Press.

Hinton, G.E. & Nowlan, S.J. (1987). How learning can guide evolution. *Complex Systems,* Vol 1, 495-502.

Holland, J.H. (1975). *Adaptation in natural and artificial systems.* Ann Arbor, MI: University of Michigan Press.

Huffman, M.A. (1984). Stone-play of *Macaca fuscata* in Arashiyama B troop: Transmission of non-adaptive behavior. *Journal of Human Evolution,* Vol 13, 725-735.

Kauffman, S. (1993). *The origins of order: Self-organization and selection in evolution.* Oxford: Oxford University Press.

Laland, K.N., Kumm, J, & Feldman, M.W. (1995). Gene-culture coevolutionary theory: A test case. *Current Anthropology,* Vol 36, No 1, 131-156.

Laland, K.N., Richerson, P.J., & Boyd, R. (1996). Developing a theory of animal social learning. In C.M. Heyes & B.G. Galef Jr., *Social learning in animals: The roots of culture.* San Diego, CA: Academic Press.

Lamarck, J.B. (1914). *Zoological Philosophy.* London: Mcmillan.

Lloyd Morgan, C. (1896). On modification and variation. *Science,* Vol 4, 733-740.

Lumsden, C. & Wilson, E.O. (1981). *Genes, mind, and culture.* Cambridge, MA: Harvard University Press.

MacLennan, B. (1991). Synthetic ethology: An approach to the study of communication. In C.G. Langton, C. Taylor, H.D. Farmer, & S. Rasmussen (Eds.). *Artificial life II: Santa Fe Institute studies in the sciences of complexity, volume X.* Redwood City, CA: Addison-Wesley.

Mayley, G. (1996). The evolutionary cost of learning. In P. Maes, M.J. Mataric, J. Meyer, J. Pollack, & S. Wilson (Eds.), *From animals to animats 4: Proceedings of the fourth international conference on simulation of adaptive behavior* (pp. 458-467). Cambridge, MA: MIT Press.

Meltzoff, A.N. (1996). The human infant as imitative generalist: A 20-year progress report on infant imitation with implications for comparative psychology. In B.G. Galef & C.M. Heyes (Eds.), *Social Learning in Animals: The Roots of Culture* (pp. 347-370). New York: Academic Press.

Mitchell, M. (1996). *An introduction to genetic algorithms.* Cambridge, MA: MIT Press.

Noble, J. & Cliff, D. (1996). On simulating the evolution of communication. In P. Maes, M.J. Mataric, J. Meyer, J. Pollack, & S. Wilson (Eds.), *From animals to animats 4: Proceedings of the fourth international conference on simulation of adaptive behavior* (pp. 608-617). Cambridge, MA: MIT Press.

Piaget, J. (1962). *Play, dreams and imitation in childhood.* New York: Norton.

Pocklington, R. & Best, M.L. (1997). Cultural evolution and units of selection in replicating text. *Journal of Theoretical Biology,* Vol 188, No 1, 79-87.

Saunders, G.M. & Pollack, J.B. (1996). The evolution of communication schemes over continuous channels. In P. Maes, M.J. Mataric, J. Meyer, J. Pollack, & S. Wilson (Eds.), *From animals to animats 4: Proceedings of the fourth international conference on simulation of adaptive behavior* (pp. 580-589). Cambridge, MA: MIT Press.

Steels, L. (1996). Emergent adaptive lexicons. In P. Maes, M.J. Mataric, J. Meyer, J. Pollack, & S. Wilson (Eds.), *From animals to animats 4: Proceedings of the fourth international conference on simulation of adaptive behavior* (pp. 562-567). Cambridge, MA: MIT Press.

Takahasi, K. (1998). Evolution of transmission bias in cultural inheritance. *Journal of Theoretical Biology,* Vol 190, 147-159.

Tooby, J. & Cosmides, L. (1989). Evolutionary psychology and the generation of culture, part I: Theoretical considerations. *Ehthology and Sociobiology*, Vol 10, 29-49.

Werner, G.M. & Dyer, M.G. (1991). Evolution of communication in artificial organisms. In C.G. Langton, C. Taylor, H.D. Farmer, & S. Rasmussen (Eds.). *Artificial life II: Santa Fe Institute studies in the sciences of complexity, volume X*. Redwood City, CA: Addison-Wesley.

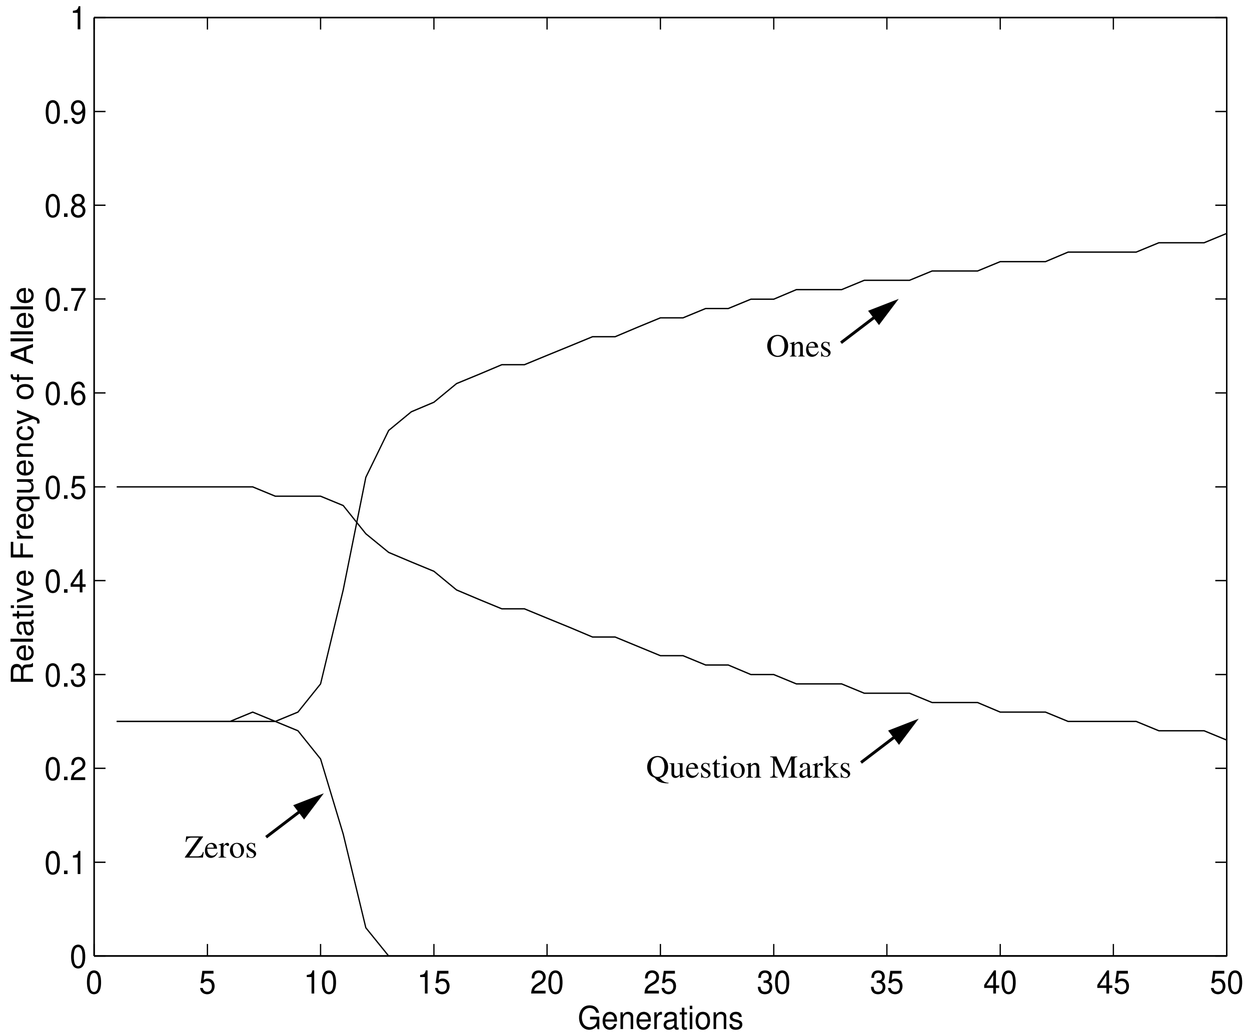

Figure 2.1: Our re-implementation of the original Hinton and Nowlan experiment. Quickly the frequency of “0’s” drops to zero and the “1’s” take off. Note the slow tapering off of the number of question marks. In Hinton and Nowlan’s original result the frequency of question marks remains fairly steady.The population size is 1000 agents.

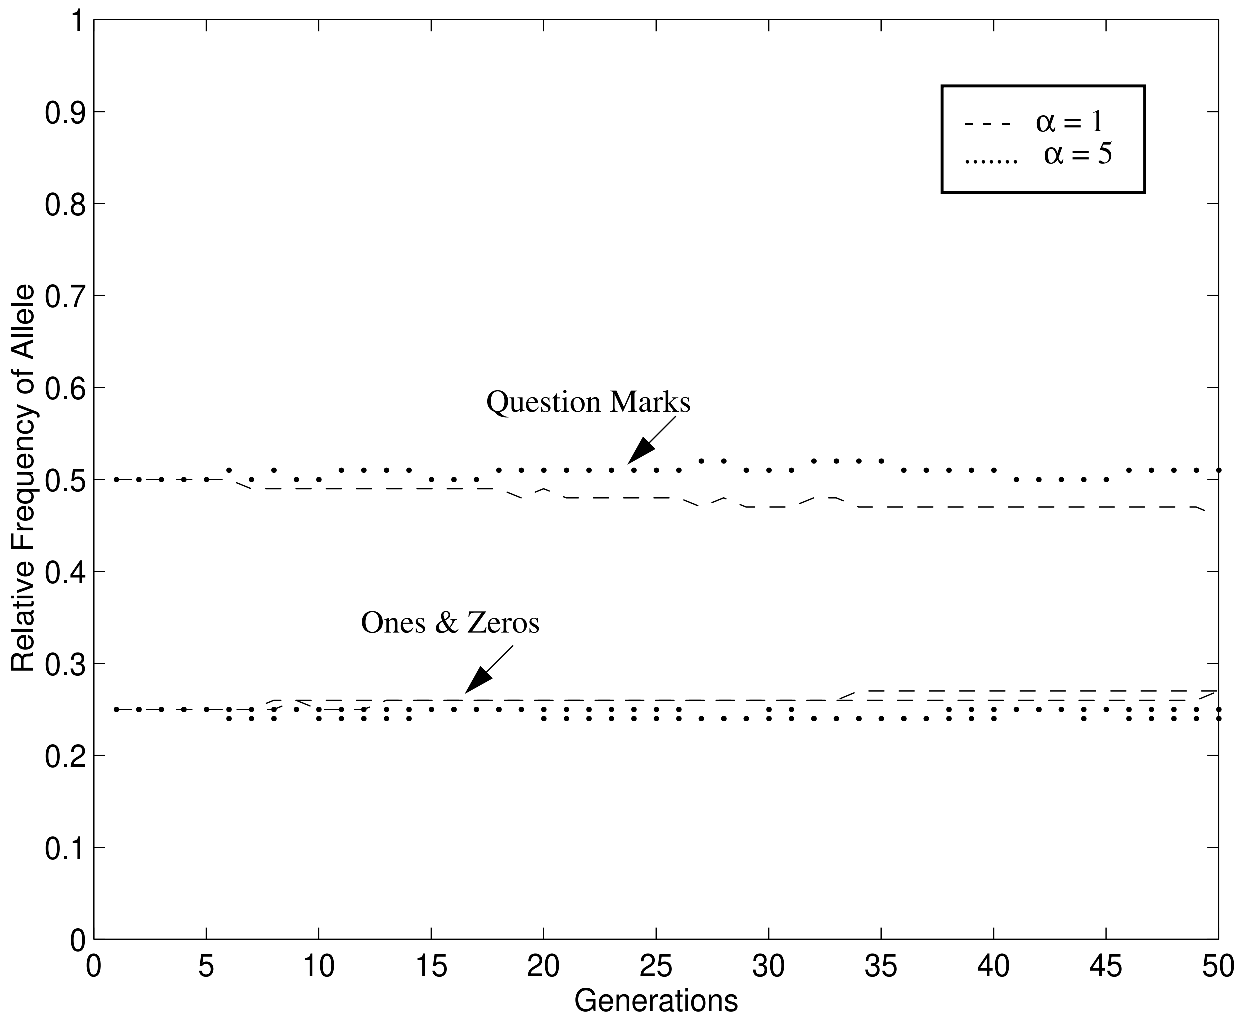

Figure 4.1: Using fitness-proportionate selection, agents are selected as models for social transmission during each learning round. For the dashed lines, α = 1, the expected number of times a winner (should one exist) is selected as a model per round is once (E(win)=1) and the expected number of times a loser is selected is.002 (E(lose)=.002). For the dotted line, E(win)=10 and E(lose)=.02. At these low learning rates the adaptive goal is not realized. The population size is 1000 agents.

\

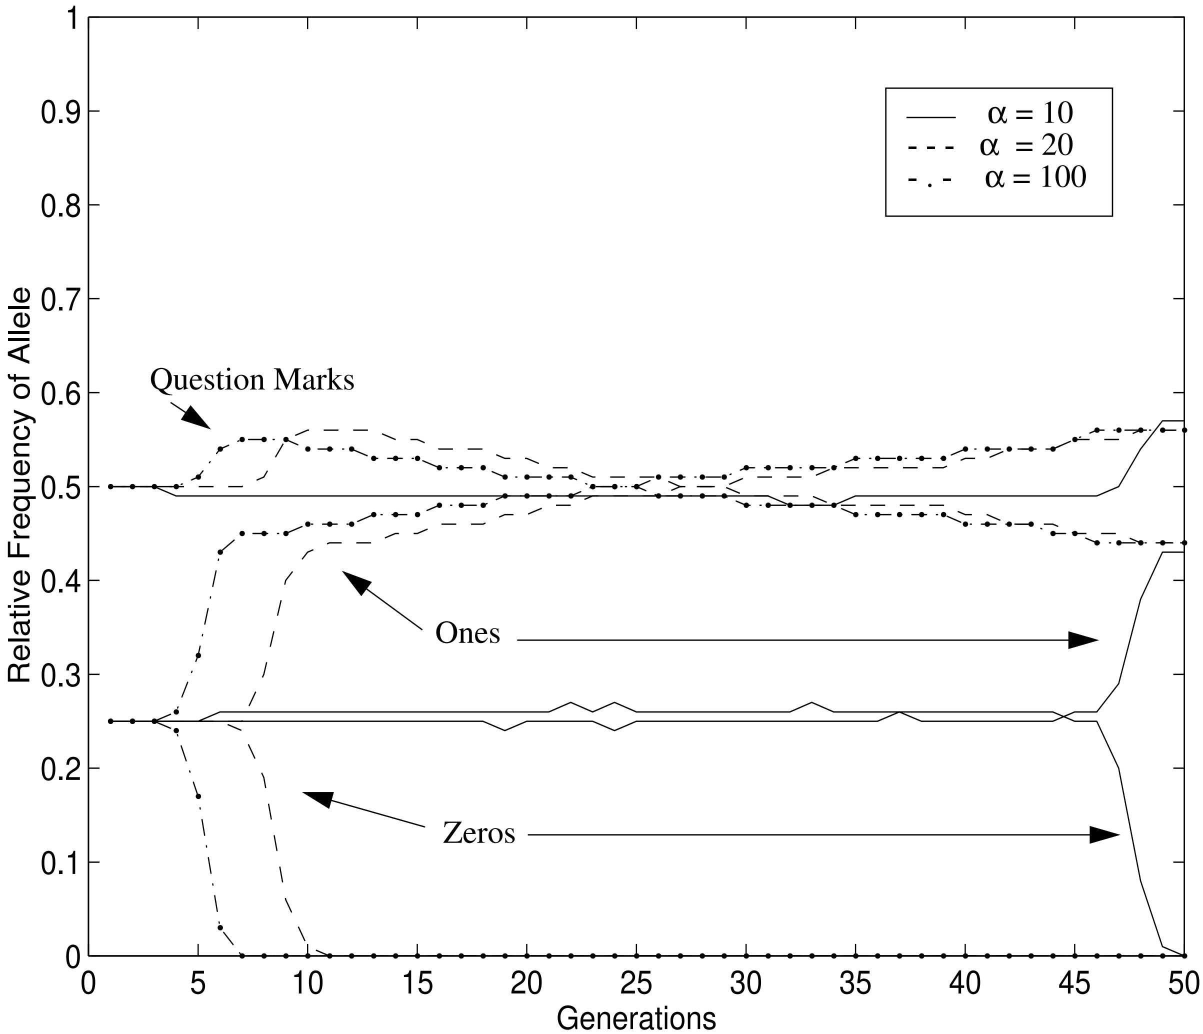

Figure 4.2: Same graphs as in Figure 4.1 but with increasing transmission rates, α. At these higher rates, the adaptive goal is realized. The population size is 1000 agents.

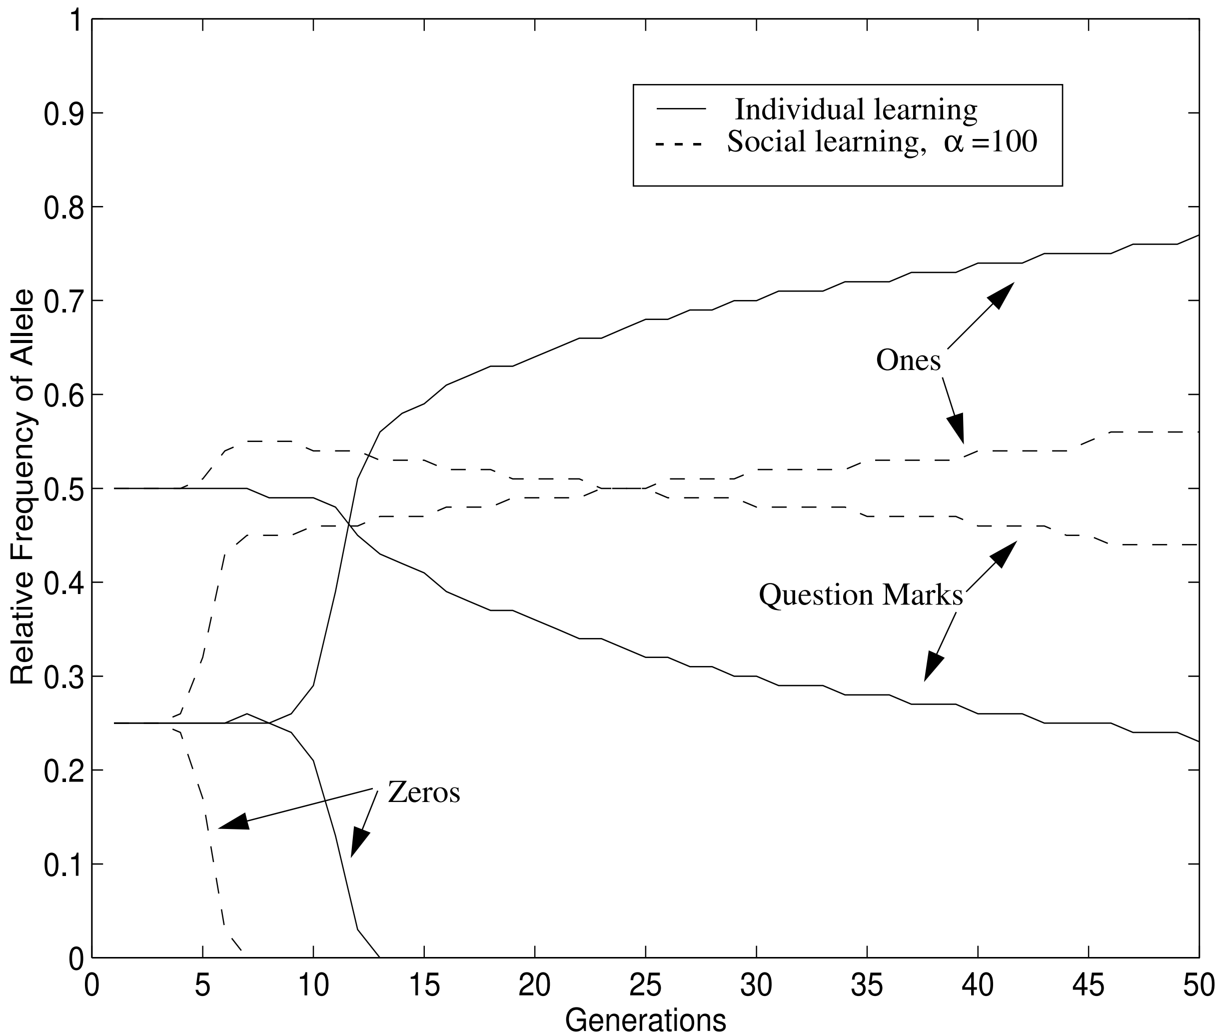

Figure 4.3: Direct comparison of individual and social learning. With a social learning transmission rate of α=100, the adaptive goal is converged on more swiftly then under individual learning. This in spite of the smaller number of agents actually engaged in learning per round. Note that under social learning there exists less pressure to assimilate “1’s” and eliminate “?’s”. The population size is 1000 agents.

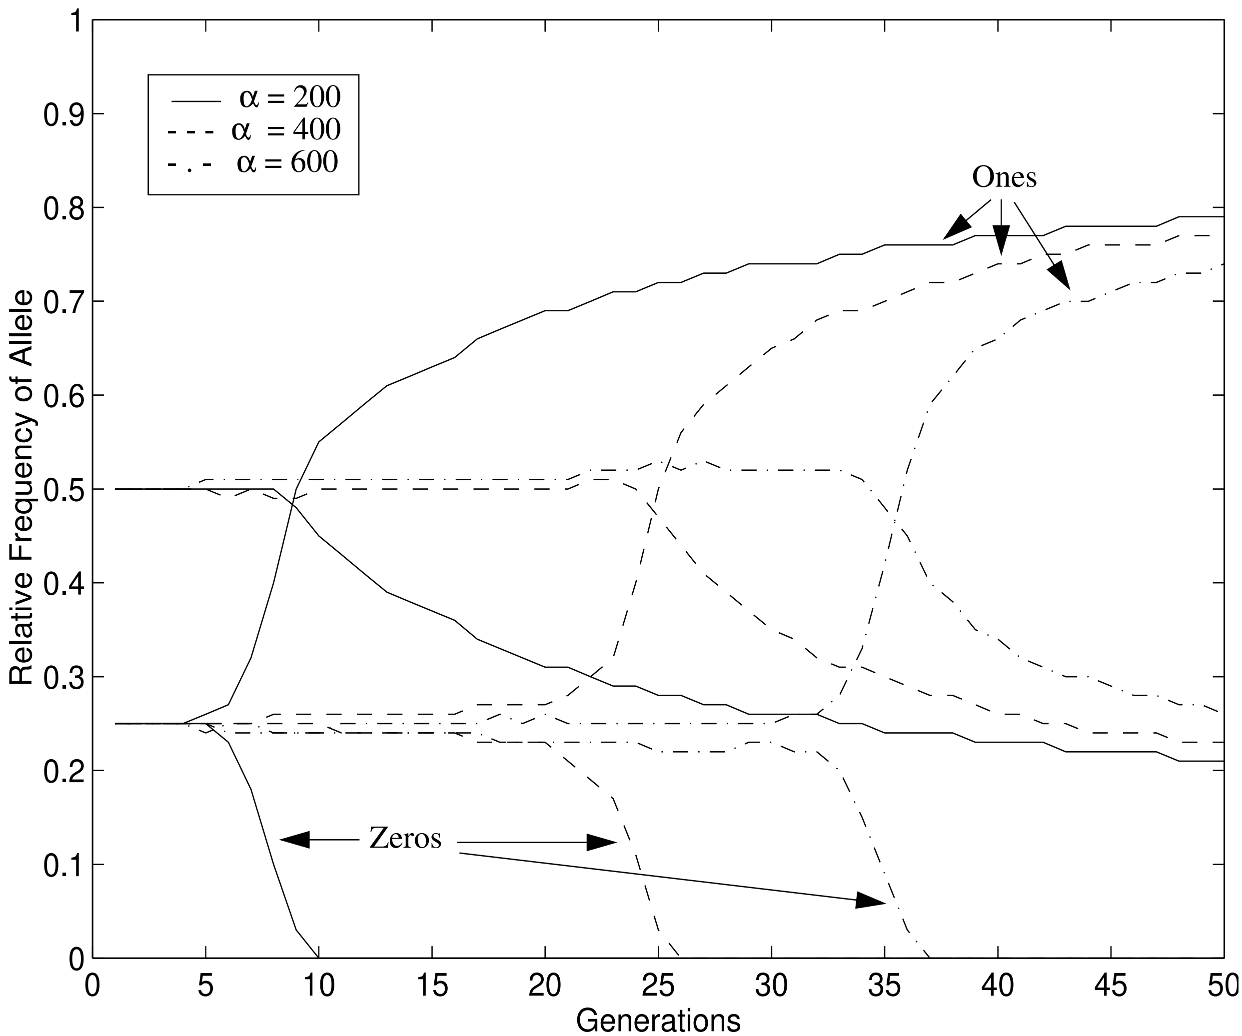

Figure 6.1: Social learning in opposition to organic evolution with individual learning. While the fitness landscape for organic evolution rewards the all-one configuration, social learning rewards and transmits the all-zero configuration. Agents are selected as models for social learning using fitness-proportionate selection. The expected number of times an all-zero winner is selected is shown for $\alpha = 200$, 400, and 600. Greater values of α cause “disruption” in organic evolution and delay convergence towards its all-one goal. The population size is 1000 agents.

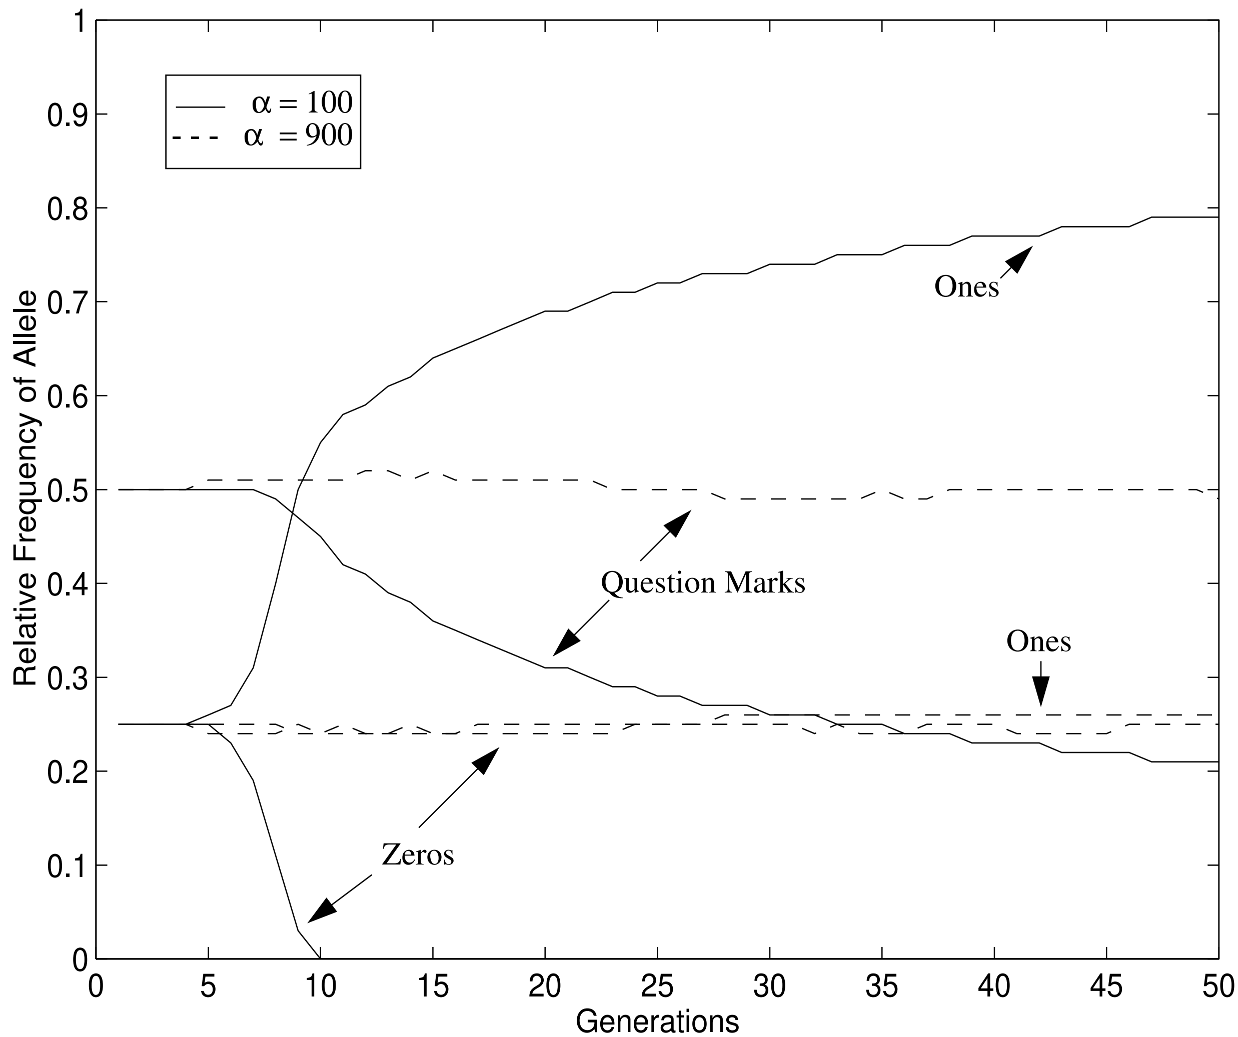

Figure 6.2: Same graphs as in Figure 6.1, but here we compare strong versus weak social learning directly. For a low value, $\alpha = 100$, social learning is not able to change the normal convergence pattern of organic evolution and individual learning. With a high value, $\alpha = 900$, social learning is able to keep organic evolution from making any progress, even when run for many generations. The population size is 1000 agents.