Effect of phenotypic plasticity on adaptation and evolution: A genetic algorithm analysis

Narayan Behera

Theoretical Sciences Unit, Jawaharlal Nehru Centre for Advanced Scientific Research, Jakkur Campus, Jakkur Post, Bangalore 560 064, India
Present address: Centre for Ecological Research, Kyoto University, Kyoto 606-01, Japan

Phenotypic plasticity is the change in the expressed phenotype of a genotype as a function of the environment. However, due to stochasticity in the developmental process, different phenotypes can result from identical genotypes even in the fixed environment. The evolutionary consequences of phenotypic plasticity is examined within the genetic algorithm framework. A genetic algorithm is an important computational tool to model any evolving system. We consider a population of genetically haploid individuals of fixed size. Genotypes are represented by one-dimensional strings of three allelic states, designated 0, 1 and X. 0 and 1 stand for fixed states with predetermined effects on the phenotype. But X is a plastic state whose phenotypic effect can be equivalent to that of a 0 or 1, the actual choice being determined by the regulatory genes. If the regulatory genes are absent, the phenotypic effect of X is determined on a random basis. The fitness values of the individuals are arranged in descending order and viability selection is performed. Then selected individuals are allowed to mate randomly to generate progeny for the next generation. The population size is kept constant from generation to generation. The population evolves by means of genetic recombination and natural selection. The process is repeated until fixation is reached. We find that phenotypic plasticity, up to a certain optimal level, slows down the rate of evolutionary change but improves the degree of adaptation finally reached. There exists an optimum plasticity for which the population is best adapted. There is a synergestic effect of regulatory genes on plastic alleles: the frequency of such alleles increases when regulatory loci are present. The regulatory genes, under certain conditions, can improve the adaptation of the population and speed up the rate of evolution.

One goal of evolutionary theory is to build a quantitative model that can encompass the complexity of how a genotype becomes translated into a phenotype. Such a model would need to incorporate external environmental effects as well as internal constraints based on a detailed knowledge of developmental programs[1]. The environment plays a dual role in the evolutionary process, affecting the developmental process and setting the fitness function. The extent to which phenotypes are subject to environmental modification is commonly termed phenotypic plasticity[2,3]. Plasticity 'masks' less-than-fit genotypes; in other words it enlarges the range of selective environments within which a genotype can function effectively. Phenotypic plasticity can have significant consequences for evolution[4–8]. A possible reason for the existence of plasticity is that the effect of an allele at one genetic locus is modifiable by alleles at other loci. By causing a gradual lowering of the environmental threshold necessary to elicit an appropriate response from the modified locus, the optimal phenotype can come to be expressed constitutively ('genetic assimilation')[9]. Plasticity can also be expressed in the same environment, leading to phenotypic variation within a population of identical genotypes. In other words, plasticity implies that even in a fixed environment the relation of phenotype to genotype is not one-to-one. This would imply that there is a stochastic element to development. Hinton and Nowlan showed in a computational study that, under specified conditions, phenotypic plasticity can modify evolutionary pathways and accelerate the course of evolution[10].

Plasticity is not directly dependent on the environment. The known instances of this come from studies on multicellular development and differentiation. For example, in the social amoeba *Dictyostelium discoideum*, genetically identical cells raised in the same environment can choose, on a cell-autonomous basis, either a stalk-like or a spore-like alternative[11]. In other situations, a comparable phenomenon has been termed 'stochastic differentiation'[12,13]. The concept of a mixed strategy in animal behaviour is yet another example of phenotypic plasticity in the sense used here[14]. The familiar observation that genetically-based traits display varying levels of penetrance in different individuals even when genetic variation is essentially absent is also a manifestation of phenotypic plasticity. The intermediate forms of *Antirrhinum majus* discovered by Darwin in his breeding experiments serve as an example[15].

Behera and Nanjundiah[16,17] have examined the evolutionary consequences of phenotypic plasticity in computational models of haploid genotypes represented by one-dimensional strings. These studies assumed that individual genotypes consisted of structural genes only, meaning genes with direct consequences for the phenotype. They have extended the model by incorporating

a qualitatively new feature, namely gene regulation[18]. The consequence of a potentially 'plastic' allele at a structural locus is assumed to be modifiable, in a probabilistic fashion, by alleles at the regulatory loci. Regulatory genes influence the phenotype only indirectly, via structural genes. Our interest is restricted to asking, (a) in what manner is genetic equilibrium (fixation) reached? (b) how are the rate of evolution and equilibrium population fitness affected by the possibility of phenotypic plasticity? and (c) how does the presence of regulatory loci affect the evolution of genotypes and phenotypes? The rest of the paper consists of a description of the model followed by the results of computations and, finally, a discussion.

Model

No gene regulation

The model is based on that of Hinton and Nowlan[10]. Populations consist of individuals of haploid genotype. Each genotype is represented by linear strings of 0s, 1s and Xs, where each symbol stands for one of the three possible alleles at a genetic locus. A 0 or a 1 means that the allele at that locus has a constitutive expression and contributes in a pre-determined fashion (as described below) to the overall phenotype. An X on the other hand means that the expression of the gene at that locus is facultative; X can act either like a 1 or a 0 allele in terms of its contribution to the phenotype. Which of the two alternatives it adopts is determined by a 'coin-tossing' procedure, in other words by making use of a random number generator. A complete series of coin tosses performed in one generation on a single genotype is referred to as a *trial*. Operated on each genotype in every generation, each *trial* produces a set of uniformly distributed (pseudo) random numbers between 0 and 1. By carrying out 'coin-tossing' independently for each X in a genotype, we decide whether that X mimics a 0 or a 1. At each 'toss' every X has a 50% probability of becoming functionally equivalent to a 1 (and with the same probability, of becoming equivalent to a 0). Equivalence refers to the phenotypic effect of the X and applies just for that generation. The *a priori* probabilities of the 0, 1 and X alleles (i.e. the probabilities used to create the starting population of genotypes) are denoted by p_0, p_1 and p_x respectively, with $p_0 + p_1 + p_x = 1$. Given N_s loci (N_s refers to the number of structural loci), the mean number of Xs in any genotype is $m = N_s p_x$. The maximum number of coin-tosses for each genotype is restricted to 2^m. Coin-tossing is stopped after a particular trial if all the Xs in a genotype becomes 1s, otherwise it continues until 2^m *trials* are exhausted. (Notice that after the coin-tossing procedure is completed, every genotype is indistinguishable, in terms of how it relates to the environment, from some genotype made up solely of 0s and 1s.) Finally, each genotype is assigned a fitness that depends both on its attained phenotype and on the number of tosses of the coin needed to reach that phenotype. Fitness is in part a measure of how closely the final phenotype resembles the phenotype of an ideal or 'target' genotype. The ideal phenotype is taken – for the sake of simplicity – to be a string of 1s, i.e. {1, 1, ... 1}. The target genotype has the maximum possible fitness (equal to 1).

Thus the fitness of a genotype has two components. The first, which can be thought of as reflecting in part the 'hard-wired' or deterministic aspect of its phenotype, depends only on the degree of matching between the attained phenotype and that of the target. This component of fitness is defined as

$$W_D = \frac{1}{N_s} \sum_{i=1}^{N_s} d_i , \tag{1}$$

where i stands for a genetic locus and $d_i = 1$ if the ith locus has a 1 allele; and $d_i = 0$ otherwise. The second component of fitness, denoted by W_p, is related to the dynamical, 'plastic' or 'soft-wired' aspect of the phenotype. The assumption is that the larger the number of coin-tossing trials undertaken in the course of attempting to reach the target, the smaller the value of fitness. This component of fitness is defined as

$$W_p = \frac{2^m - \text{Actual number of trials}}{2^m} . \tag{2}$$

Note that the maximum number of allowed trials is the same for every genotype and is independent of the number of Xs it carries. The larger the initial number of Xs in any given genotype, the less likely it is that all of them will become 1-like within the 2^m trials and so the lower the likely value of W_p. In a sense, W_p reflects a cost associated with plasticity. If all the 2^m trials are gone through, $W_p = 0$ irrespective of whether the Xs become 1-like on the last trial or not. If all structural locus alleles are 1s to begin with, $W_p = 1$ because no trial needs to be undertaken. Also, $W_p = 0$ whenever a genotype has no X allele and at least one 0 allele. On the other hand, if coin-tossing is stopped before the full quota of trials is exhausted, W_p takes a value between 0 and 1. The total fitness is defined by

$$W_T = f W_D + u(1 - f) W_p , \tag{3}$$

where f is a positive number lying between 0 and 1 and represents the fractional weightage assigned to the deterministic part of the fitness. By varying f one can vary the significance of the cost of plasticity. A value

of $f=0$ would imply a high cost whereas $f=1$ would imply a zero cost. Note that fitness is a characteristic of an individual (phenotype) rather than of a genotype. The larger the plasticity, the more is the allowed number of coin-tossing trials. The higher number of trials require more expenditure of energy, so the effective fitness of W_p is reduced. Hence the factor

$$u = \frac{a}{a + 2^m} \tag{4}$$

is inserted in the second term of the expression for W_T. This is dependent on the level of plasticity. Here a equals 2^{10} and is the maximum number of possible trials in any model, obtained when the *a priori* probability of X takes its highest value of 0.56 (in our simulations).

The only elements of the model that remain to be described are mating and recombination. Mating is at random and is assumed to be followed by a single obligatory crossover between parental genomes. One offspring is generated at each mating by choosing a crossover point at random and copying all alleles from the first parent up to the crossover point, and from the second parent beyond the crossover point. After fitnesses are computed for 1035 haploid individuals, genotype frequencies are weighted in proportion to their relative numbers and arranged in descending order of fitness. Truncation selection is performed by restricting the number of individuals that go on to mate to 46. Random mating generates 1035 ($=^{46}C_2$) individuals for the next generation. We continue the analysis for as many generations as necessary until fixation is reached.

Gene regulation

The model is described in Figure 1. Each genotype is represented by two chromosomes, idealized as one-dimensional strings of lengths N_s and N_r respectively. Of these, N_s stands for the number of structural genetic loci and N_r for the number of regulatory loci. A structural locus has three possible allelic states, represented by 0, 1 and X; a regulatory locus has just two possible alleles, represented by 0 or 1. In the case of a regulatory locus, a 0 or 1 contributes to regulating the expression of X alleles in structural loci in a manner to be explained below. Thus structural loci contribute directly to fitness but regulatory loci do so indirectly, by influencing the phenotypic contribution of Xs in the structural loci. The probability that an X becomes equivalent to a 1 (or to a 0) in terms of its phenotypic effect is equal to p_{1r} (or p_{0r}), the frequency of 1s (or 0s) at regulatory loci. Regulation can be mediated either by closely-linked DNA sequences (*cis* regulation) or by diffusible intermediates that are the products of genes that may be unlinked to the gene that is regulated-*trans* regulation. In terms of its operational consequences *vis-à-vis* our model, the central difference between *cis* and *trans* regulation is that recombination can modify the arrangement of structural genes and their *cis* regulatory elements. For reasons of simplicity, in the present study we restrict ourselves to an examination of the consequences of *trans* regulation. Note that because of this, for a given value of p_{1r}, about $(1/p_{1r})^m$ trials ought to be sufficient to convert a fair proportion of Xs to 1s. The entire sequence of 2^m coin-tossing trials is carried through unless one of the intermediate trials results in all the Xs in a genome becoming 1-like.

The fitness of the individuals is computed as before (without gene regulation). For simplicity, we have taken $u=1$ in equation (3). After the starting 1056 genotypes have been ordered in terms of their relative fitnesses, truncation selection is performed and the 1056 individuals are reduced to a mating population of 33 adults. Mating takes place between every pair of individuals. We assume that two obligatory crossovers occur at each mating, one each within the regulatory loci and the structural loci. The crossover points are chosen at random. After

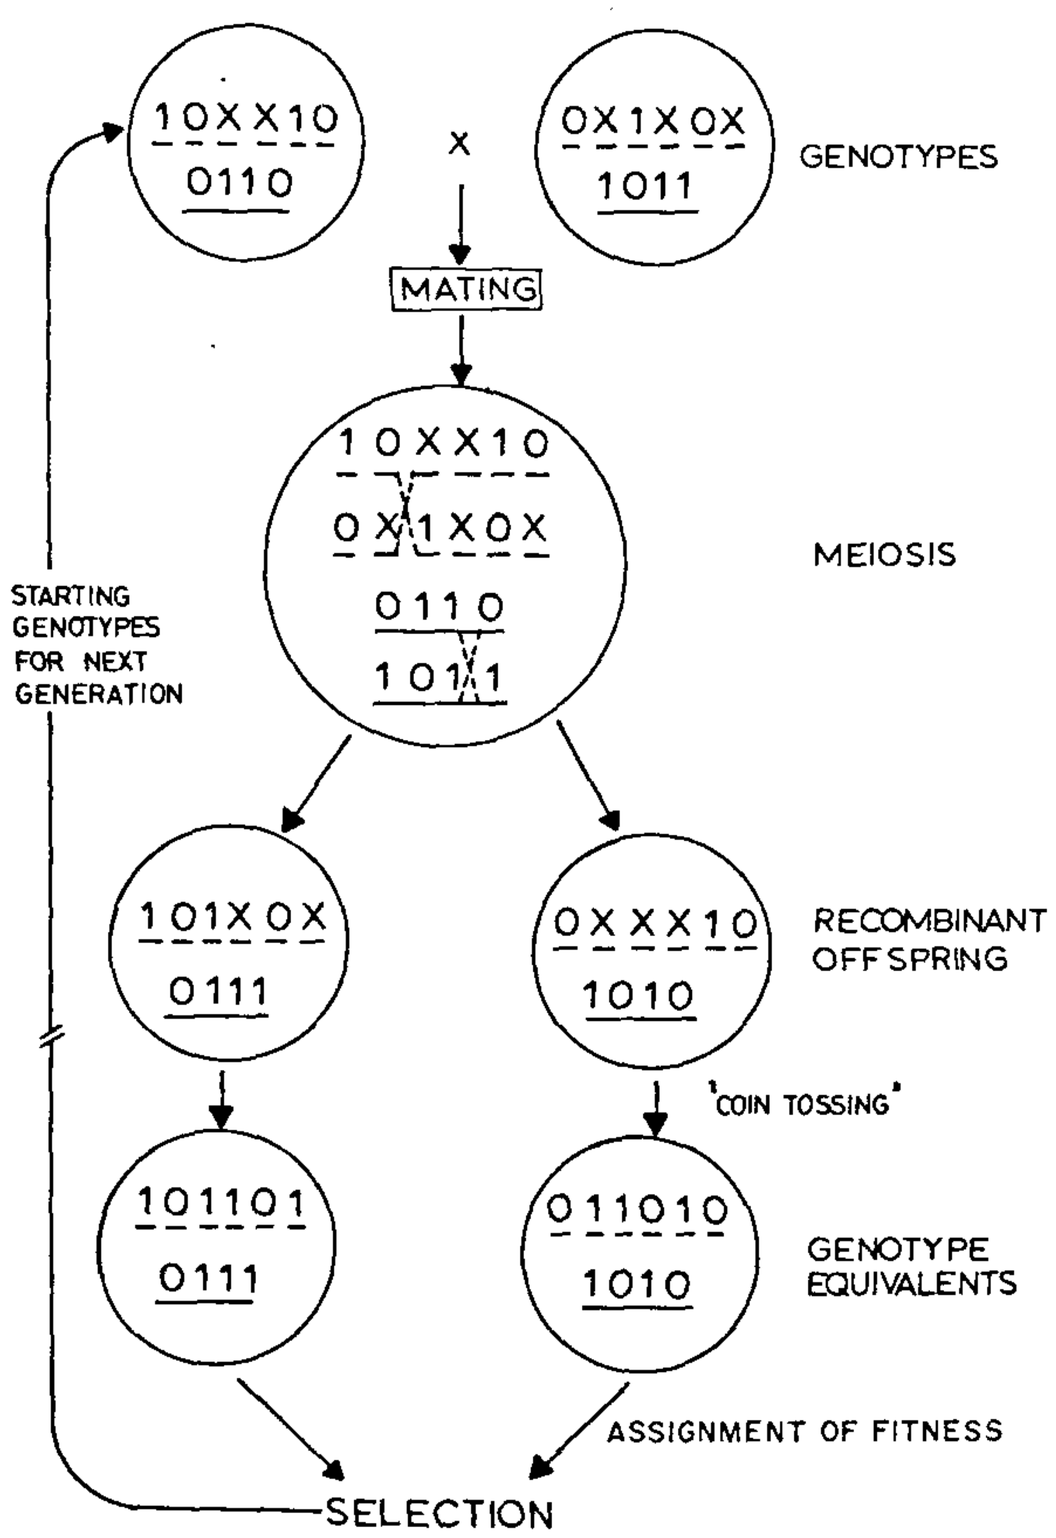

Figure 1. A diagrammatic sketch of the model. Each genotype is idealized as consisting of two haploid chromosomes, one of which contains structural genetic loci (dashed lines) and the other contains regulatory loci (continuous lines). Meiosis is accompanied by a single obligatory crossover in each chromosome.

meiosis one should expect there to be four haploid recombinant genotypes. Of these, we pick just two by applying the following rule (Figure 1): the genotype of the first offspring is chosen by copying all alleles from the first parent starting from the left up to either crossover point, and from the second parent beyond crossover point. Correspondingly, the genotype of the second offspring is made up of all alleles from the second parent to the left of either crossover point and from the first parent beyond the crossover point. Consequently the 33 pairs of genotypes that mate give rise to 1056 individuals which constitute the next generation. The simulation is continued until fixation is reached at all loci so that there can be no further genetic evolution in the absense of mutational change.

Results

The outputs that we monitor are (a) allele frequencies from generation to generation, (b) the number of generations needed for fixation to be attained, and (c) the mean population fitness during the course of evolution as well as at fixation. It might appear that a finite population size, and the fact that we pick just one out of two (without gene regulation) or two out of four (with gene regulation) recombination products, predisposes the system to drift. In fact the major source of variation in the outcome is the initial random choice of genotypes and truncation selection (see Discussion). Each simulation has been repeated at least six times with the same initial conditions.

We have put limitations on population size and the number of genetic loci because of constraints on computer time. Our choice of parameters is governed by the following considerations. When W_D is the only contributor to fitness (see equation (1)), we choose $p_1 = 0.03$. For smaller values of p_1, equilibrium is reached very far from the target, meaning that the population remains poorly adapted; for higher values of p_1 the target is reached extremely rapidly, i.e. with virtual certainty and within a very short time. Finally, where the fitness has two components, we have retained $p_1 = 0.03$ (and $p_0 + p_x = 0.97$) in order to compare the fitness function W_T (equation (3)) with that of W_D. With p_x much larger than 0.5, limits on computer time become significant because the total number of coin tosses rises faster. In the presence of regulatory genes, we have chosen the same set of initial conditions in order to make it easier to compare outcomes with and without gene regulation.

No gene regulation

Allele frequencies. In the absence of any X allele, that is, in the absence of plasticity, drift is the sole cause of change of allelic frequencies (unless the target genotype happens to be present at the start or appears as a result of recombination). The causes of drift are two-fold: first, the restriction to a small population size and second, the element of randomness in deciding which of the two offspring of a mating survives to reproduce. The chances of the target ever being reached are very low (we never found this happening in our simulations). Solely on account of chance, one or the other genotype gets fixed eventually and this occurs on average around the 33rd generation (Figure 2 *a*). (No meaning should

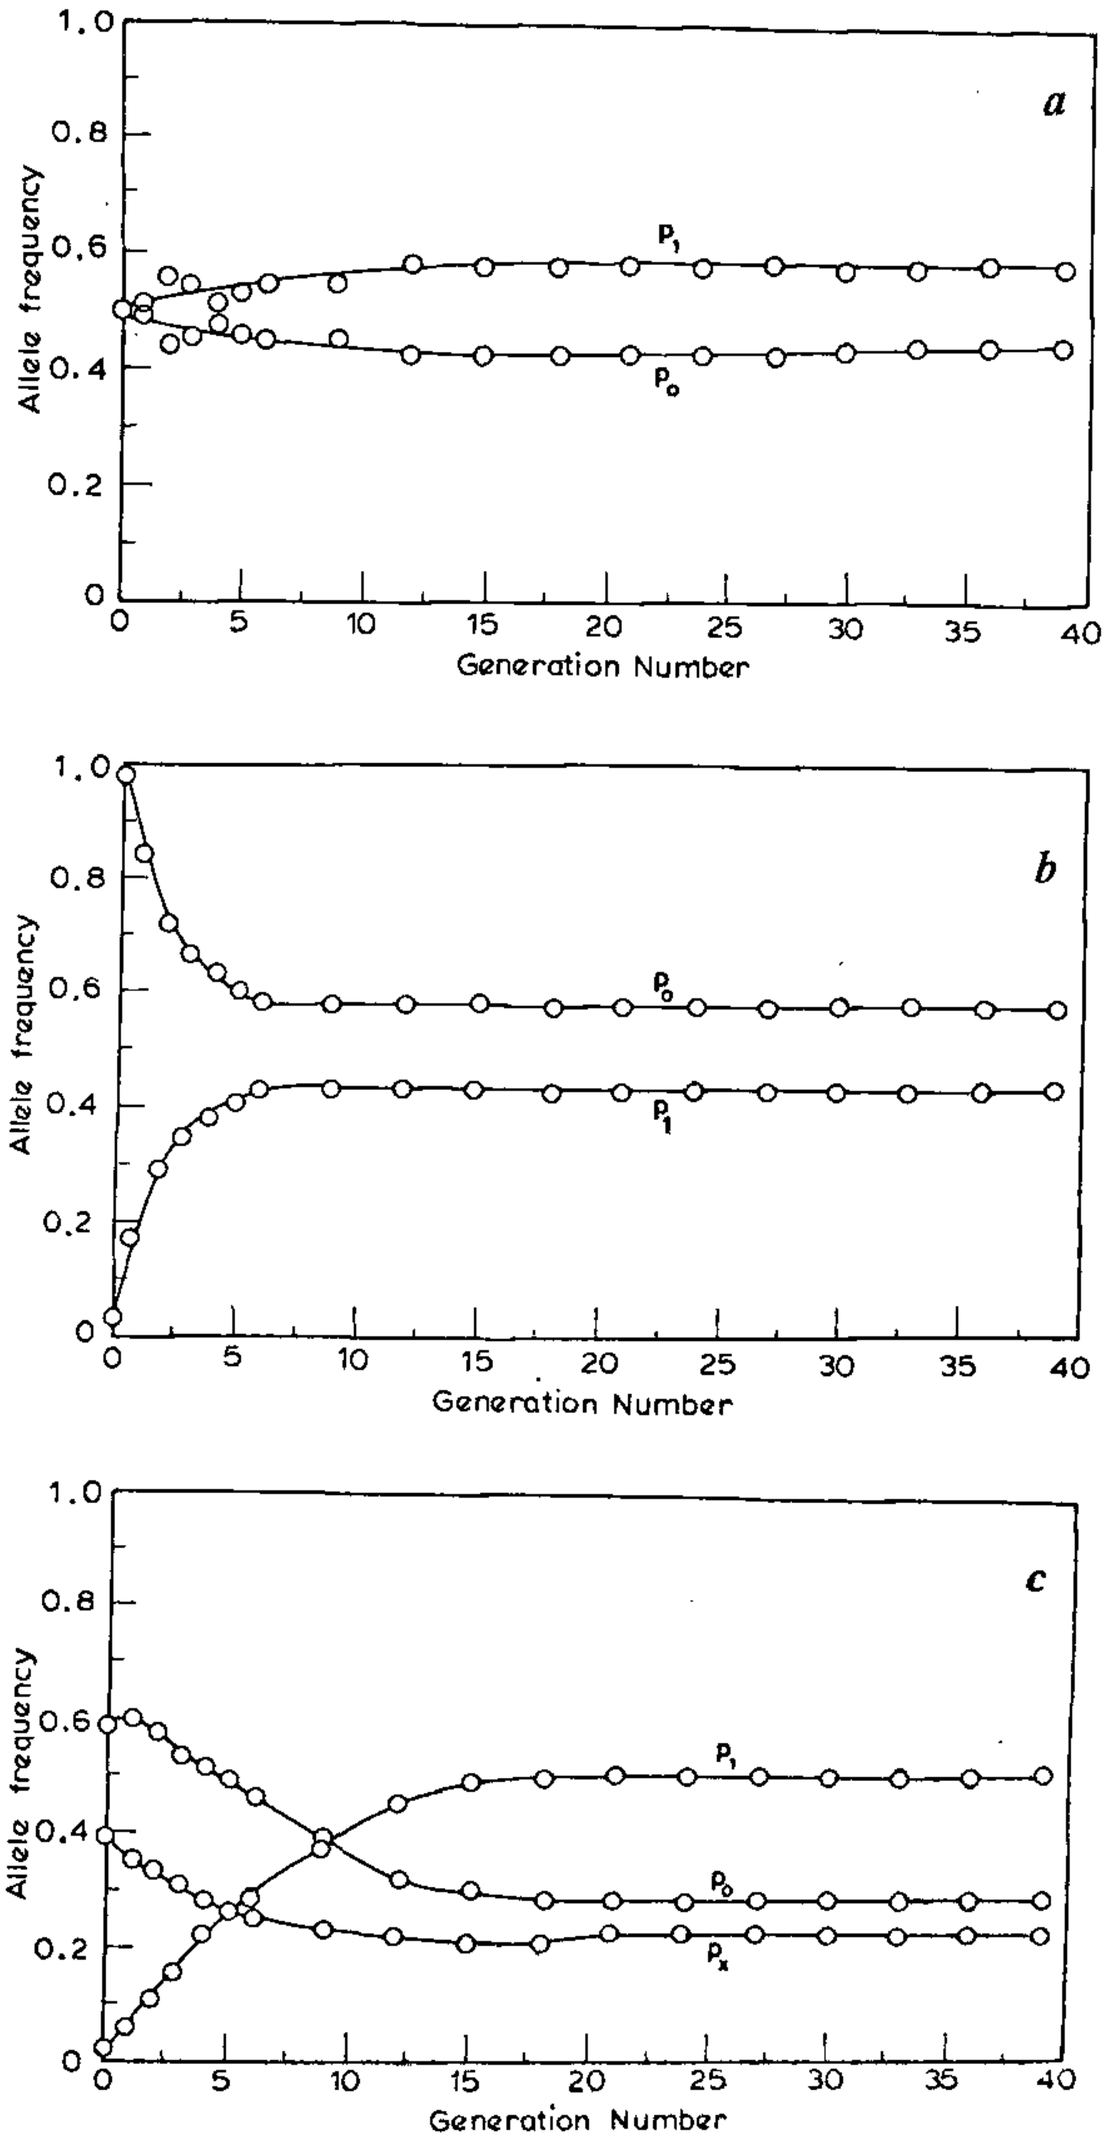

Figure 2. Change in allelic frequency with generation. Generation numbers start with zero. ***a***, Without plasticity, and no fitness differences, allelic frequencies change due to recombination and drift only. ***b***, Without plasticity, the fitness function is W_D only. The initial frequency of 1 allele is 0.03. ***c***, The fitness function is W_T.

be attached to the fact that in Figure 2 *a* there is a small increase in frequency of the 1 allele.)

The introduction of plasticity dramatically changes the situation. Even if the target genotype never appears in the population, it is possible that some genotypic combination of 1s and Xs acquire the same phenotype as that of the target. Once this happens, natural selection acts very strongly in its favour – or more correctly, acts against 0 alleles. Here fixation takes place faster than it does in the absence of plasticity (by about 20th generation instead of 33rd – compare Figure 2 *a* and 2 *c*). A similar speeding up in the rate of evolution can occur without plasticity if fitnesses are scaled with respect to distance from the target (Figure 2 *b*). However, the increase in the rate of evolution is balanced by the fact that the final genotype is at a greater distance from the target than in the case wherein plasticity is present (Figure 2 *b* and 2 *c*).

It is interesting to monitor evolutionary rates as a function of the degree of plasticity (by which we mean p_x, the *a priori* probability of an X allele at any locus). Plasticity slows down evolution instead of speeding it up (data not shown), a finding that is independent of whether or not we take into account a cost factor associated with coin-tossing (see equation (3)). At constant plasticity the generation number at which fixation occurs is highest when $f = 0.5$ and it slowly decreases as f tends to 0 or 1. An unequal weightage for W_D and W_P speeds up the rate of evolution in comparison with the situation of equal weights (not shown).

Mean fitness. The mean population fitness increases steadily with generation number until 'equilibrium' is reached (Figure 3). Note that in the presence of phenotypic plasticity, mean and maximum fitnesses differ even when a state of equilibrium has been attained. The explanation is that the population may consist of a single genotype, but if that genotype has one or more X alleles, the fitness of individual members of the population can differ depending on whether any given X mimics a 0 or a 1. The mean fitness at equilibrium increases with the degree of plasticity till the optimum p_x, beyond which it decreases (Figure 4). If fitnesses are combined with unequal weightages, the mean fitness at equilibrium is a minimum at $f = 0.5$ (not shown) but the reasons for its rise as f approaches 0 or 1 are different (see Discussion).

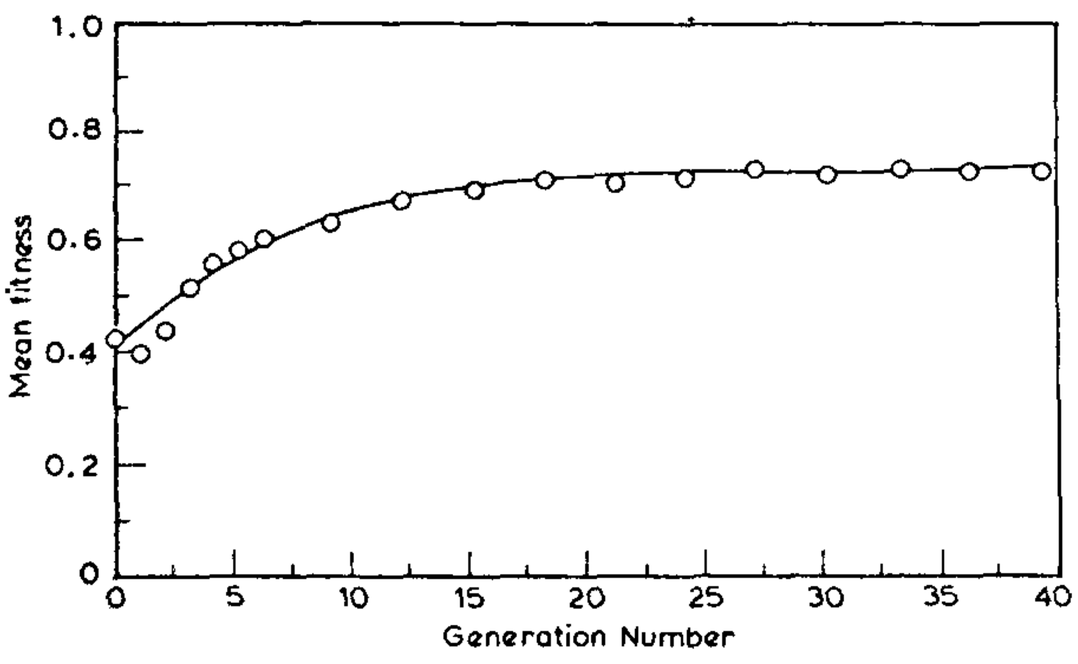

Figure 3. The change of mean fitness with generation number (fitness function is W_T, $p_x = 0.5$).

Gene regulation

Change of allele frequencies. In general, depending on initial conditions, the frequency of the 1 allele averaged over structural loci, p_{1s}, can either increase or decrease in the course of time (Figure 5) but that of the X allele always increases. On the other hand, the frequency of

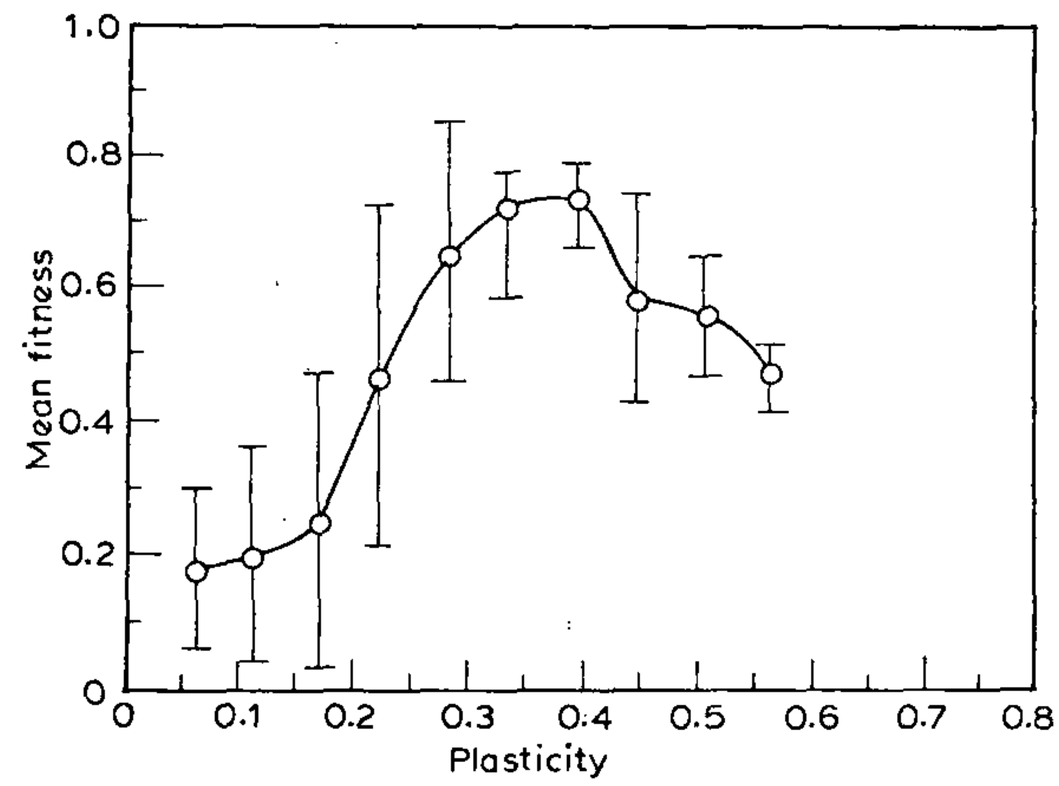

Figure 4. The change of mean fitness with plasticity (fitness function is W_T, $f = 0.85$ and cost function included). The actual value of u depends on the degree of plasticity p_x (see equations (3) and (4)). Note that the mean fitness is a maximum at $p_x = 0.39$.

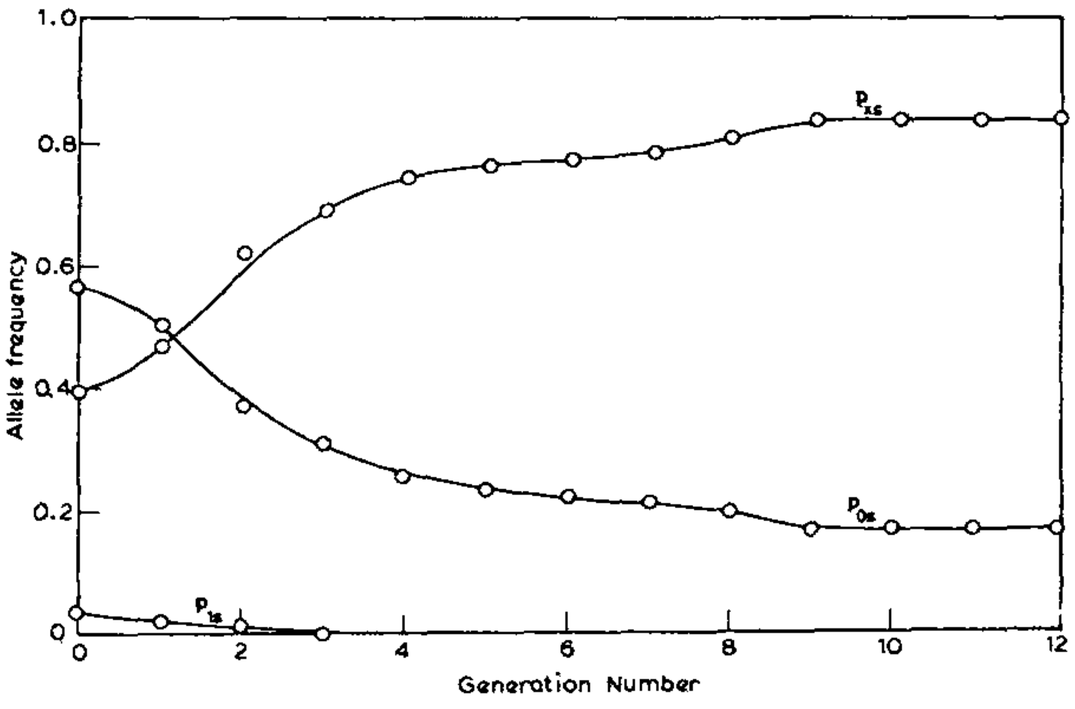

Figure 5. Change in allelic frequencies at structural loci as a function of generation number (single simulation). Initial conditions: $p_{xs} = 0.39$, $p_{1s} = 0.03$, $p_{1r} = 0.20$ and $N_r = 10$.

the 1 allele at regulatory loci, p_{1r}, increases steadily in spite of the fact that selection does not act directly on regulatory loci (Figure 6 *a*). The increase in p_{1r} is most rapid, and the value reached at fixation is highest, when the number of regulatory loci N_r is small. Indeed, when N_r is very large, p_{1r} barely rises above its initial value before fixation is attained. On the whole, the rate of evolution, as estimated from the number of generations needed for fixation, decreases both as a function of the starting value of p_{1r} and as a function of N_r (data not shown). Finally, the rate of evolution is affected by the relative importance of the two components of fitness and peaks at an intermediate value (0.37) of the weight f (not shown).

Mean fitness. As should be expected for a population evolving under natural selection, the mean fitness $\langle W_T \rangle$ increases over the course of successive generations (Figure 6 *b*). The value of $\langle W_T \rangle$ at fixation rises with the initial value of p_{1r} (Figure 7 *a*). For a given initial value of p_{1r}, $\langle W_T \rangle$ at fixation is highest when there is just one regulatory locus (Figure 7 *b*). $\langle W_T \rangle$ varies only weakly with f; indeed, $\langle W_T \rangle$ is a minimum at $f = 0.34$ (not shown). The mean fitness at equilibrium decreases in a nonlinear fashion as a function of the ratio of N_s to N_r (Table 1).

Discussion

The most common argument in favour of plasticity is

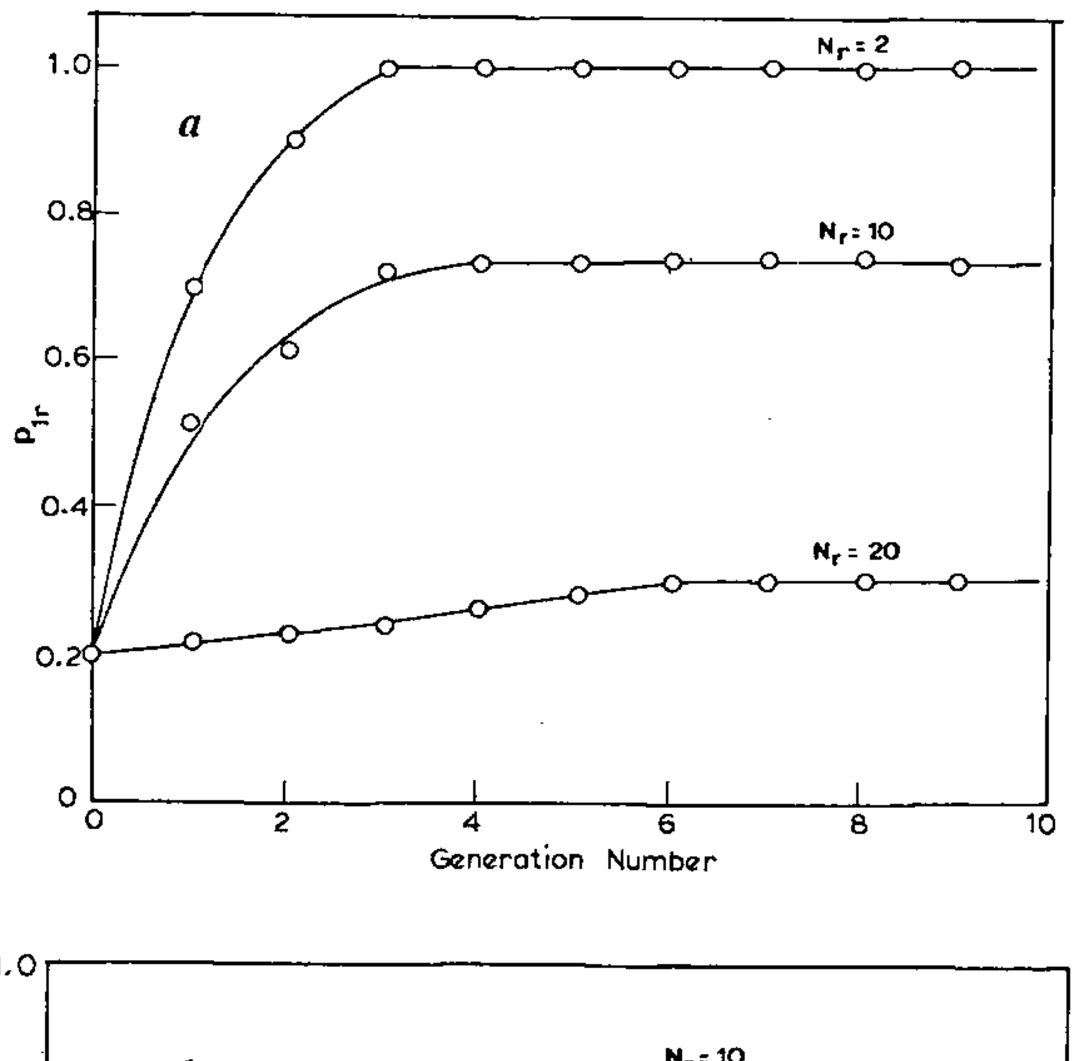

Figure 6. Evolution at regulatory loci via hitch-hiking (*a*) and the consequent changes in mean population fitness (*b*) as a function of generation number (single simulations). *a* and *b* show how the changes in p_{1r} and mean fitness depend on N_r. Initial conditions: *a*, $p_{ss}(0) = 0.39$, $p_{1s}(0) = 0.03$ and $N_r = 10$; *b*, $p_{ss}(0) = 0.39$ and $p_{1s}(0) = 0.03$; *c*, $p_{ss}(0) = 0.39$, $p_{1s}(0) = 0.03$ and $p_{1r}(0) = 0.2$.

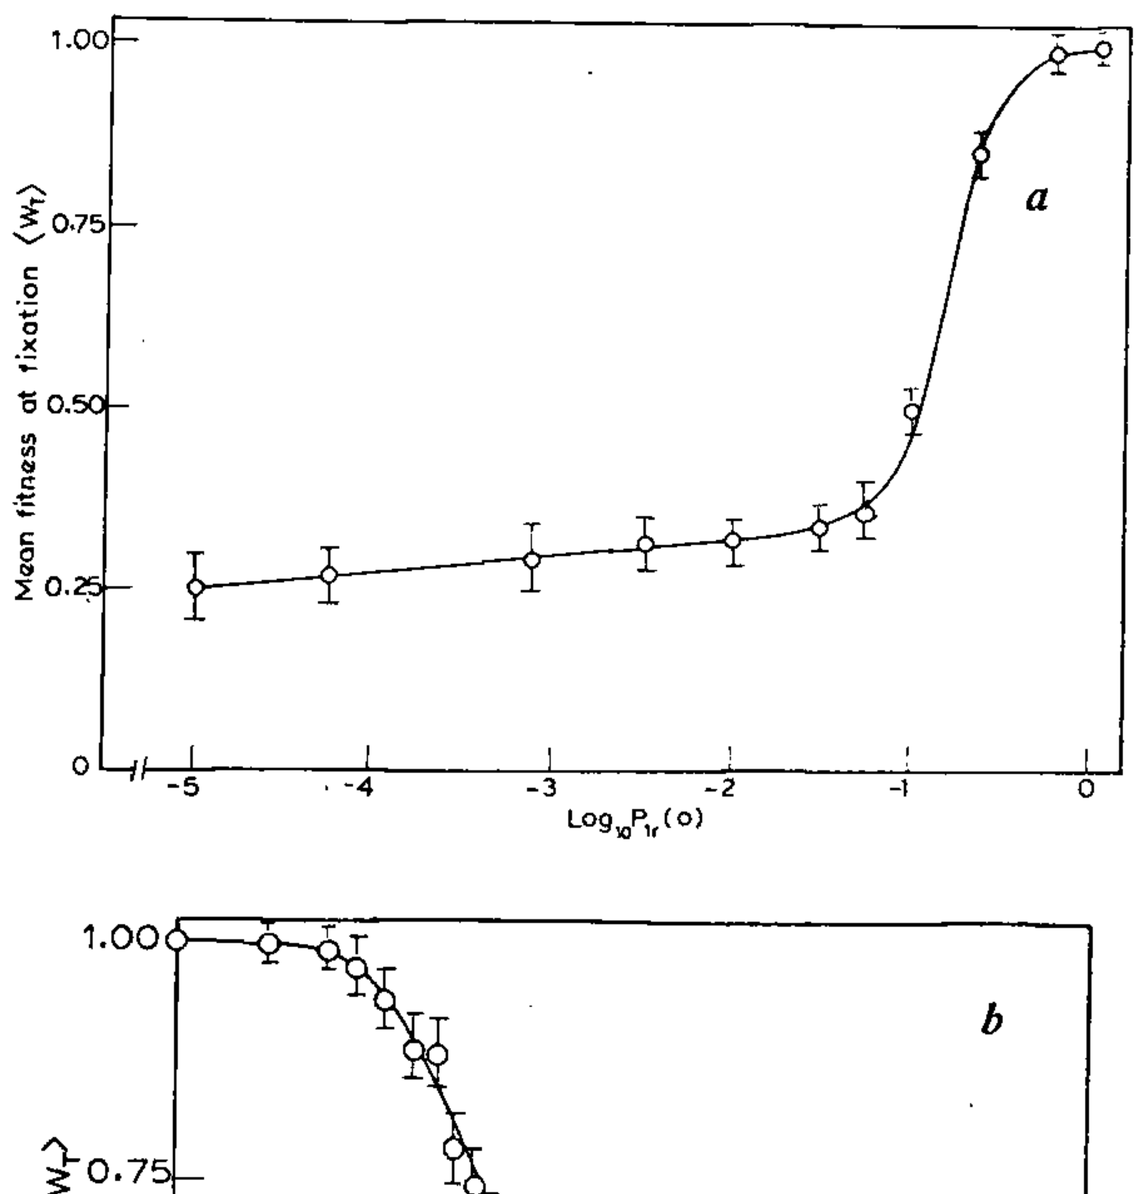

Figure 7. Population mean fitness at fixation as a function of the initial strength of regulation $p_{1r}(0)$ and regulatory loci N_r. Initial conditions: *a*, $p_{ss}(0) = 0.39$, $p_{1s}(0) = 0.03$ and $N_r = 10$; *b*, $p_{ss}(0) = 0.39$ and $p_{1s}(0) = 0.03$ and $p_{1r}(0) = 0.2$.

Table 1. Equilibrium mean fitness as a function of the ratio of structural and regulatory loci at the transition point

No. of structural loci	No. of regulatory loci	Ratio	Equilibrium mean fitness
10	10	1.00	0.57 ± 0.07
15	21	0.71	0.52 ± 0.06
18	28	0.66	0.35 ± 0.03
20	200	0.10	0.31 ± 0.04
25	625	0.04	0.28 ± 0.02

The transition point is defined as that value of the number of regulatory loci beyond which gene regulation reduces the mean fitness of the population at equilibrium. Equilibrium mean fitness is essentially independent of the number of regulatory loci when the number is below the transition point. Initial conditions; $p_{xs}(0) = 0.40$, $p_{1r}(0) = 0.2$ and $p_{1x}(0) = 0.03$ (except when there are 18 structural loci, when $p_{xs}(0) = 0.39$).

Data corresponding to 10 and 15 structural loci are averages of twelve simulations while others are due to six simulations. Constraints on computer time became forbidding when the structural loci numbered more than 25.

that some aspects of the environment are unpredictable. For specifying a phenotype, such unpredictability makes it advantageous to leave some decisions to physiologically adaptive process rather than specifying them genetically. Each coin-tossing trial can be almost as helpful to the evolutionary search as the production and evaluation of a whole new organism. This greatly increases the efficiency of evolution.

Hinton and Nowlan's model is realistic for the specific situation of the formation of precise synaptic connections in a developing nervous system. In such a situation, the 'correct' connections can be strongly context-dependent and also have a correspondingly high premium on being precisely made. One can then see that it makes sense genetically to specify gross features of connectivity and leave fine-tuning to trial and error. However, the evolutionary picture at the back of our mind is rather different and an illustrative example could be the following. Suppose the pattern on an insect's wing helps in camouflage and that a large number of genes act to specify the pattern. If the insect has a wide geographic range, it is plausible that while the optimal pattern differs in detail from location to location, all possible patterns confer some advantage (relative to none). Our interest is in asking whether it helps the individual to fine-tune the pattern by physiological adaptation over and above a genetically specified, common 'basal' pattern.

It is possible to interpret the elements that enter into W_T along the following lines: W_D is that portion of total fitness which is based on the genetically-specified component of the phenotype and W_P represents that portion which is dependent on physiological adaptation based on a random search; f is the relative importance of W_D. A value of $f = 1$ would imply that the phenotype is completely specified in advance, $f = 0$ implies that it is entirely plastic, that is, left entirely to the hazards of random chance. One can justify the parameter u in a different way. It seems reasonable to assume that a completely prespecified genotype extracts a lower metabolic cost than one capable of responding suitably to a range of environments. For example, the latter situation could demand the presence of regulatory genes that can be dispensed with in the former. Translated to fitness, u stands for the price paid by the genotype in return for being permitted a choice of phenotypes.

The degree of adaptation 'lost' on account of a changing environment can to some extent be compensated by decreasing the total number of genetic loci[17]. Decreasing the number of loci amounts in turn to decreasing one measure of complexity of the system. The implication is that other things being equal, systems with many genes – and in that sense 'more complex' – take longer to adapt than those with fewer genes (and so 'less complex'). Therefore, the larger the degree of uncertainty with regard to the external environment, the better off is a system with few genes *vis-à-vis* one with many. The decrease in adaptation caused on account of change in the environment can be partially or completely compensated, not only by reducing the size of the genome but by increasing the plasticity of the genetic system. This means that phenotypic plasticity must have been increasingly advantageous as genome sizes increased during evolution[17]. We note that the inference remains valid until an optimal level of plasticity is attained. Thereafter, an increase in the level of plasticity can slow down the efficiency of evolution[16]. The possible roles of phenotypic plasticity in evolution, and of randomness uncorrelated with the environment ('developmental noise'), have been explored by a number of authors[1,7,8,19]. Broadly, what these earlier authors concluded was that (a) given the right circumstances, phenotypic variation could be maintained in a population; (b) different forms of environmental variation could lead to different evolutionary dynamics (possible with the same end-point); and (c) stabilizing selection on a quantitative trait could reduce the variance in that trait.

It might appear that we have made other compromises

with biological realism. Our model for recombination ignores two of the four possible products of a crossover; the model does not guarantee a stable phenotype (because 'switching' persists whenever there are one or more X alleles in structural loci); and the way we look at evolution is different from the standard population genetics approach (because we generate genotypes containing many Xs right at the start instead of asking what happens when a single X invades a population that contains only 1 and 0 alleles). Interestingly, the first compromise leads to the consequence that each mating preserves the allele frequencies of the parents, besides which it produces exactly two offspring. Therefore the primary forces driving gene change are initial conditions, truncation selection and linkage disequilibrium/hitch-hiking effects but *not* genetic drift. Because of this, selection plays a stronger role in influencing the course of gene frequency change than one might imagine at first glance. The fact that a stable phenotype does not have to result means that there is almost always a residual of developmental noise, a not unrealistic outcome. As for the last point, it is an interesting question as to how an X allele would fare if it is introduced as a rare, newly-arisen variant, and we are looking into the possible consequences.

Finally, a comment regarding the meaningfulness of the distinction between *cis* and *trans* effects that we have tried to draw. Regulatory and structural loci are supposed to reside on different chromosomes and undergo independent assortment. However, just two of the potential four haplotypes are used after recombination. In effect, this means that the two chromosomes behave like a single chromosome that undergoes three crossovers upon mating, and one might legitimately wonder whether it does not follow that the structural and regulatory loci behave as one large chromosome. Indeed this is a valid point, and our response is to say that we use '*trans*' to mean sufficiently distant from a regulatory locus that it can most economically be imagined to act via a diffusible intermediate.

The mean population fitness at fixation decreases as the number of regulatory loci increases (Figure 7 *b*), and a certain minimum starting value of p_{1r} is necessary – under the conditions of Figure 7 *a*, this is about 0.06 – for the mean fitness to increase appreciably during evolution. From our simulations it turns out that the optimal number of regulatory loci N_r is just one, but this finding is a consequence of the special assumptions we have made, in particular the assumption that the $X \rightarrow 1$ transition probability at any structural locus is exactly equal to p_{1r}. With a different functional dependence, there could be a different outcome. For example, if a minimum of two regulatory gene products were needed to interact with structural sites for a $X \rightarrow 1$ transition to be possible, or if regulatory gene products interact with structural loci in a combinatorial fashion, the optimal number of regulatory loci N_r would be larger than 1.

By using 2^m coin-tossing trials instead of $(1/p_{1r})^m$ we implicitly ensure that the following holds good. When $p_{1r} > 0.5$, the population fitness is raised above the value it would have had in the absence of gene regulation, and when $p_{1r} < 0.5$ it is lowered below the corresponding value. Because of this the variance of fitness at equilibrium can be either lower (when $p_{1r} > 0.5$) or higher (when $p_{1r} < 0.5$) than the variance expected in the absence of regulatory loci (work in preparation).

Gene regulation can speed up evolution significantly: For example, at $f = 0.3$ and when there is no gene regulation, the average number of generations for fixation to be reached is 26.4 (ref. 16) whereas it is 9.5 in the presence of gene regulation (not shown)[18]. Other things being equal, gene regulation can accelerate the rate of evolution.

An increase in the frequency p_{xs} of the 'plastic allele' X from one generation to the next can take place only when regulatory genes are present, because in their absence p_{xs} stays approximately constant or decreases[16]. Thus a genetic capability for modifying the phenotype can itself act as a selective force in favour of alleles with potentially variable phenotypic effects. However, for p_{xs} to increase in the course of evolution, the initial value of p_{1r} has to be above a certain minimum and the number of regulatory loci N_r has to be below a certain maximum (see below).

Figure 5 shows that p_{1s} decreases with generation number, which appears counter-intuitive (because the value of p_{1s} should be a measure of fitness). The point is that concomitant with the decrease in p_{1s} there is a rapid increase in p_{xs} as well as of p_{1r}. Therefore the probability of an $X \rightarrow 1$ transition is high; and because of this, the mean fitness rises even as p_{1s} decreases. We mention that if the initial value of p_{1s} is sufficiently high, the 1 allele need not be eliminated (not shown).

Regulatory genes evolve although there is no direct selection acting on them. This is because selection does act, albeit indirectly, by favouring those combinations of regulatory alleles that predispose an increase in fitness (as determined by alleles in the structural loci). One might say that regulatory genes hitch-hike on structural genes. Hitch-hiking is very rapid to begin with, but slows down over the course of time (Figure 6 *a*). A consequence of hitch-hiking is linkage disequilibrium between the two sets of loci. Hitch-hiking is relatively ineffective when the number of regulatory loci is large, because then the rate of change in the frequency of alleles due to recombination in successive generations is relatively slow. Similarly, when the number of structural loci is small (much less than 18), fixation is reached so rapidly that there is no appreciable increase in frequency of the 1 allele at regulatory loci.

There are situations in which the existence of regulatory loci offers no evolutionary advantages. When $N_s = 18$ and $N_r = 26$, the mean fitness at equilibrium $\langle W_T \rangle$ (= 0.35) has the same value that it would have had in the absence of gene regulation[16]. As N_r increases further, $\langle W_T \rangle$ decreases, implying that there is no further advantage to be gained by increasing the number of regulatory loci. It has been pointed out that the optimal number of regulatory loci is $N_r = 1$; the higher the value of N_r, the lower is the value of $\langle W_T \rangle$ (Figure 4 *b*) and the slower the rate of evolution. When does an increase in N_r lower $\langle W_T \rangle$ to such an extent that it drops even below the value that it would have had in the absence of regulation? The answer depends on the number of structural loci. With V standing for that ratio of N_s to N_r below which no regulation is better than regulation, we find that $V = 1$ when $N_s = 10$ and V decreases to 0.04 as N_s increases to 25 (Table 1). Thus the number of regulatory loci that the system can 'tolerate' increases disproportionately with the number of structural loci.

When compared to a model in which a genetic locus can be switched 'on' or 'off' entirely at random, biasing the relative probabilities of switching with the help of regulatory genes offers mixed advantages. The rate of evolution is speeded up; but the population is not always as well adapted. A less-than-ideal level of adaptation is the price paid for rapid evolution. The most obvious advantage of the presence of 'plastic' alleles and regulatory genes is that purely phenotypic selection becomes effective[18].

Many years ago, C. H. Waddington showed that by carrying out phenotypic selection on a population of *Drosophila* he could generate true-breeding phenotypes that were absent in the starting population ('genetic assimilation')[9]. The phenotypes on which selection was performed were not caused by mutation but instead were generated as a response to specific treatments (e.g. raised temperature, ether) applied at critical stages of development. Waddington[9] suggested that the seemingly Lamarckian results of his experiments might be explained in conventional Darwinian terms by postulating that they were the outcome of indirect selection for regulatory genes which did not, by themselves, exert direct effects on the phenotype. The present work constitutes part of an ongoing exercise aimed at exploring Waddington's explanation in terms of an idealized computational model.

Artificial selection experiments suggest that plasticity is a trait that can evolve[1]. A range of conditions that favour the evolution of phenotypic plasticity also favours the development of cultural transmission ('the transfer of information by behavioural means')[20]. The picture of a highly plastic species with elaborate systems of cultural transmission seems to fit the human species rather well.

1. Scheiner, S. M., *Annu. Rev. Ecol. Syst.*, 1993, **24**, 35–68.
2. Gause, G. F., *Trans. Conn. Acad. Sci.*, 1947, **37**, 17–68.
3. Bradshaw, A. D., *Adv. Genet.*, 1965, **13**, 115–155.
4. Gupta, A. P. and Lewontin, R. C., *Evolution*, 1982, **36**, 934–948.
5. Via, S. and Lande, R., *Evolution*, 1985, **39**, 505–522.
6. West-Eberhard, M. J., *Annu. Rev. Ecol. Syst.*, 1989, **20**, 249–278.
7. Gommulkiewicz, R. and Kirkpatric, M., *Evolution*, 1992, **46**, 39–411.
8. Gavrilets, S. and Hastings, A., *Evolution*, 1994, **48**, 1478–1486.
9. Waddington, C. H., *Adv. Genet.*, 1961, **10**, 257–293.
10. Hinton, G. E. and Nowlan, S. J., *Complex Systems*, 1987, **1**, 497–502.
11. Nanjundiah, V. and Saran, S., *J. Biosci.*, 1992, **17**, 353–394.
12. Till, J. E., McCulloch, E. A and Siminovitch, L., *Proc. Natl. Acad. Sci. USA*, 1964, **51**, 24–36.
13. Nakahata, T., Gross, A. and Ogawa, M., *J. Cell. Physiol.*, 1982, **113**, 455–458.
14. Maynard Smith, J., *Evolution and The Theory of Games*, Cambridge University Press, Cambridge, 1982.
15. Whitehouse, H. L. K., *Towards an Understanding of the Mechanism of Heredity*, Edwin Arnold, London, 1973.
16. Behera, N. and Nanjundiah, V., *J. Theor. Biol.*, 1995, **172**, 225–234.
17. Behera, N. and Nanjundiah, V., *J. Theor. Biol.*, 1996, **178**, 135–144.
18. Behera, N. and Nanjundiah, V., *J. Theor. Biol.*, 1997, **188**, 153–162.
19. Bull, J., *Evolution*, 1987, **41**, 303–315.
20. Brandon, R. N., *Concepts and Methods in Evolutionary Biology*, Cambridge University Press, Cambridge, 1996.

ACKNOWLEDGEMENT. I thank V. Nanjundiah (who has been a collaborator) for numerous helpful discussions.