EVOLUTION
INTERNATIONAL JOURNAL OF ORGANIC EVOLUTION
PUBLISHED BY
THE SOCIETY FOR THE STUDY OF EVOLUTION

Vol. 58 February 2004 No. 2

Evolution, 58(2), 2004, pp. 211–221

ADAPTATION AND SPECIES RANGE

JOEL R. PECK[1,2] AND JOHN J. WELCH[1]
[1]*Centre for the Study of Evolution, School of Life Sciences, University of Sussex, Brighton BN1 9QG, East Sussex, United Kingdom*
[2]*E-mail: j.r.peck@sussex.ac.uk*

Abstract.—Phase III of Sewall Wright's shifting-balance process involves the spread of a superior genotype throughout a structured population. However, a number of authors have suggested that this sort of adaptive change is unlikely under biologically plausible conditions. We studied relevant mathematical models, and the results suggest that the concerns about phase III of the shifting-balance process are justified, but only if environmental conditions are stable. If environmental conditions change in a way that alters species range, then phase III can be effective, leading to an enhancement of adaptedness throughout a structured population.

Key words.—Adaptation, epistasis, group selection, shifting balance, species range.

Received June 3, 2003. Accepted September 15, 2003.

Some genotypes confer a relatively high level of fitness. However, because of epistasis (fitness interactions among loci) the alleles that constitute a high-fitness genotype may not confer high fitness when they occur in combination with other alleles. In this type of situation, we say that the high-fitness genotype is "coadapted."

Sewall Wright spent much of his working life investigating the evolutionary consequences of coadapted genotypes (Wright 1931, 1932, 1982). He suggested that superior coadapted genotypes can arise and spread through a population by means of a three-phase process, which he called the shifting-balance process. The shifting-balance depends on the population being subdivided into a number of demes (local breeding populations). In phase I of the process, a coadapted genotype with a relatively high fitness arises in one of the demes by means of genetic drift. (For brevity, we will simply call this genotype the "high-fitness genotype.") In phase II, selection causes the high-fitness genotype to become common within the deme where it arose. Finally, a process of interdemic selection sees the high-fitness genotype become common throughout the entire metapopulation (i.e., throughout the entire collection of demes).

Recently, Wright's theory has been the focus of substantial negative criticism (Barton 1992; Gavrilets 1996; Coyne et al. 1997, 2000). The critics conclude that problems with phase III, the process of interdemic selection, constitute "the central weakness of the theory" (Coyne et al. 1997). This echoes the thoughts of Haldane (1959) in one of the earliest critiques of the shifting-balance process.

The central difficulty with phase III, as recognized by Wright, is that individuals with a coadapted high-fitness genotype invading an area where other genotypes are common will find that, when they mate, their genomes are broken apart by segregation and recombination (Wright 1931, 1932, 1949, 1970, 1977, 1982; Barton and Hewitt 1989; Coyne et al. 1997). Thus, as a consequence of their genetic coadaptation, migrants with the high-fitness genotype may produce low-fitness offspring. This can prevent the spread of the high-fitness genotype.

In a defence of phase III, Crow et al. (1990) used a model designed to be unfavorable to the efficacy of the shifting-balance process. They concluded that phase III was remarkably successful, and that quite low levels of migration from the deme with the high-fitness genotype were sufficient to ensure its spread, even in the face of reverse migration. However, these results were reanalyzed and reinterpreted by Barton (1992). Barton's results suggested that, unless selection is very strong or the number of loci is very small, the outcome of evolution depends only weakly on any fitness advantage of the coadapted genotype, and more strongly on factors unconnected with adaptation. The most important of these factors are the dominance relationships of the alleles that comprise the coadapted genotype, and the relative rates of migration between the demes. Barton showed that excessive migration from a deme could lead to the spread of the genotype it contained, even if this genotype conferred a selective disadvantage. Crow et al.'s results, he suggested, were an example of migration overwhelming recombination *and* selection. Barton's conclusion was that while a phase III type of process could occur, the result was not likely to be adaptive (Barton 1992; Phillips 1993; Rouhani and Barton 1993).

In Wright's original formulation of the shifting-balance

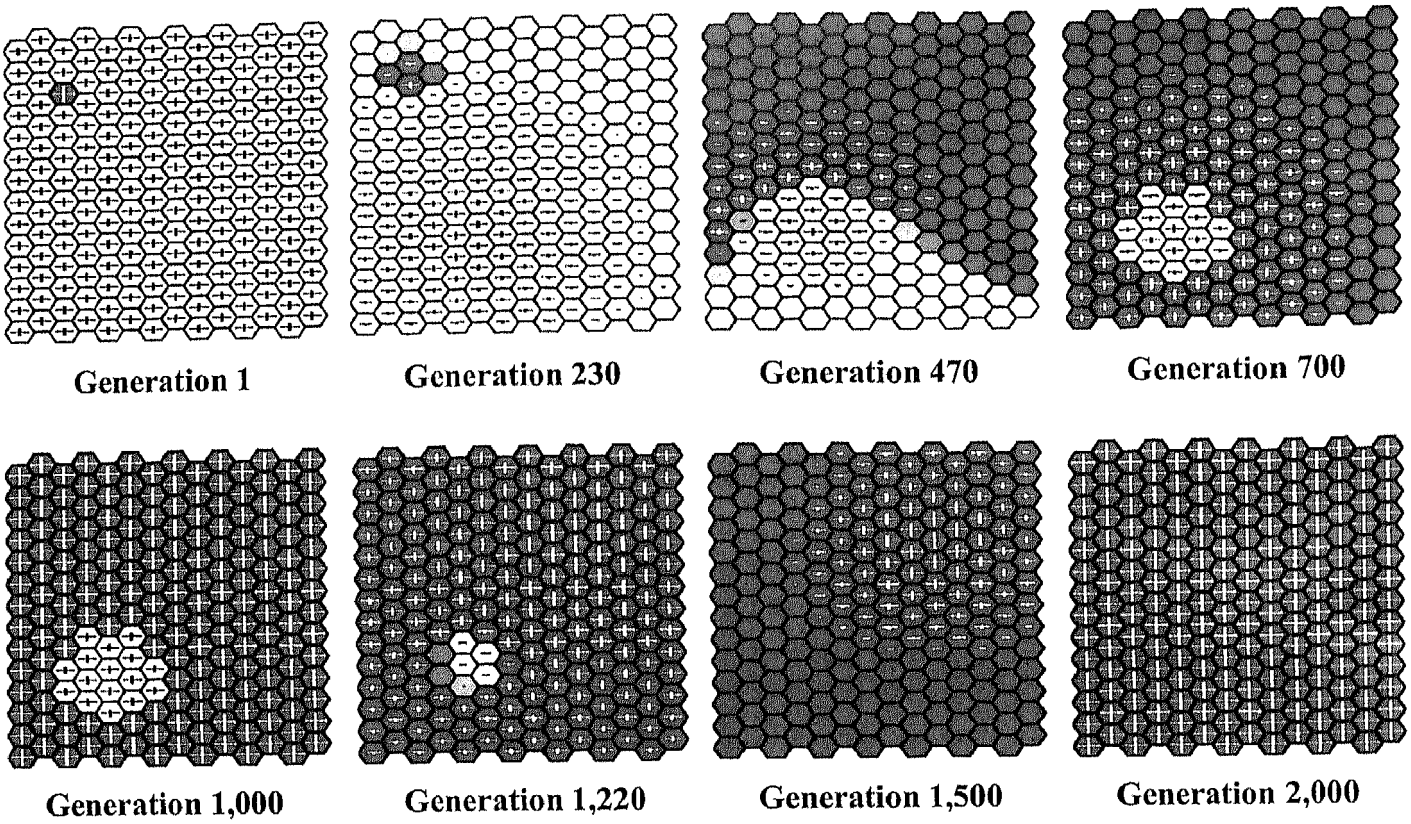

FIG. 1. Results from a numerical study showing the spread of a high-fitness coadapted genotype through a structured population. Each hexagon represents a deme within the metapopulation. The length of the vertical line within the territory of each deme is proportional to the population size of that deme, and the horizontal lines are proportional to the natural logarithm of population density, plus a constant. The color of each deme's territory depends on the proportion of *A* alleles in the deme, with the darkest demes consisting purely of the high-fitness genotype (which is an *AA* homozygote at all loci) and the lightest consisting purely of the wild type (which is a *BB* homozygote at all loci). Parameter values (see text and Appendix 1) were set as follows: $L = 3$, $m = 0.01$, $V_1 = 1$, $V_2 = 0.7$, and $V_3 = 0.4$. During the first 500 generations the environment deteriorates (i.e., the Θ_i values decline). There is a refuge centred at the fifth deme from the bottom of the metapopulation in the fifth column from the left side of the metapopulation. The degree of environmental deterioration increases with distance from the center of the refuge. During the next 500 generations the environment gradually improves, and the Θ_i values attain their original levels at generation 1001. A similar cycle of deterioration and improvement takes place during the last 1000 generations, but the position of the center of the refuge for these generations is the fifth deme from the top in the fifth column from the right side of the metapopulation. Additional details are given in Appendix 2.

process this was not considered a problem, because the differential migration necessary for the spread of the high-fitness genotype was assumed to result directly from its high fitness. High-fitness demes were assumed to export more individuals, allowing individual and interdemic selection to work in tandem.

Such a state of affairs requires that hard selection operates within demes, such that increasing the individual fitness of deme members increases the size of the deme. It also requires that increasing deme size lead to increased emigration from the deme (Wright 1982). This set of assumptions has been attacked on several grounds. First, the positive relationship between individual fitness and deme size, though not implausible, is by no means guaranteed (Prout 1980; Coyne et al. 1997). Second, as stressed by Gavrilets (1996), if the high-fitness genotype is initially present in a single deme that is surrounded by several other demes, then there is likely to be an inherent source of differential migration *into* the deme containing the high-fitness genotype resulting from this fact alone (Barton 1992; Gavrilets 1996). Any differences in deme size due to individual fitness would have to be extremely large to compensate. (This point had been obscured by a reliance on two-deme models when studying phase III: Crow et al. 1990; Barton 1992; Kondrashov 1992; Phillips 1993.) Third, while most previous authors have allowed the migration rates between the demes to enter as free parameters, results from explicit and plausible models coupling demography and fitness have suggested that the effect of individual fitness on deme size tends to be rather small (Rouhani and Barton 1987, 1993; Barton 1992; Gavrilets 1996). As a result, it is argued, interdemic selection is weak, and any differences in the migration rates that stem from other factors would be likely to determine which gene combination spreads. Again, this suggests that phase III is unlikely to be adaptive.

The criticisms mentioned so far concern phase III in isolation. However, there is another set of criticisms that concerns the shifting balance as a whole. These criticisms flow from the fact that the conditions required for each phase to succeed in isolation are quite different from one other, and in some cases seem mutually exclusive. For example, phase I requires that selection is weak enough to be overcome by drift, and yet phase II involves the operation of mass selec-

tion. Similarly, phase III requires migration of sufficient magnitude for the high-fitness genotype to spread, and yet phase I, the formation of this genotype, would seem to require a highly isolated deme. The few explicit models of the entire shifting-balance process, though they show that the process can work, do indeed suggest that parameter values must be very finely balanced (Rouhani and Barton 1993; Moore and Tonsor 1994; Bergman et al. 1995; Gavrilets, 1996; Coyne et al. 1997).

In this work, we demonstrate that the force of some of these criticisms depends on the existence of stable environmental conditions. In particular, our results suggest that, in times of environmental change, phase III can be adaptive, particularly when such change produces alterations in species range. Furthermore, at such times of change, phase III can be effective even if the migration rate is extremely low, as required for the between-deme genetic differentiation that characterises phase I of the shifting-balance process.

To see why changes in species range can lead to the spread of high-fitness genotypes, it helps to consider a situation like the one pictured in Figure 1 (generation 1), where a high-fitness genotype is common in one deme, and where this initial-high-fitness deme is not on the edge of the metapopulation. If the environment deteriorates in the region of the initial-high-fitness deme, then we can expect this deme to persist longer than its neighbors, which are populated by the wild type, a genotype that confers a lower level of fitness in any given environment. Once the wild type goes extinct (or nearly extinct) in the vicinity of the initial-high-fitness deme, the high-fitness genotype can spread into the territory of neighboring demes without too much mating with other genotypes. Of course, this success may be short lived, as the environment may continue to deteriorate to the point where the species becomes rare (or extinct) in the vicinity of the initial-high-fitness deme. However, if environmental decay is slow enough, the high-fitness genotype may spread beyond the doomed region before local extinction occurs. This can allow the high-fitness genotype to spread widely when the environment improves again, or when the species invades new territory. To test the logic of this idea and to explore its consequences, it is useful to study a mathematical model.

Model

Consider a diploid and hermaphroditic organism that lives in a subdivided metapopulation. Each deme occupies a hexagonal area that is adjacent to the areas occupied by six other demes, unless it is on the edge of the metapopulation (as shown in Fig. 1). We will consider the case where the population size of each deme is always sufficiently large that genetic drift is negligible. (This is in line with the majority of previous studies, which treat phase III as a deterministic process: Crow et al. 1990; Barton 1992; Kondrashov 1992; Phillips 1993; Gavrilets 1996.) Let us number the demes as 1,2,3,. . . and let N_i represent the density of adults in deme i (i.e., N_i is the average number of adults per unit area in deme i). Generations are discrete, so that all adults in the population die at about the same time as the offspring are born.

Adults produce offspring by means of female effort, which means contributing one gamete to an offspring, plus most of the resources required for the offspring to become independent. Adults also reproduce by means of male effort, which means contributing one gamete to an offspring, but little in the way of resources. The adult chosen to contribute to a given offspring by means of male effort is selected at random from among all the adults in the deme concerned. Gametes form by means of normal Mendelian segregation and recombination, and there is free recombination between all loci that contribute to variation in fitness. Fertility is independent of genotype, but depends on the density of adults within each deme and on the quality of the environment occupied by the deme. In particular, the expected number of offspring born to an adult via female effort within deme i is denoted by $\Theta_i F(N_i)$. The first component of this birth rate, Θ_i, is an indicator of the environmental quality of deme i. We choose Θ_i such that $0 \leq \Theta_i \leq 1$. The second component, the function $F(N_i)$, determines the dependence of fertility on population density. To provide general results, we leave this function unspecified, but we will make some assumptions about its form.

Let us assume that $F(N_i)$ is a continuous and monotonically decreasing function that tends toward zero as N_i tends toward infinity. As a result, birth rate declines as population density increases; that is, we have negative density dependence. Let us also assume that the quantity $N_i F(N_i)$ is an increasing function of N_i, so that the number of offspring produced in a deme will increase as the number of parents increases. These assumptions about $F(N_i)$ are not very restrictive, and a wide variety of plausible birth-rate-regulation mechanisms satisfy these requirements. Note that under these assumptions, the maximum possible birth rate is achieved in the limit as N_i goes to zero. Because this maximum birth rate will prove to be important, we introduce the following notation.

$$\lim_{N \to 0} F(N) = F_0, \tag{1}$$

where F_0 is a finite positive constant.

Immediately after birth, offspring undergo viability selection. We assume that viability is determined by L loci (where $L \geq 2$) . At each of these L loci there are two possible alleles (denoted as A and B) . We follow the selection scheme of Crow et al. (1990) and Barton (1992). Under this scheme, the viability of any given individual takes one of three possible values. These values are denoted as V_1, V_2, and V_3, where $V_1 \geq V_2 \geq V_3$. The highest of these three values, V_1, is taken to be the probability of surviving viability selection for individuals that are AA homozygotes at all L loci (this is the high-fitness genotype). The second highest value, V_2, is the probability of survival for individuals that are BB homozygotes at all L loci (this genotype is the wild type). We assume that the wild type is entirely dominant to the high-fitness genotype. Thus, any genotype that has at least one B allele at every locus is assumed to have a viability of V_2, just like the wild type. All other genotypes (which are AA homozygotes as some, but not all loci) have a probability of surviving viability selection given by V_3. Thus, V_3 represents a valley in between the fitness peaks constituted by the high-fitness genotype and the wild type.

Just before maturation, some of the surviving juveniles emigrate from the demes in which they were born. The probability of attempting emigration for each surviving juvenile is given by m (where $0 < m < 1$). Emigrants enter into the territory of another deme by crossing a randomly selected side of the territory of their home deme. If no deme exists across the side (because the deme is on an edge of the metapopulation), then the juvenile attempting emigration returns to its home deme. After migration, mating takes place, and the life cycle begins again.

We assume that the metapopulation is initially in the situation pictured in Figure 1 (generation 1) where the A allele (and the high-fitness genotype) is initially present only in the initial-high-fitness deme (which has six neighbors) and absent in all other demes.

Note that several features of the model are designed to be actively unfavorable to the success of phase III. First, the alleles forming the high-fitness genotype are fully recessive, and this hinders their spread (Crow et al. 1990; Barton 1992; Phillips 1993; Gavrilets 1996). Second, we have assumed free recombination, and this undermines the integrity of the high-fitness genotype (Crow et al. 1990; Phillips 1993; Gavrilets 1996). Third, the high-fitness genotype is initially present in a single deme that is surrounded by six other demes (Barton 1992; Gavrilets 1996). Finally, migration takes place after selection (Crow et al. 1990; Phillips 1993). All of these points militate against the success of phase III, and as such we offer a conservative test of its plausibility. This is in line with the general strategy of Crow et al. (1990) and Kondrashov (1992).

Evolutionary Outcomes When There Is No Environmental Variation

What will be the evolutionary fate of the high-fitness genotype? In general, this question is difficult to answer without the use of a computer simulation. However, there are some special cases that are both tractable and illuminating. Let us begin with the case where, at any given time, all of the Θ_i values are exactly the same. This means that, at any given time, the quality of the environment does not vary from one part of the metapopulation to another (except for those aspects of the environment that depend on local population density).

Say that the quality of the environment throughout the metapopulation (the Θ_i values) begins to decline slowly. Let us arbitrarily choose a number, denoted as z, that lies between V_1 and V_2 ($V_1 > z > V_2$). We shall assume that the value of m is sufficiently small so that, while $\Theta_i \geq 1/(zF_0)$, the high-fitness genotype will be preserved within the initial-high-fitness deme in substantial numbers, and with a non-negligible frequency. (Note that, for any set of initial conditions and any allowable value of z, it is always possible to choose a nonzero value of m sufficiently small to satisfy this requirement so long as the high-fitness genotype is initially sufficiently common in the initial-high-fitness deme; Barton 1992; Gavrilets 1996.) Finally, let us assume that the value of m satisfies the following inequality:

$$m < 1 - \frac{z}{V_1}. \tag{2}$$

The case just specified is inclusive of situations where the value of m is small, which is important for addressing some of the criticisms of the shifting-balance theory (Barton and Hewitt 1989; Crow et al. 1990; Coyne et al. 1997).

In the case specified, the high-fitness genotype will eventually become fixed throughout the metapopulation and the B allele will be completely eliminated. This will occur so long as the decline in the Θ_i values is sufficiently slow. Elimination of the B allele will take a very long time because the metapopulation is assumed to contain a very large number of individuals. Nevertheless, elimination is certain once the B allele is sufficiently rare in every deme. Furthermore, as long as the environmental decline is sufficiently slow, this point of no return will be reached before the Θ_i values decline to the point where they are equal to $1/(zF_0)$.

To understand these results, it is worthwhile to consider a deme where $\Theta_i < 1/(V_2F_0)$. For each of the wild-type adults in such a deme the expected number of offspring produced by female effort that will survive to reproductive age is less than one. Because of this low rate of successful reproduction, the B allele must eventually disappear from any deme where $\Theta_i < 1/(V_2F_0)$, unless the B allele is maintained by immigration. In other words, all such demes are demographic sinks (e.g., Shmida and Ellner 1984; Holt 1996). Thus, once the universal value of Θ_i falls below $1/(V_2F_0)$, the B allele will undergo a long-term decline in density throughout the metapopulation, and it will eventually become vanishingly rare in all demes. At the same time, in the initial-high-fitness deme, we know (by assumption) that the high-fitness genotype will remain common, at least until the universal value of Θ_i falls below $1/(zF_0)$. That is, the high-fitness deme will act as a demographic source. If the decline in the Θ_i values is sufficiently slow, then we can ensure that the B allele falls to an arbitrarily low density in every deme, while the high-fitness deme is still able to act as a source. Thus, eventually, because the high-fitness deme is a demographic source and the surrounding demes are demographic sinks, the proportion of individuals in surrounding demes that are immigrants from the high-fitness deme can be increased to arbitrary levels. This guarantees that the high-fitness genotype will become common in all six demes that are adjacent to the high-fitness deme (Barton 1992). (The assumption embodied by inequality [2] ensures that this will happen despite the loss of migrants from these demes.) The demes adjacent to the initial-high-fitness deme can then act as sources for the colonization of new demes by the high-fitness genotype, and this process eventually leads to the takeover of the entire metapopulation by the high-fitness genotype.

Note that the foregoing results involved two requirements on m. The first requirement is that m is sufficiently small so that, while $\Theta_i \geq 1/(zF_0)$, the high-fitness genotype will be preserved within the initial-high-fitness deme in substantial numbers, and the second requirement is that inequality (2) holds. It is worth noting that, for a wide range of parameter values, the second of these requirements will be satisfied if the first is satisfied. To see why, it is important to recognize that the greatest threat to preservation of the high-fitness

genotype within the initial high-fitness deme arises when the population densities of the demes surrounding the initial-high-fitness deme are similar to the density of the initial-high-fitness deme. Typically, this situation will occur prior to the decline in the universal value of Θ_i. A relevant approximation has been derived by Barton (1992), and this suggests that the high-fitness genotype will be maintained at a substantial density in the initial-high-fitness deme as long as:

$$m < \frac{1 - (V_3/V_2)}{2eL}, \tag{3}$$

where e is the base of natural logarithms. This approximation requires that selection against the unfit hybrids is not too strong ($V_3/V_2 > 0.9$), the number of loci is not too small ($L > 3$), and that differences in the densities of adjacent demes are not too large (the derivation assumes equally sized demes). Comparison of inequalities (2) and (3) shows that, for a wide range of parameter values, inequality (2) will be satisfied if m is small enough to satisfy inequality (3). Thus, if the high-fitness genotype can survive within the initial-high-fitness deme while environmental conditions are favorable, then, typically, the high-fitness genotype will take over the metapopulation during a sufficiently slow decline in the quality of the environment.

Evolutionary Outcomes When Environmental Variation Is Allowed

The preceding analysis may provide some insight, but it is not biologically reasonable to expect a total absence of between-deme variation in the aspects of environmental quality that are described by the Θ_i values. Furthermore, the lack of environmental variation makes the metapopulation prone to total extinction. In particular, if the universal value of Θ_i drops below $1/(V_1F_0)$, then the size of the metapopulation will fall during every generation. On the other hand, if between-deme variation in the Θ_i values is allowed, then, during a global environmental decline, the metapopulation can resist extinction so long as the degree of environmental variation is sufficient to ensure that at least one deme is sufficiently hospitable to allow the survival of a stable local population.

Our preliminary studies showed that the spatial scale of environmental variation is a critical determinant of evolutionary outcomes. A simple way to characterize the scale of environmental variation is to specify the minimum distance between two demes that one must travel to observe a certain amount of environmental variation. In particular, let us calculate the absolute value of the difference in Θ_i values between the initial-high-fitness deme and every other deme that is within a distance d of the initial-high-fitness deme. (A particular deme is considered to be within a distance d of the initial-high-fitness deme if some (or all) of its territory is within a distance d of the center of the territory of the initial-high-fitness deme.) Let Δ represent the largest of these differences in Θ_i values. If the Θ_i values change over time, then the value of Δ depends on the moment in time at which the Θ_i values are measured. We will assume that, regardless of when the measurement is made, Δ is always less than a certain value, which we denote as Δ_{max}.

The results of the model can be quite complicated if Δ_{max} is too large. We will therefore assume that:

$$\Delta_{max} < \frac{V_1 - V_2}{2V_1V_2F_0}. \tag{4}$$

We also assume that that the value of m is sufficiently small so that if, in the initial high-fitness deme, we have $\Theta_i \geq (1/V_2F_0 - \Delta_{max})$, then the high-fitness genotype will definitely be preserved within the initial-high-fitness deme in substantial numbers, and with a nonnegligible frequency. Finally, let us assume that:

$$m < 1 - \frac{V_2}{V_1 - 2\Delta_{max}V_1V_2F_0}. \tag{5}$$

Note that it always possible to find a nonzero value of m sufficiently small to satisfy these requirements (Barton 1992; Gavrilets 1996). Also, note that these assumptions are relaxed in the simulation studies presented below.

With these assumptions in place, let us consider the immediate neighborhood of the initial high-fitness deme (or the immediate neighborhood, for short). The immediate neighborhood is defined to include all demes that have at least some territory within a distance, d^*, of the center of the territory of the initial-high-fitness deme. The value of d^* may be chosen arbitrarily, as long as it is finite and sufficiently large so that the immediate neighborhood contains multiple demes.

What will happen if there is a global environmental deterioration, such that there is a decrease in all of the Θ_i values within the metapopulation? Let Θ_H denote the value of Θ_i in the initial-high-fitness deme. The preceding assumptions imply that, if the decline in Θ_H is sufficiently slow and the value of d is sufficiently large, then the high-fitness genotype will become fixed (or nearly fixed) in all demes that are in the immediate neighborhood of the initial-high-fitness deme. (Here, near fixation means that a tiny frequency of B alleles persists as a consequence of migration from outside of the immediate neighborhood.) The takeover of the immediate neighborhood by the high-fitness genotype will occur before Θ_H becomes low enough so that $\Theta_H = (1/V_2F_0 - \Delta_{max})$. Note that this is true regardless of the size of the immediate neighborhood (the radius of which is denoted by d^*).

The reasoning behind these results is similar to the reasoning for the preceding case, which allowed for no between-deme environmental variation. The extended neighborhood is the collection of demes that have some territory within a distance d of the center of the initial-high-fitness deme. Because of the assumed upper limit on the degree of environmental variation within the extended neighborhood, we can be sure that, before Θ_H becomes low enough so that $\Theta_H = (1/V_2F_0 - \Delta_{max})$, it must be the case that $\Theta_i < 1/(V_2F_0)$ for every deme within the extended neighborhood. This observation has implications for the fertility of adults that have a genotype that contains at least one B allele, and that live in a deme within the extended neighborhood. In particular, for each of these adults, the expected number of offspring produced by female effort that will survive to reproductive age will fall below one before Θ_H becomes low enough so that $\Theta_H = (1/V_2F_0 - \Delta_{max})$. Because of this low rate of successful

reproduction, the B allele must disappear from all demes within the extended neighborhood, unless this allele is sustained by inflow from outside of this area. By increasing the value of d, we can reduce, to an arbitrarily low level, the flow of B alleles to the immediate neighborhood. We know that the high-fitness genotype can survive in substantial numbers in the initial-high-fitness deme as long as $\Theta_H \geq (1/V_2F_0 - \Delta_{max})$. The initial-high-fitness deme can thus provide a source of the high-fitness genotype for the immediate neighborhood. Furthermore, the inequalities (4) and (5) collectively imply that, as long as $\Theta_H \geq (1/V_2F_0 - \Delta_{max})$, we have $\Theta_iF_0V_1(1 - m) > 1$ in all demes within the extended neighborhood. This means that the high-fitness genotype can be expected to increase to substantial numbers in all demes within the extended neighborhood, except where the B allele is not very rare. As the B allele can be made arbitrarily rare in the immediate neighborhood by increasing the value of d, we can be sure that the high-fitness genotype will rise to substantial numbers in all demes within the immediate neighbourhood if d is sufficiently large, and if the decline in Θ_H is sufficiently slow.

Once the high-fitness genotype has spread throughout the immediate neighborhood of the initial-high-fitness deme, the spread may continue to new areas that are adjacent to this neighborhood. However, the spread will stop if the high-fitness genotype approaches a new neighborhood where the Θ_i values are too low to sustain the high-fitness genotype. In this case, the spread can continue if the environment in the new neighborhood improves sufficiently (i.e., if there is a sufficient increase in the Θ_i values). The spread may also stop if the high-fitness genotype approaches a neighborhood where the wild type is present in substantial numbers. In this case, the spread may continue if the local environment deteriorates (i.e., a fall in the Θ_i values). Thus, over the long term, the high-fitness genotype may take over the entire metapopulation, if this takeover is enabled by changes in environmental conditions.

If, during a global environmental decline, the value of Θ_i remains quite high in some areas (above $1/[V_2F_0]$), then these areas may act as refuges where the wild type is preserved. However, if the high-fitness genotype spreads to most areas of the metapopulation during the environmental decline, then these refuges will tend to be surrounded by demes in which the high-fitness genotype is common. After recovery of the environment, this situation can make it impossible for the wild type to escape from its refuges in the event of a second environmental decline. If the demes that retain the highest Θ_i values are different in the second environmental decline, as compared to the first, then the wild type may go extinct, as the now-inescapable refuges in which it survived become uninhabitable. A numerical study, which illustrates the sequence of events just described, is presented in Figure 1 (see also Mitteldorf and Wilson 2000). Appendices 1 and 2 give details of all the simulations presented in this work.

Computer Simulation Studies

The process we have described can amplify the evolutionary consequences of viability differences between genotypes. This amplification, in turn, can lead to an adaptive phase III of the shifting-balance process. However, it is clear that an environmental decline could also amplify nonfitness-based differences between demes. If a deme has some intrinsic environmental advantage, then it will have an enhanced chance of remaining a demographic source longer than neighboring demes, regardless of the genotypes of its members. Factors that might give some demes an intrinsic advantage are frequently observed in natural habitats, and they include a relatively high input of resources, superior shelter, and a low density of predators.

Our analytic treatment does allow for intrinsic differences between demes (which are reflected by variation in the Θ_i values). However, to make analytic progress, we placed limits on the degree of variation among the Θ_i values within any given area. To consider the impact of between-deme differences that are larger than allowed by the analysis, we have made use of computer simulations. Simulations also allow us to consider other questions, such as the length of time that it takes for a substantial spread of the high-fitness genotype through the metapopulation, and the effect of environmental changes that alter the position of the demes with the most favorable habitats.

To provide a mathematical description of between-deme differences in environmental quality, we assume that the Θ_i values are determined by two quantities. The first of these is $C(t)$, which is a function that describes time-dependent changes in environmental quality (t gives the number of generations since the start of a particular simulation trial). The second quantity that determines the values is Q_i. The value of Q_i depends on the particular deme, i, in question and represents the intrinsic environmental quality of that deme. We express the Θ_i values as follows:

$$\Theta_i(t) = C(t)Q_i, \tag{6}$$

where $\Theta_i(t)$ represents the value of Θ_i during generation t, and both $C(t)$ and Q_i are nonnegative.

Before each simulation trial the Q_i values are chosen at random from a probability distribution that is described in detail in Appendix 1. The crucial point about this distribution is that it depends on a single parameter, ν (where $\nu > 1$). This parameter acts as a simple and useful measure of the environmental differences between adjacent demes. Specifically, ν measures the environmentally induced differential migration expected between a pair of randomly chosen demes. This follows from the following relationship (which is derived in Appendix 1):

$$E\left(\left|\ln\left(\frac{\Theta_i(t)}{\Theta_j(t)}\right)\right|\right) = \ln(\nu), \tag{7}$$

where i and j denote two randomly chosen demes, and $E(\bullet)$ denotes statistical expectation.

Using ν to measure environmental variation allows us to make quantitative comparisons with the fitness differences between genotypes. Indeed, ν compares directly with the relative fitness of the higher peak, denoted $w_1 \equiv V_1/V_2$. So, if we were to take $\nu = w_1$, then we could plausibly claim that selection and environmental variation were of equivalent strengths.

In each simulation trial the metapopulation consists of 100 demes, with 10 demes on each side of the metapopulation.

Trials were initiated with a single deme fixed for the A allele (the initial-high-fitness deme). The position of this deme was chosen at random from among the demes that are not on the edge of the metapopulation. All other demes were fixed for the B allele at all loci (the wild type). The initial population densities in all demes were set to the equilibrium values that would be achieved in the absence of migration, under optimal environmental conditions (i.e., when $C[t]$ is at its maximum value) in all demes.

Each simulation trial consisted of 10 cycles of environmental decline and recovery. During each of these cycles $C(t)$ starts at its maximal value, and remains there for 50 generations. The value of $C(t)$ then declines for 100 generations until it reaches a nadir. The value of $C(t)$ remains at this nadir for T generations (where T is a positive integer). Next, the value of $C(t)$ rises again over the course of 100 generations and stays at its maximal value for another 50 generations. This completes one environmental cycle (see Appendix 1). After each of these cycles the Q_i values are reassigned, using the same method that was used at the beginning of the simulation trial (with one exception, as noted below). This means that there are occasional changes in the intrinsic quality of demes. These changes can be expected as a consequence of localized changes in environmental quality caused by changes in the spatial distributions of, for instance, predators, disease organisms, and resources. Of course, changes of this sort are unlikely to occur everywhere at once, as in our simulation. However, by making changes in the Q_i values all at once, it is possible to separate the effects of this sort of environmental change from the global changes represented by changes in the value of $C(t)$.

Let p_A represent the frequency of A alleles in the metapopulation. In other words, if we consider all the alleles present in all individuals in all demes and at all of the loci that are under selection, p_A is the proportion of these alleles that are A alleles. At the start of each trial p_A is always approximately equal to 0.01. This makes sense, as we start with one deme fixed for A alleles at all of the loci under selection, and there are 100 demes. (The value of p_A is never exactly equal to 0.01 because of variation in the Q_i values, and because, when $V_1 \neq V_2$, the initial population densities differ between the high-fitness deme and other demes.) We record the value of p_A at the end of each environmental cycle. Thus, for each trial, we record the value of p_A at 10 different times (see Tables 1, 2).

The model is unavoidably complex, with a substantial number of parameters, and so a systematic exploration of parameter space would be computationally prohibitive. Instead, we use a standard set of parameter values (given in the figure captions and tables), and we examine the effects of varying these. Our studies focus on the effects of altering the relative strengths of selection and the effects of small-scale environmental variation. We also explore the effects of changing migration rate, m, and the number of generations during each environmental cycle that the environment is at its worst, T. The rationale for our parameter choices and full details of the birth function, $F(N)$, and other more technical aspects of simulation procedures can be found in Appendix 1.

Simulation results

For each choice of parameter values we ran 1000 simulation trials. Each set of 1000 simulation trials with the same parameter values will be called a simulation set. Fig. 2 shows results from a simulation set using the standard parameter values. Fig. 2a shows a histogram of p_A after a single environmental cycle. Clearly, the high-fitness genotype approaches extinction in a substantial fraction of trials and attains a wide range of frequencies in the remaining trials. The substantial rate of loss of the high-fitness genotype is unsurprising given that we have chosen the case most unfavorable to the spread of this genotype in several respects. Although the eventual frequency attained is small in many of the remaining cases, the full significance of this partial spread can be gauged by considering what happens during multiple cycles of decline and regeneration. To see this, compare Figure 2a with Figure 2b, which shows data from the same simulation trial after two environmental cycles. As can be seen, an additional cycle of environmental decline and recovery generally allows the high-fitness genotype to spread further through the metapopulation, at least in those cases where there was some spread in during the initial cycle. Figures 2c and 2d show the corresponding data after three and four environmental cycles, respectively, and it can be seen that the trend continues until, after four cycles, the high-fitness genotype is either nearly lost or nearly fixed in almost every case.

Given the pattern of the data shown above, in what follows, it will most useful to report two statistics. First, we report the proportion of trials in which the high-fitness genotype was preserved with appreciable frequency (defined as the proportion of trials for which $p_A \geq 0.0005$ after 10 environmental cycles). This statistic will be denoted Y. Second, we report the mean frequency attained by the A allele in those trials (i.e., the mean value of p_A in trials for which $p_A \geq 0.0005$ after 10 cycles). This statistic will be denoted $\overline{p_A^*}$. In the simulation set using the standard parameter values (shown in Fig. 2) $Y = 0.450$. After one cycle, $\overline{p_A^*} = 0.360$, while after 10 cycles, $\overline{p_A^*} = 1.000$.

As explained above, the Q_i values change after each environmental cycle. It is reasonable to ask whether these changes make much difference. With this in mind, we ran a simulation set using the standard parameter values (as above), but we did not change the Q_i values after they were set at the beginning of each trial. At the end of 10 cycles, this resulted in $Y = 0.424$ and $\overline{p_A^*} = 0.517$. Comparison with the results reported above shows that, if the Q_i values do not change, then repeated environmental cycles allow for some further spreading of the high-fitness genotype, but this effect is not dramatic.

In Table 1 we examine the effects of changing the selective advantage of the high-fitness genotype, w_1, and the degree of variation in the intrinsic quality associated with demes, v. We report Y and $\overline{p_A^*}$ after one, five, and 10 cycles. In general, the results show that increasing w_1 tends to increase both Y and $\overline{p_A^*}$, while increasing v tends to decrease both of these quantities. The results also show that, even when w_1 is only one-third the size of v (selection much weaker than variation in deme-habitat quality) the frequency of the high-fitness

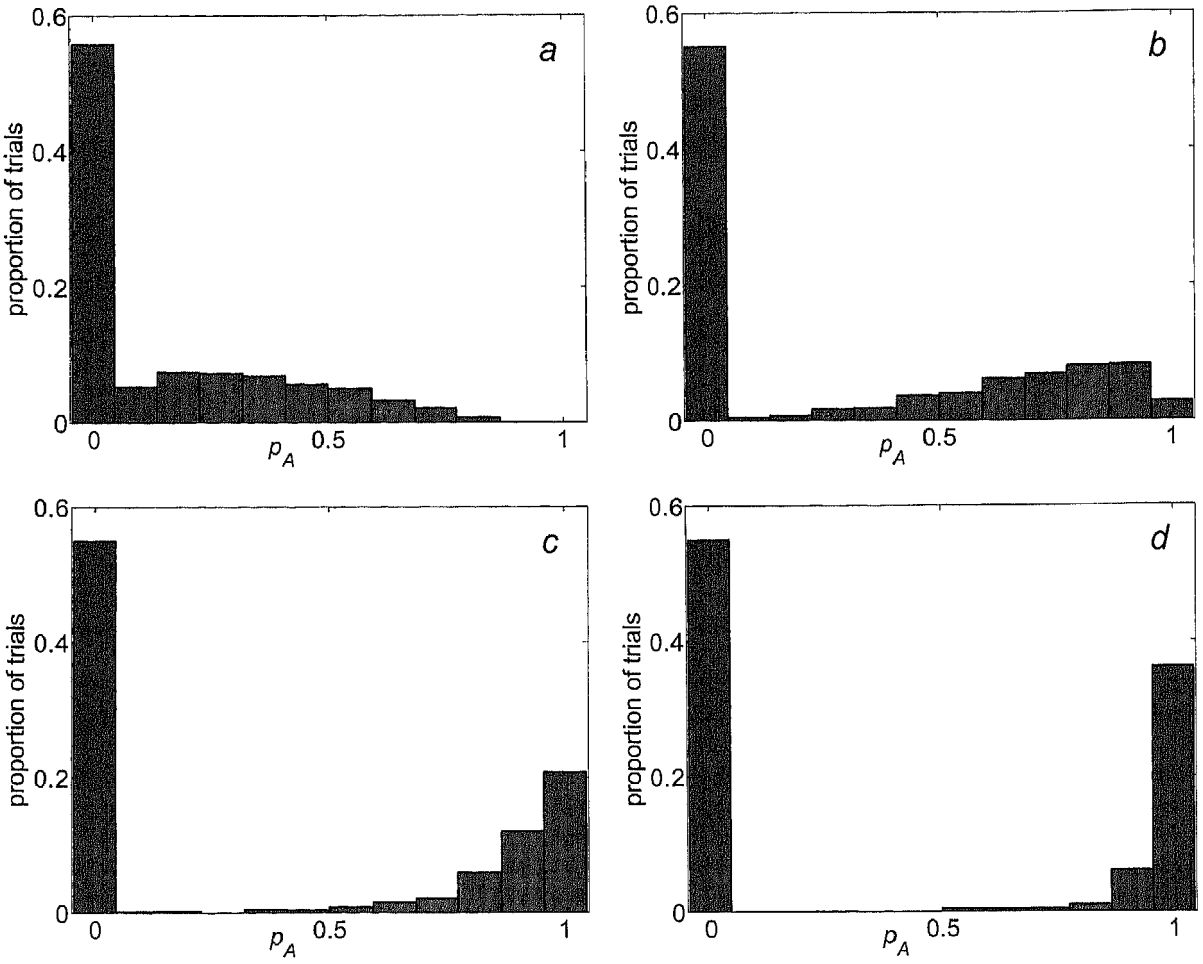

FIG. 2. Histograms of the frequency of the A allele within the metapopulation, p_A, in 1000 replicate trials. (a) The results after a single cycle of environmental decline and regeneration; (b, c, d) The results after two, three, and four complete cycles (each commencing immediately after the last). A new set of environmental qualities, the Q_i values, were chosen for each cycle (see text). For all plots, the parameter values chosen were as follows. The metapopulation contained 10×10 demes, $L = 5$ loci affect fitness, the migration rate was $m = 10^{-3}$, the viability of the wild type was $V_2 = 0.7$, the maximum birth rate was $F_0 = 4$, and all cycles included $T = 1000$ generations spent at the environmental nadir. The relative fitnesses of the high-fitness genotype and the unfit recombinants were $w_1 = 1.10$ and $w_3 = 0.9$. The Q_i values were generated at random, using the parameter $v = 1.10$. As such, selection in both directions and the strength of environmental variation were all set at 10%. In addition, the depth of the environmental cycle was set so that 10% of demes populated by the wild type were expected to remain demographic sources at the environmental nadir. Further details are given in the text and in Appendix 1.

genotype in the metapopulation can reach quite high levels, given a sufficient number of environmental cycles. However, when there is no advantage to the high-fitness genotype ($w_1 = 1$) the high-fitness genotype persists in less than 2% of trials. On the rare occasions when the high-fitness genotype does persist, it tends to reach only intermediate frequencies, even after 10 environmental cycles. When the high-fitness genotype was associated with a selective disadvantage ($w_1 = 0.95$), the high-fitness genotype was lost from the metapopulation in all 1000 trials. (Of course, for $w_1 \leq 1$ the term "high-fitness genotype" is a misnomer.)

Next, let us consider how the duration of an environmental cycle affects the results. Preliminary studies suggested that the largest effects of environmental-cycle length are achieved by changing the amount of time that the population spends at the environmental nadir (when environmental conditions are at their worst). As stated above, this period of time is denoted by T. We noticed that the effect of manipulating T depended on the rate of migration, m. Therefore, in Table 2 we provide data on the joint effects of altering T and m.

As can be seen from Table 2, neither changing the value of m nor changing the value of T has much effect on Y. However, increasing the value of m or T seems to increase $\overline{p_A^*}$. This makes sense, as increasing m means that the high-fitness genotype can move through the metapopulation more easily, and increasing T means that there is more time for movement of this sort.

DISCUSSION

The analytic and simulation results presented in this study clearly demonstrate that phase III of Wright's shifting-balance process can lead to enhanced adaptation within a metapopulation. The key ingredient that seems to be missing from

TABLE 1. Results from simulation trials. Four statistics are listed for each parameter-value combination. The first, shown in parentheses, is the proportion of runs in which the high-fitness genotype was preserved at substantial frequency after 10 complete environmental cycles (the statistic denoted Y in the text). The remaining three numbers (from top to bottom and left to right) show the mean frequency of the A allele achieved in these runs, $\overline{p_A^*}$, after one, five, and 10 complete environmental cycles. The table shows results for various values of $w_1 = V_1/V_2$ (the relative fitness of the high-fitness genotype) and v (the measure of environmental variation). Other parameters in all cases are the standard set listed below Figure 2 and in Appendix 1.

	v		
w_1	1.15	1.10	1.05
1.20	(0.572) 0.642 1.000 1.000	(0.800) 0.906 1.000 1.000	(1.000) 1.000 1.000 1.000
1.15	(0.452) 0.418 0.993 1.000	(0.631) 0.668 1.000 1.000	(1.000) 0.993 1.000 1.000
1.10	(0.341) 0.252 0.916 1.000	(0.450) 0.360 0.992 1.000	(0.792) 0.736 1.000 1.000
1.05	(0.184) 0.164 0.586 0.900	(0.265) 0.178 0.740 0.984	(0.454) 0.247 0.976 1.000
1.00	(0.011) 0.132 0.298 0.259	(0.013) 0.132 0.252 0.326	(0.017) 0.132 0.312 0.259
0.95	(0.000) —	(0.000) —	(0.000) —

most previous treatments of the process is the occurrence of occasional, widespread, temporary deterioration in environmental conditions. However, environmental deterioration is only effective if the coupling between fitness and demography allows local isolated populations to crash towards zero population density when fitness falls below a certain level. Such a coupling is congruent with much recent work on the coevolution of demography and fitness (e.g., Gomulkiewicz and Holt 1995; Holt 1996; Kirkpatrick and Barton 1997; Gomulkiewicz et al. 1999; Ronce and Kirkpatrick 2001). Furthermore, there are many potential causes of environmental deterioration of the sort required, some of which will affect many species, while others will affect only a small number of species. These causes include changes in climate, the evolution of more virulent parasites, and an increase in the density of predators.

Some demographic models are unlikely to yield the sort of results presented here. For example, a model by Rouhani and Barton (1993; see also Barton and Rouhani 1987, 1993) does not allow any local population to become a demographic sink, no matter how low mean fitness becomes. This difference between the models explains why the conclusions drawn by Rouhani and Barton (1993) about the efficacy of phase III, differ from those drawn here.

It is intuitively (and mathematically) clear that widespread environmental deterioration can cause the total extinction of a species, unless there are favorable areas where the deterioration is not so severe. These favorable areas can act as refuges where the species can persist when environmental conditions are at their worst. However, the refuges can also act as places where inferior genotypes are preserved so that these genotypes will remain in the metapopulation after environmental conditions improve again. Thus, the spread of superior genotypes during times of environmental decline may be incomplete if there is substantial small-scale spatial variation in environmental quality. Nevertheless, if different episodes of environmental decline occur for different reasons, it is plausible that there will be changes in the position of the most favorable local habitats between one such episode and the next. Our results suggest that these changes can allow a superior genotype eventually to outcompete other genotypes throughout a metapopulation. This conclusion is in accord with results produced by Mitteldorf and Wilson (2000).

The impact of large-scale environmental deterioration was considered by Coyne et al. (1997) in their comprehensive critique of the shifting-balance process. However, they dismissed the importance of such environmental changes by saying, "If large regions go extinct the variation in success of different regions is likely to be determined largely by extrinsic factors such as climate, leaving even less scope for the differential spread of particular gene combinations that increase fitness." The current study has shown that, contrary to this intuition, widespread environmental deterioration may magnify rather than mask the importance of between-deme differences in viability.

Note that the work presented here relates to environmental changes over large regions, as suggested by Coyne et al. (1997). The present study follows up on work carried out by Lande (1979, 1985), which focused on the extinction of local

TABLE 2. Results for various values of T (the number of generations spent at the environmental nadir) and m (the migration rate). See Table 1 for additional details.

T m	10^{-3}	10^{-4}	10^{-5}	10^{-6}
4000	(0.459) 0.543 0.998 1.000	(0.423) 0.408 0.989 1.000	(0.448) 0.372 0.985 1.000	(0.460) 0.343 0.975 1.000
2000	(0.434) 0.501 0.998 1.000	(0.455) 0.365 0.987 1.000	(0.460) 0.295 0.970 1.000	(0.452) 0.262 0.962 1.000
1000	(0.457) 0.366 0.991 1.000	(0.470) 0.245 0.951 1.000	(0.444) 0.213 0.923 1.000	(0.470) 0.167 0.857 0.999
500	(0.438) 0.229 0.953 1.000	(0.469) 0.150 0.824 0.999	(0.441) 0.123 0.703 0.988	(0.468) 0.098 0.594 0.970
250	(0.479) 0.131 0.739 0.995	(0.450) 0.087 0.529 0.938	(0.446) 0.073 0.434 0.875	(0.449) 0.064 0.364 0.793
125	(0.500) 0.075 0.435 0.872	(0.504) 0.056 0.304 0.698	(0.528) 0.049 0.246 0.596	(0.509) 0.043 0.218 0.521

populations. Lande's work, in turn follows suggestions initially made by Wright (1931, 1940, 1941, 1970).

Our results suggest that, over the course of several cycles of large-scale environmental decline and recovery, a superior coadapted genotype can often go from being common in a single deme to being common throughout an entire metapopulation. Furthermore, this can take place even when local between-deme variation in environmental quality has a much larger effect than the differences between alternative genotypes (see Table 1). Presumably, as a superior genotype begins to spread, between-deme differences in environmental quality tend to become averaged out among demes, so that the differences between genotypes become increasingly important.

The results presented here may lessen the force of some of the criticisms levelled at Wright's shifting-balance process. For example, critics have suggested that, during phase III, barriers to migration will prevent the spread of the high-fitness type—even if these barriers are leaky (Barton and Hewitt 1989; Gavrilets 1996; Coyne et al. 1997). (Leaky barriers might include a small stream for terrestrial species, or a narrow area of grassland for woodland species.) The present study suggests that these sorts of stalemates may often be overcome at times when the species range is shifting or contracting. Our results also address the supposed ineffectiveness of phase III when migration rates are very low, a state that may be common in nature (Ehrlich and Raven 1969).

However, the present study certainly does not constitute a thoroughgoing defense of the shifting-balance process; such a defense would have to deal with the origin and local initial increase of a high-fitness genotype (phases I and II). This notwithstanding, the present results may have a wider relevance. This is because superior coadapted genotypes may arise by means other than those specified by Wright. For example, as Dobzhansky pointed out, there may be ridges in the adaptive landscape, and local populations may move along these ridges as they traverse from the region of one adaptive peak to another (Dobzhansky 1937; Gavrilets 1997). Transitions between peaks might also happen as a result of local environmental fluctuations (Whitlock 1997). Regardless of how they arise, new coadapted high-fitness genotypes should be able to spread through a metapopulation by means of a process similar to the one described here. Thus, changes in species range may lead to significant adaptive changes in the genetic constitution of a species, regardless of whether genetic drift is the means by which coadapted genotypes arise.

Acknowledgments

This work was supported by a grant from the Biology and Biotechnology Research Council (U.K.) and the University of Sussex under its Graduate Teaching Assistantship scheme. We are grateful to N. H. Barton, R. Lande, D. Waxman, L. Nunney, and anonymous reviewers for helpful comments and advice.

Literature Cited

Barton, N. H. 1992. On the spread of new gene combinations in the third phase of Wright's shifting-balance. Evolution 46: 551–557.

Barton, N. H., and G. M. Hewitt. 1989. Adaptation, speciation and hybrid zones. Nature 341:497–503.

Barton, N. H., and S. Rouhani. 1987. The frequency of shifts between alternative equilibria. J. Theor. Biol. 125:397–418.

———. 1993. Adaptation and the 'shifting balance'. Genet. Res. 61:57–74.

Bergman, A., D. B. Goldstein, K. E. Holsinger, and M. W. Feldman. 1995. Population structure, fitness surfaces, and linkage in the shifting balance process. Genet. Res. 66:85–92.

Coyne, J. A., N. H. Barton, and M. Turelli. 1997. Perspective: A critique of Sewall Wright's shifting balance theory of evolution. Evolution 51:643–671.

———. 2000. Is Wright's shifting balance process important in evolution? Evolution 54:306–317.

Crow, J. F., W. R. Engels, and C. Denniston. 1990. Phase three of Wright's shifting-balance theory. Evolution 44:233–247.

Dobzhansky, T. 1937. Genetics and the origin of species. Columbia Univ. Press, New York.

Ehrlich, P. R., and P. H. Raven. 1969. Differentiation of populations. Science 165:1228–1232.

Gavrilets, S. 1996. On phase three of the shifting-balance theory. Evolution 50:1034–1041.

———. 1997. Evolution and speciation on holey adaptive landscapes. Trends Ecol. Evol. 12:307–312.

Gomulkiewicz, R., and R. D. Holt. 1995. When does evolution by natural selection prevent extinction? Evolution 49:201–207.

Gomulkiewicz, R., R. D. Holt, and M. Barfield. 1999. The effects of density dependence and immigration on local adaptation and niche evolution in a black-hole sink evolution. Theor. Popul. Biol. 55:283–296.

Haldane, J. B. S. 1959. Natural selection. Pp. 101–149 *in* P. R. Bell, ed. Darwin's biological work: some aspects reconsidered. Wiley, New York.

Holt, R. D. 1996. Adaptive evolution in source-sink environments: direct and indirect effects of density-dependence on niche evolution. Oikos 75:182–192.

Kirkpatrick, M., and N. H. Barton. 1997. Evolution of a species' range. Am. Nat. 150:1–23.

Kondrashov, A. S. 1992. The 3rd phase of Wright's shifting-balance: a simple analysis of the extreme case. Evolution 46: 1972–1975.

Lande, R. 1979. Effective deme sizes during long-term evolution estimated from rates of chromosomal rearrangement. Evolution 33:234–251.

———. 1985. The fixation of chromosomal rearrangements in a subdivided population with local extinction and colonization. Heredity 54:323–332.

Mitteldorf, J., and D. S. Wilson. 2000. Population viscosity and the evolution of altruism. J. Theor. Biol. 204:481–496.

Moore, F. B. G., and S. J. Tonsor. 1994. A simulation of Wright's shifting balance process: migration and the three phases. Evolution 48:69–80.

Phillips, P. C. 1993. Peak shifts and polymorphism during phase-3 of Wright's shifting-balance process. Evolution 47:1733–1743.

Prout, T. 1980. Some relationships between density independent selection and density dependent population growth. Evol. Biol. 13:1–68.

Ronce, O., and M. Kirkpatrick. 2001. When sources become sinks: migrational meltdown in heterogeneous habitats. Evolution 55: 1520–1531.

Rouhani, S., and N. H. Barton. 1987. Speciation and the "shifting balance" in a continuous population. Theor. Popul. Biol. 31: 465–492.

———. 1993. Group selection and the 'shifting balance'. Genet. Res. 61:127–135.

Shmida, A., and S. Ellner. 1984. Coexistence of plant species with similar niches. Vegetatio 58:29–55.

Shpak, M., and A. S. Kondrashov. 1999. Applicability of the hypergeometric phenotypic model to haploid and diploid populations. Evolution 53(2):600–604.

Whitlock, M. C. 1997. Founder effects and peak shifts without genetic drift: Adaptive peak shifts occur easily when environments fluctuate slightly. Evolution 51:1044–1048.

Wright, S. 1931. Evolution in Mendelian populations. Genetics 16: 97–159.
———. 1932. The roles of mutation, inbreeding, crossbreeding and selection in evolution. Proc. Sixth Int. Congress Genet. 1: 356–366.
———. 1940. Breeding structure of populations in relation to speciation. Am. Nat. 74:232–248.
———. 1941. On the probability of fixation of a reciprocal translocation. Am. Nat. 75:513–522.
———. 1949. Population structure in evolution. Proc. Am. Philos. Soc. 93:471–478.
———. 1970. Random drift and the shifting balance theory of evolution. Pp. 1–31 *in* K. Kojima, ed. Mathematical topics in population genetics. Springer-Verlag, Berlin.
———. 1977. Evolution and the genetics of populations. Vol. 3. Univ. of Chicago Press, Chicago, IL.
———. 1982. The shifting balance theory and macroevolution. Annu. Rev. Genet. 16:1–19.

Corresponding Editor: L. Nunney

Appendix 1

This appendix sets out some of the details of the simulation procedures and provides the rationale behind the choice of the standard parameter values. The following also applies to the numerical study shown in Figure 1, except as detailed in Appendix 2.

For the intrinsic differences in the environmental quality of demes, Q_i (eq. 6) we take $Q_i = \exp(-x_i)$, where x_i is a random number drawn from a uniform distribution on the range $(0, 3\ln[v])$. This means that the logarithm of the ratio of two randomly chosen Θ_i values—the quantity whose expectation is taken in eq. (7)—has a symmetrical triangular distribution. The range of this distribution is $v^{-3} \leq Q_i \leq 1$, and so varying v alters its lower bound, but leaves other important properties unaltered.

For the cycling component of environmental quality, $C(t)$, (eq. 6), we assume that the function is on the range $(C_{\min}, 1)$ where $1 > C_{\min} > 0$. The functional form is as described in the text, with a sinusoidal decline. We allow $C_{\min}$ to vary to fix the proportion of wild-type demes that are demographic sources, in the absence of migration, at the environmental nadir; that is the expected proportion of demes for which $C_{\min}Q_iV_2F_0 > 1$. Given the method of generating the Q_i explained above, this proportion, denoted $P_{\min}$, is given by $P_{\min} = \ln(C_{\min}F_0V_2)/[3\ln(v)]$. (As expected, this is only defined for $v > 1$.) In the simulation trials, we fix $P_{\min} = 0.1$ and so, on average, 10% of demes populated by the wild type remain as demographic sources at the environmental nadir.

For the birth function, we take $F(N) = R_0\,[1 - \exp(-\alpha N)]/N$. This function yields conventional sigmoidal growth, but unlike the more traditional logistic function, it does not exhibit chaotic dynamics and is derived from an explicit model of resource consumption.

To derive this function, assume that each deme member consumes some resource at a rate αr, where α, $r > 0$ and r is the density of the resource. This means that the rate of consumption declines as the resource becomes more scarce. Say, for instance, that resources are replenished at the beginning of each generation, and that, each generation, the initial value of r is R_0. If we scale time so that resource consumption lasts one unit of time per generation, then each individual will consume a total amount of resources given by $R_0[1 - \exp(-\alpha N)]/N$. We obtain the birth function by assuming that the measurement of resources is scaled so that the expected number of offspring produced by female effort increases by one for each unit of resources consumed. With this function, we have $F_0 = R_0\alpha$ (see eq. 1). In all simulations detailed in Tables 1 and 2, we take $R_0 = \alpha = 2$, and so $F_0 = 4$. For the viability of the wild type, we choose $V_2 = 0.7$.

With these parameter values, for the range of v values used in the simulation trials, every deme populated by the wild type is guaranteed to be a demographic source when environmental conditions are optimal (i.e., when $C[t] = 1$). Note that the assumptions of Rouhani and Barton (1993) regarding population regulation are equivalent to letting $\alpha \to \infty$ (see Discussion).

In all simulations, we take $L = 5$ loci. This takes us into the regime where fitness differences have negligible effects on the outcome of phase III in a model without explicit demographics (Barton 1992). Although 528 distinct genotypes are possible with five loci, we can drastically reduce the number of necessary variables by making use of the hypergeometric model introduced in the appendix of Barton (1992; see also Shpak and Kondrashov 1999).

Appendix 2

The numerical study shown in Figure 1 followed the simulation procedures described in the main text and Appendix 1, with the exception of the way the environmental quality values, $\Theta_i(t)$, were calculated. The trial consisted of two consecutive sets of 1000 generations, and for each set there was a refuge. The center of the refuge for the first set of 1000 generations appears in Figure 1 as the fifth deme from the bottom of the metapopulation in the fifth column from the left side of the metapopulation. This deme was used to calculate the distance from the refuge for all other demes. The point at the center of each deme was used to calculate the distance between demes. Distance was calculated on a scale so that the left-right distance from the center of the demes furthest to the left to the center of demes furthest to the right is equal to 1.0. Using this scale, the distance of each deme from the center of the refuge was calculated. For deme i this distance is denoted as D_i. For a particular generation, t, the value of Θ in deme i was then calculated as $\Theta_i(t) = \exp[-C(t)D_i]$, where $C(t)$ is a sinusoidal function that cycles between zero, at generations $t = 1$ and $t = 1001$ and unity at generation $t = 500$. This same procedure was repeated with the second set of 1000 generations. The only difference was that the center of the refuge was moved so that, in Figure 1, it appears as the fifth deme from the top of the metapopulation in the fifth column from the right side of the metapopulation. The horizontal line in each deme in Figure 1 is proportional to $5 + \ln(N_i)$. The parameters for the birth function were set at $R_0 = 2$, and $\alpha = 1$.