J. theor. Biol. (1999) **196,** 197–209
Article No. jtbi.1998.0833, available online at http://www.idealibrary.com on IDEAL

A Quantitative Model of the Simpson–Baldwin Effect

LAUREN W. ANCEL*

Department of Biological Sciences, Stanford University, Stanford, CA 94305, *U.S.A.*

(*Received on* 26 *February* 1998, *Accepted in revised form on* 10 *September* 1998)

G. G. Simpson was the first to explain the Baldwin Effect completely in terms of the theory of natural selection. A genetic version of a seemingly non-hereditary adaptation may arise when natural selection acts on the likelihood of having an adaptive trait not just on the trait itself. We present a quantitative model of the Simpson–Baldwin Effect. Organisms in the model have mutable ranges of phenotypic plasticity. The distribution of phenotypes in a population depends largely on the extent of environmental stochasticity. When the environment undergoes intermediate rates of fluctuation, the Simpson–Baldwin effect arises through the interaction of natural selection and mutation on norms of reaction. In a highly volatile environment, organisms benefit from plasticity, and consequently do not experience a Simpson–Baldwin channeling of phenotypic possibility.

© 1999 Academic Press

Introduction

Just over a century ago J. Mark Baldwin postulated a conciliatory alternative to both Neo-Lamarckism and Neo-Darwinism. His influential treatise, *A New Factor in Evolution*, introduced organic selection as the evolution of hereditary adaptations by way of non-hereditary innovations (Baldwin, 1896). Organisms that make non-hereditary physical or behavioral modifications to survive environmental stresses will have better representation in future generations than less versatile organisms. Through natural selection then, the capacity for such adaptation along with the beneficial acquired traits will become universal. Although it initially received much attention (Morgan, 1896; Osborn, 1896), this line of theory faded with the rediscovery of Mendel and the subsequent development of the modern evolutionary synthesis.

In 1953, G. G. Simpson reintroduced the *Baldwin Effect* as genetical reinforcement of advantageous but initially non-hereditary traits. Until Simpson, the Baldwin Effect was seen as an evolutionary process distinct from natural selection. Simpson's critical insight was that the Baldwin Effect can be explained by the theory of natural selection. A genetic version of a seemingly non-hereditary adaptation may arise when natural selection acts on the likelihood of having an adaptive trait, not just on the trait itself.

Inspired by Schmalhausen (1949), Simpson outlined the following three conditions under which natural selection may produce a Baldwin Effect.

[S1] The ability to acquire a character has in itself a genetical basis.

*E-mail: ancel@charles.stanford.edu

0022–5193/99/002197 + 13 $30.00/0

© 1999 Academic Press

[S2] . . . selection for the ability to acquire an adaptive character so narrowed the developmental range that the character would usually or invariably appear.
[S3] There is . . . a certain balance between lability and stability of developmental ranges and norms in evolution (Simpson, 1953).

This rendering generalizes Baldwin's original focus on traits learned within an individual life. A genetically determined norm of reaction is, broadly speaking, a range of phenotypic possibility. The actual trait or traits that an individual possesses may be acquired via learning throughout life, or they may be determined through physical interaction with the environment during early development.

Simpson's ideas brought the Baldwin Effect back to the forefront of evolutionary thought. In particular, it inspired the Hinton and Nolan legacy of models in which learning guides the genetical evolutionary trajectory (Hinton & Nolan, 1980). For various modes of environmental interaction and response, learning is tantamount to a search of phenotype space, and the population evolves towards those phenotypes that find higher fitness faster. Learning initially finds optimal phenotypes, and genetical selection follows with reinforcement.

In this paper, we present a quantitative model of Simpson's simple formulation of the Baldwin Effect. Individuals have genetically determined degrees of phenotypic plasticity which undergo mutation and selection. The tension between intrinsic costs of phenotypic plasticity and the benefits of flexibility in a dynamic environment steers evolution. Sometimes the balance of these forces narrows individual norms of reaction. This is what we call the Simpson–Baldwin effect. On the other hand, in volatile environments we observe the evolution of greater plasticity from ancestral inflexibility. Through the model, we explore the relationship between environmental fluctuation, fitness regimes, mutation rates and the occurrence of the Simpson–Baldwin effect.

The Model

The following is a simple model for the evolution of norms of reaction under a fluctuating environment. The parameters and variables are outlined in Table 1. We assume a constant finite population size N of asexual individuals. Each individual is defined by a range of phenotypic plasticity. The space of all possible phenotypes is $[0, 1]$. An individual i has a genetically determined norm of reaction $\rho_i = [l_i, u_i]$, which specifies its interval of phenotypic possibilities.

Mutation occurs at a rate μ as modifications to the upper and lower bounds of the ρ_is. Each mutation in individual i is either an increment or decrement of size v to either l_i or u_i. Mutation acts on upper and lower bounds with equal frequency. The direction of each mutation is

TABLE 1
Model parameters and variables

	Parameters & variables	Description	Units
N	$N \in \mathbf{Z}^+$	Total population size	Individuals
P	$P = [0, 1]$	Range of all possible phenotypes	Phenotype units
ρ_i	$\rho_i = [l_i, u_i] \subseteq [0, 1]$	Range of phenotype space accessible to individual i, $1 \leqslant i \leqslant N$	Phenotype units
v	$0 \leqslant v \leqslant 1$	Size of increment or decrement to upper or lower bound of ρ_i upon mutation	Phenotype units
μ	$0 \leqslant \mu \leqslant 1$	Mutation rate	Mutations/Individuals·unit time
σ	$0 \leqslant \sigma \leqslant 1$	Rate at which mutations decrease the norm of reaction	Individuals/Mutation
ϵ	$0 \leqslant \epsilon$	Rate of environmental fluctuation	Changes/Unit time
Ω_t	$0 \leqslant \Omega_t \leqslant 1$	Optimal phenotype at generation t	Phenotype units
κ	$0 \leqslant \kappa$	Fitness benefit to an individual when Ω_t lies within its norm of reaction	

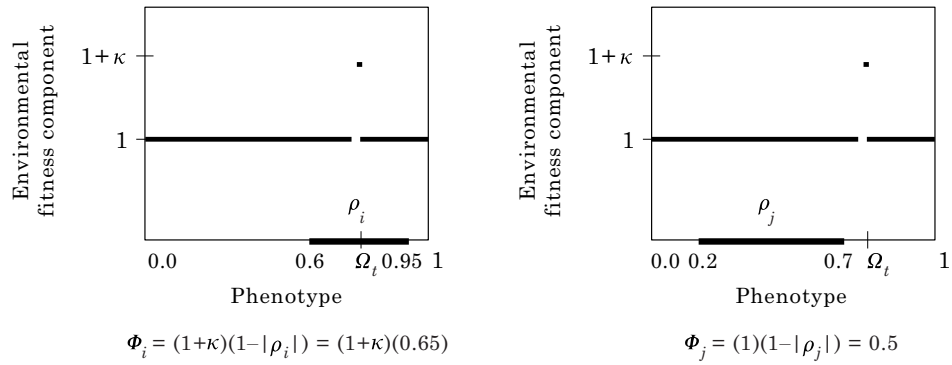

FIG. 1. Sample fitnesses: Ω_t is the optimal phenotype at time t. Individuals i and j have norms of reaction $\rho_i = [0.6, 0.95]$ and $\rho_j = [0.2, 0.7]$, respectively.

governed by σ and $1 - \sigma$, the proportion of mutations that cause a norm of reaction to shrink and expand, respectively.

The fitness of an individual is determined by both the state of the environment and the length of its norm of reaction. An environmental fitness component at time t is a function from phenotype space $P = [0, 1]$ to $\mathbf{R}^+ \cup \{\mathbf{0}\}$. All phenotypes except the optimal phenotype Ω_t are mapped to 1, while Ω_t is mapped to $1 + \kappa$. Then the fitness Φ_i of individual i with norm of reaction ρ_i is determined according to:

$$\Phi_i = \begin{cases} (1 + \kappa)(1 - |\rho_i|) & \text{if } \Omega_t \in \rho_i, \\ 1 - |\rho_i| & \text{otherwise,} \end{cases}$$

where $|\rho_i|$ denotes the length of the interval ρ_i. According to this fitness regime, all individuals with phenotypic ranges containing the optimal phenotype benefit. Figure 1 illustrates the fitnesses for two different Ω_t–ρ_i combinations. There is a trade-off between the intrinsic costs of phenotypic plasticity and improving the probability of reaching the optimal phenotype. The former selects for shortened norms of reaction while the latter improves with enlarged ranges.

This fitness function is motivated by the costs of phenotypic flexibility apparent in many biological systems. In bacteria, for example, plasticity requires genetic machinery whose costs are in terms of increased replication time. Learning-based plasticity may entail energetic costs of searching. RNA molecules with multiple configurations may trade accuracy for the potential for variable binding. This simple map from phenotype to fitness also has the benefits of analytical tractability and of simplifying the conditions under which the Baldwin Effect may arise. Continuous and multiple peaked fitness functions give rise to interesting phenomena worthy of future analysis.

In each generation, we perform the following algorithm.

- Determine the environment: $\text{Prob}(\Omega_t = \Omega_{t-1}) = 1 - \epsilon$, where $0 \leqslant \epsilon \leqslant 1$ is the number of environmental transitions per generation. If $\Omega_t \neq \Omega_{t-1}$, then choose Ω_t randomly from a uniform distribution [0, 1].
- Calculate relative fitnesses as described above.
- Asexually produce offspring in proportion to fitness, keeping total population size constant.
- Mutate offspring at the following rates.

Old p_i	New ρ_i	Effect on ρ_i	Rate	
$[l_i, u_i]$	$[l_i + v, u_i]$	Shrink	$\frac{1}{2}\mu\sigma$	
$[l_i, u_i]$	$[l_i - v, u_i]$	Extend	$\frac{1}{2}\mu(1 - \sigma)$	if $l_i \geqslant v$
			0	if $l_i = 0$
$[l_i, u_i]$	$[l_i, u_i + v]$	Extend	$\frac{1}{2}\mu(1 - \sigma)$	if $u_i \leqslant 1 - v$
			0	if $u_i = 1$
$[l_i, u_i]$	$[l_i, u_i - v]$	Shrink	$\frac{1}{2}\mu\sigma$	

In the following three-part analysis, we seek parametric conditions that result in the Simpson–Baldwin effect. First we describe the population dynamics and equilibrium conditions in a static environment. Then we find that the rate of environmental fluctuation dramatically affects the possibility for narrowing norms of

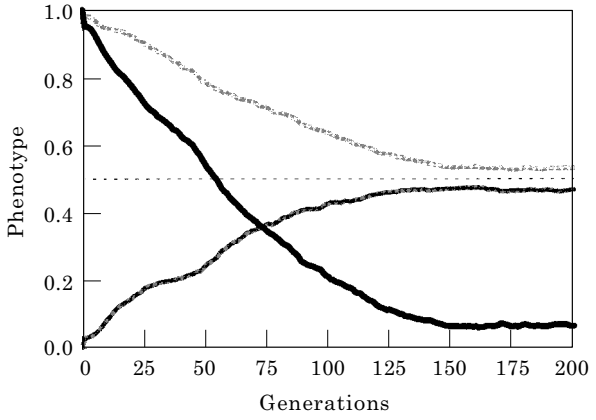

FIG. 2. Constant environment equilibrium: beginning with $\rho = [0, 1]$ for all i, and assuming $\Omega_t = 0.5$, $\mu = 0.1$, $\sigma = 0.25$ and $v = 0.05$, this illustrates the Simpson–Baldwin narrowing of norms of reaction in a static environment. The solid lines depict the population averages for norm of reaction lower bounds $(\Sigma_{i=1}^{1000} l_i)/N$, upper bounds $(\Sigma_{i=1}^{1000} u_i)/N$ and range lengths $(\Sigma_{i=1}^{1000} |\rho_i|)/N$, and the broken line shows Ω_t. (Dark curve), average range; (medium curve), average lower bound; (light curve), average upper bound; (- - -), optimal phenotype.

reaction: rapid environmental shifts select for individuals with large norms of reaction capable of quickly taking advantage of diverse phenotypic optima, while infrequent changes enable populations to search phenotype space through mutation, and then hone in on the optimal phenotype via selection for small norms of reaction. By calculating the time between subsequent equilibria, we determine the environmental fluctuation–mutation combinations in which the Simpson–Baldwin effect is realized.

will occur. After some time, the average size of the norms of reaction will converge to an equilibrium. Figure 2 illustrates this trend. The equilibrium average norm of reaction length results from a balance between mutational pressure extending the phenotypic ranges and selection pressure towards shorter ranges.

Given values for μ and σ, we can determine analytically the expected size of the average phenotypic range at equilibrium. The following analysis assumes that every individual has a norm of reaction containing the optimal phenotype. Hence the fitness benefit κ is felt by all, and is therefore ignored in calculations of relative fitness in the following analysis. Since each mutation adjusts the size of a phenotypic range by exactly v, the range of any individual necessarily has length $L \cdot v$ for some integer L where $0 \leqslant L \leqslant 1/v$. We construct a matrix $\mathbf{M}_{ij}$, $0 \leqslant i, j \leqslant 1/v$ in which the ij-th entry is the product of the probability with which a norm of reaction of length iv will mutate in a single time step to have length jv, and the relative fitness of norms of reaction length iv. For example, if $v = 0.05$, then there are 21 possible range lengths; that is, for all i, $|\rho_i| \in \{0, 0.05, 0.1, \ldots, 1\}$. Then $\mathbf{M}_{2,3}$ would be the relative fitness of a norm of reaction length 0.10 times the probability that an individual i with $|\rho_i| = 0.10$ will mutate to $|\rho_i| = 0.15$ in a single time step. This probability is $\mu(1 - \sigma)(1 - 0.10)$.

In terms of the model parameters, the matrix is the following.

$$\mathbf{M} = \begin{bmatrix} (1-\mu+\mu\sigma) & \mu(1-\sigma) & 0 & 0 & \cdots & 0 & 0 \\ \mu\sigma(1-v) & (1-\mu)(1-v) & \mu(1-\sigma)(1-v) & 0 & \cdots & 0 & 0 \\ 0 & \mu\sigma(1-2v) & (1-\mu)(1-2v) & \mu(1-\sigma)(1-2v) & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & \mu(1-\sigma)(2v) & 0 \\ 0 & 0 & 0 & 0 & \cdots & (1-\mu)(v) & \mu(1-\sigma)(v) \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix}.$$

Equilibrium in a Constant Environment

To assess equilibrium conditions in a constant environment, we assume $\epsilon = 0$. Suppose our population has 1000 individuals, each of which begins with a norm of reaction containing the entire phenotype space, i.e. for all i, $\rho_i = [0, 1]$. Initially then only range-shrinking mutations

Although $\mathbf{M}$ reflects the inter-generation mutation–selection dynamics, it is not exactly a transition matrix. In our model, we maintain a constant population size. We can modify the matrix to accurately reflect population transitions by multiplying M by the constant $1/\lambda$ where λ is the leading eigenvalue of $\mathbf{M}$. The

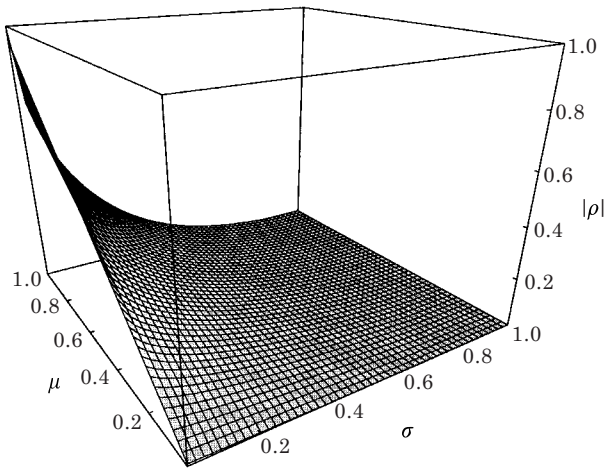

FIG. 3. Equilibrium average range: average length of norm of reaction across the population as a function of the mutation rate μ and the fraction of shrinking mutations σ ($v = 0.05$).

matrix $\mathbf{M}_\lambda = 1/\lambda \cdot \mathbf{M}$ has leading eigenvalue equal to 1, and gives the transition probabilities for a population of constant size. Although the leading eigenvalue of $\mathbf{M}_\lambda$ is different than that of $\mathbf{M}$, the associated leading left eigenvectors are identical. We therefore consider $\mathbf{M}$ in lieu of $\mathbf{M}_\lambda$ in the following discussion.

The leading left eigenvector of $\mathbf{M}$ describes the equilibrium distribution of phenotypic range lengths. From this distribution we calculate the average range at equilibrium. Since such calculations are computationally intensive, we are limited to solutions for specific parameter values. Figures 3 and 4 illustrate the average range at equilibrium as we vary μ and σ for $v = 0.05$.

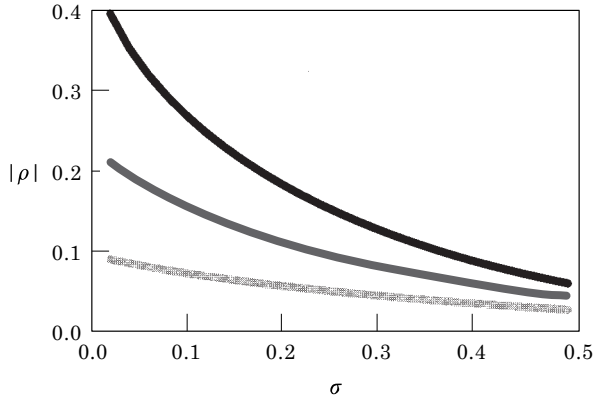

FIG. 4. Equilibrium average range for specific μ: average length of norm of reaction across the population as a function of the fraction of shrinking mutations σ for $\mu = 0.1$ (light curve), $\mu = 0.25$ (medium curve) and $\mu = 0.5$ (dark curve).

As μ increases and as σ decreases, the size of the average phenotypic range at equilibrium increases to 1. A combination of large μ and small σ results in a high rate of mutation towards enlarged norms of reaction. Hence the faster phenotypic ranges expand due to mutation, the larger the equilibrium size of these ranges. A low rate of range-expanding mutation tips the mutation–selection balance towards selection, resulting in equilibrium ranges as narrow as the single point Ω_t. In addition, we find that the equilibrium average range length obtained through simulation (see the black curve of Fig. 2) matches well the prediction of the analysis shown by the bottom curve of Fig. 4.

Population Dynamics in a Changing Environment

In a static environment, we observe shrinking of phenotypic flexibility. Population dynamics become more interesting when we add environmental variation. Figure 5(a) depicts the evolution of a population size $N = 200$ in an environment which changes on average once every 200 generations, i.e. $\epsilon = 0.005$. Assuming further that $\mu = 0.1$, $\sigma = 0.25$ and $\kappa = 0.2$, we track the dynamics in terms of population averages for lower bound ($\Sigma_{i=1}^{N} l_i/N$), upper bound ($\Sigma_{i=1}^{N} u_i/N$) and norm of reaction length ($\Sigma_{i=1}^{N} |\rho_i|/N$). The dashed line represents the value of Ω_t. Note that steps in this line mark environmental transitions.

At the onset of a new environment (mark *a*), the population is typically at equilibrium with average norm of reaction centered around the prior optimal phenotype. In most cases, a subsequent search of phenotype space is marked by jumps in the average phenotypic range size (mark *b*). Upon discovery of a new optimal phenotype, selection drives the concerted movement of the average lower and upper bounds until they surround the new optimum (mark *c*). When the distance between subsequent phenotypic optima is very large $|\Omega_{new} - \Omega_{old}| \gg 0$ (mark *d*), then the population may not have enough time to happen upon Ω_{new} before the next environmental transition.

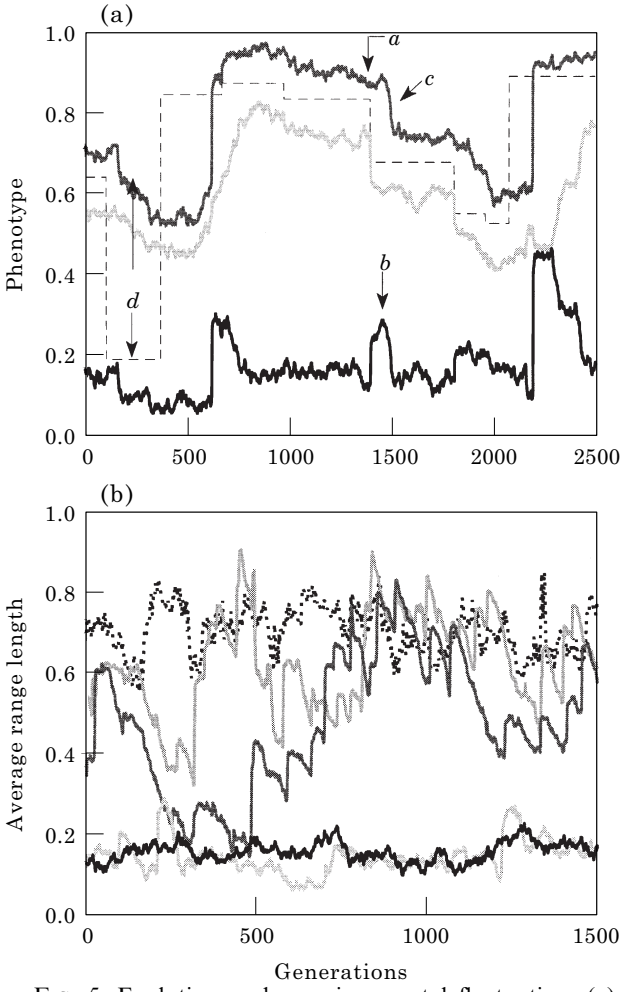

FIG. 5. Evolution under environmental fluctuation: (a) displays the average lower bound (light curve), average upper bound (medium curve), average range length (dark curve) and optimal phenotype (- - -) for a single population where $\epsilon = 0.005$; (b) shows average norm of reaction lengths under varying rates of environmental fluctuation ϵ. (▬), 0.001; (▬), 0.01; (▬), 0.05; (▬) 0.1; (▪ ▪), 1.0. In both figures, $N = 200$, $v = 0.05$, $\mu = 0.1$, $\sigma = 0.25$ and $\kappa = 2$.

Figure 5(b) compares the evolution of average norm of reaction lengths under different rates of environmental fluctuation. As the rate of environmental change ϵ increases from 0.001 to 1, the time that the population spends at equilibrium decreases dramatically. Populations with small ϵ spend most of their time in equilibrium, shifting away briefly following the infrequent environmental shifts. At the other extreme, $\epsilon = 1$, the rapidity of change prevents the population from ever coming close to equilibrium.

In order to quantify the effect of environmental variation on population evolution, we first estimate the rate at which a population searches phenotype space before the first individual encounters the new optimal state, and then estimate the rate at which a population converges on that optimal phenotype.

FROM EQUILIBRIUM TO FIRST ENCOUNTER

We calculate the time between equilibria as a function of the distance between subsequent environments. By distance we mean the difference between phenotypic optima, $|\Omega_{new} - \Omega_{old}|$. Suppose a population is at equilibrium in the initial environment, characterized by Ω_{old}. If the environment changes to Ω_{new}, and no individual norm of reaction contains Ω_{new}, then the population will search the phenotype space via a random walk until an individual first encounters Ω_{new}. Without loss of generality, suppose $\Omega_{new} > \Omega_{old}$. Then the first encounter with Ω_{new} will occcur when an individual experiences an upper bound mutation that extends the norm of reaction to include Ω_{new}.

The evolving distribution of the upper bounds (the u_is) breaks down into two components. First the population explores phenotype space by mutation since individuals no longer profit from phenotype Ω_{old}. To chart this wandering, we consider the distribution of the norm of reaction centers, i.e. for all i, $c_i = (l_i + u_i)/2$. Second we estimate the distribution of the norm of reaction lengths during this period of exploration. As long as the entire population experiences the same environmental component of fitness—that is each individual has a norm of reaction containing Ω_t or no individual experiences Ω_t—then the leading left eigenvector of $\mathbf{M}$ describes the equilibrium range lengths. See for example mark d of Fig. 5(a). Recall that $\mathbf{M}$ includes only range-size fitness effects and ignores environmental fitness variations. The leading left eigenvector is determined by the mutation–selection balance, and applies both at the first equilibrium when all ρ_i contain Ω_{old}, and immediately following the environmental switch as long as no ρ_i contains Ω_{new}. Together the evolving distribution of the c_is and the static

distribution of phenotype range lengths yield the desired estimates for the changing distribution of the upper bounds.

We track the population in phenotype space through the distribution of norms of reaction midpoints. At equilibrium, the distribution of the c_is is shown through simulation to be approximately normal as depicted in Fig. 6.

After equilibrium, the c_is move approximately according to N independent identical unbiased random walks. For any i, $c_{i_t} = c_{i_{t-1}} + b$ where b is a random variable such that

$$b = \begin{cases} -\dfrac{v}{2} & \text{with probability } \dfrac{1}{2}(\mu\alpha_\mu) \\ \dfrac{v}{2} & \text{with probability } \dfrac{1}{2}(\mu\alpha_\mu) \\ 0 & \text{with probability } 1-\mu\alpha_\mu \end{cases}$$

where $\alpha_\mu = [v\mu(\mu-1)(1-2\sigma]\,/\,[(1-|\bar{\rho}\,|) - v\mu(1-2\sigma)]$ and $|\bar{\rho}|$ is the average range size at equilibrium as given by the leading eigenvector of **M**. See the Appendix for the derivation of these transition probabilities.

The changing distribution of the norm of reaction midpoints can be computed using a matrix **C** of transition probabilities,

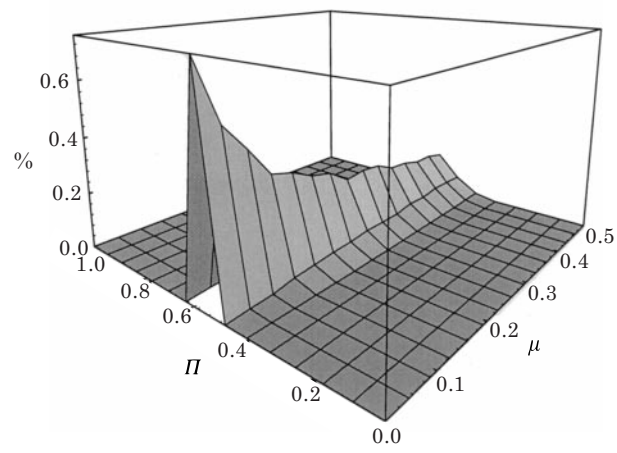

FIG. 6. Frequency distribution of norm of reaction midpoints for various μ: we assume $N = 1000$, $v = 0.05$ and $\sigma = 0.25$ and graph the distribution of norm of reaction centers, $(l_i + u_i)/2$, achieved through simulation to equilibrium. In this graph, Π denotes phenotype.

$$\mathbf{C} = \begin{bmatrix} 1-\mu\alpha_\mu+\frac{1}{2}\mu\alpha_\mu & \frac{1}{2}\mu\alpha_\mu & 0 & 0 & \dots & 0 & 0 \\ \frac{1}{2}\mu\alpha_\mu & 1-\mu\alpha_\mu & \frac{1}{2}\mu\alpha_\mu & 0 & \dots & 0 & 0 \\ 0 & \frac{1}{2}\mu\alpha_\mu & 1-\mu\alpha_\mu & \frac{1}{2}\mu\alpha_\mu & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & \frac{1}{2}\mu\alpha_\mu & 0 \\ 0 & 0 & 0 & 0 & \dots & 1-\mu\alpha_\mu & \frac{1}{2}\mu\alpha_\mu \\ 0 & 0 & 0 & 0 & \dots & \frac{1}{2}\mu\alpha_\mu & 1-\mu\alpha_\mu+\frac{1}{2}\mu\alpha_\mu \end{bmatrix}.$$

For $0 \leqslant i, j \leqslant 2/v$, $\mathbf{C}_{ij}$ gives the probability that an individual with midpoint $v/2 \cdot i$ will give rise to an individual in the next generation with midpoint $v/2 \cdot j$. For example, if $v = 0.05$, then norm of reaction midpoints will lie in the set $\{0, 0.025, 0.05, 0.075, 0.1, \dots, 1\}$ and $\mathbf{C}_{11,12}$ gives the probability that a norm of reaction centered around 0.275 mutates in one generation to be centered around 0.3. This occurs either through a mutational increment of $v = 0.05$ to the upper bound or decrement of $v = 0.05$ to the lower bound. To approximate the distribution of norm of reaction centers at t generations after the environmental switch, we apply the matrix $\mathbf{C}^t$ to the equilibrium distribution of midpoints.

Now we calculate the distribution of the $u_i - c_i$, the distance between the upper bounds and centers of the norms of reaction. Since $u_i - c_i = u_i - (l_i + u_i)/2 = (u_i - l_i)/2 = |\rho_i|/2$, then this is just the distribution of the $|\rho_i|/2$, which does not change after the environmental shift until the population encounters Ω_{new}. Recall that the distribution of the $|\rho_i|$, hence the $|\rho_i|/2$, is described by the leading left eigenvector of transition matrix **M**.

We generate the distribution of the u_i at time t by combining the two approximations with the equilibrium distribution of c_is acquired through simulation. Let P_t denote a probability at

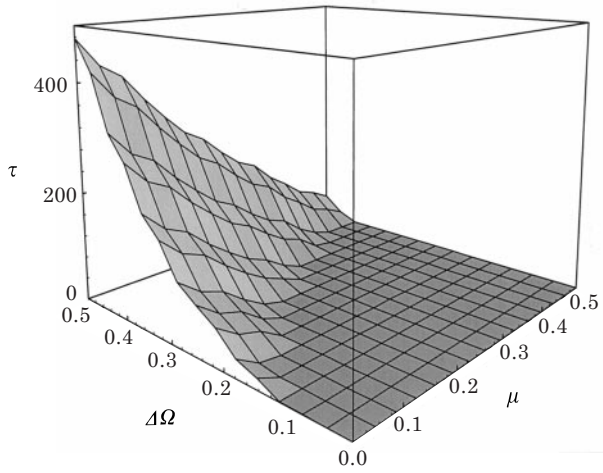

FIG. 7. Time until population first encounters new optimum: we assume $N = 1000$, $v = 0.05$ and $\sigma = 0.25$ and graph the time τ until the first individual encounters Ω_{new} in terms of the mutation rate μ and the distance between environmental optima $\Delta\Omega = |\Omega_{new} - \Omega_{old}|$.

time t. Then,

$$P_t\left[u_i = U \cdot \frac{v}{2}\right]$$

$$= \sum_{j=1}^{U} P_t\left[c_i = j \cdot \frac{v}{2}\right] \cdot P_t\left[\frac{|\rho_i|}{2} = (U-j) \cdot \frac{v}{2}\right]$$

$$= \sum_{j=1}^{U} \mathbf{C}^t\mathbf{c}_0[j] \cdot \frac{1}{2}\mathbf{l}_M[U-j],$$

where $\mathbf{c}_0$ denotes the vector describing the equilibrium distribution of midpoints, $\mathbf{C}^t\mathbf{c}_0[i]$ denotes the i-th term in the vector resulting from iterated application of $\mathbf{C}$ to $\mathbf{c}_0$, and $\mathbf{l}_M[i]$ denotes the i-th term of the leading left eigenvector of $\mathbf{M}$, that is, the frequency of individuals with range size iv.

Finally we calculate the time in generations until $1/N$-th of the population has a norm of reaction that contains Ω_{new}. Recall that N is the total population size. This is the earliest time at which the distribution of upper bounds is such that at least one individual is expected to have an upper bound greater than Ω_{new}. In other words, for a given distance $|\Omega_{new} - \Omega_{old}|$, we seek the minimum time in generations t since the environmental transition such that $P_t[u_i \geqslant \Omega_{new}] \geqslant 1/N$.

Figure 7 plots τ, the time in generations until a population first encounters the new phenotypic optimum as a function of the distance between the previous and new environments $\Delta\Omega = |\Omega_{new} - \Omega_{old}|$ and μ. Figure 8 provides two-dimensional slices at $\mu = 0.1$, 0.3, and 0.5 (from smallest to largest discs). This analysis assumes $N = 1000$, $v = 0.05$ and $\sigma = 0.25$. Intuitively, smaller μ and larger $\Delta\Omega$ yield longer search times. When $\mu \leqslant 0.1$, an environmental leap of $\Delta\Omega > 0.35$ would require search times of over 100 generations before a single individual would encounter the new optimal phenotype. When $\mu \leqslant 0.1$ and $\Delta\Omega < 0.1$, at least one norm of reaction already contains Ω_{new} at the time of the environmental transition. Note that smaller σ entails a higher proportion of range extending mutation. This, like increasing μ, would speed up the population exploration of phenotype space and thereby yield shorter search times.

FROM FIRST ENCOUNTER TO POPULATION CONVERGENCE

Next we find the expected time between the first encounter of Ω_{new} and the subsequent equilibrium. This analysis breaks down into two components: the time until all individuals realize

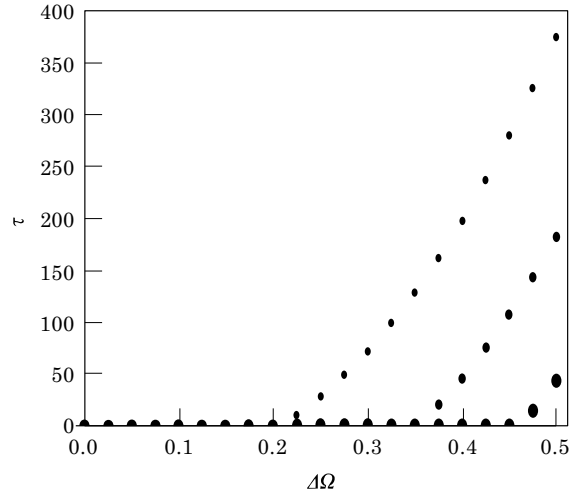

FIG. 8. Time until population first encounters new optimum for specific μ: we again assume $N = 1000$, $v = 0.05$ and $\sigma = 0.25$ and graph two-dimensional slices of Fig. 7 for $\mu = 0.1$ (small discs), $\mu = 0.3$ (medium discs) and $\mu = 0.5$ (large discs).

the optimal phenotype, and the subsequent time until the norms of reaction converge around the optimal phenotype to reach the new equilibrium. Since we know the distribution of range sizes before the population encounters Ω_{new}, we can approximate the fitness of the first individual i such that $\Omega_{new} \in \rho_i$, and the average fitness of the rest of the population. Because the first individual to reach Ω_{new} is likely to have a large norm of reaction, we assume conservatively that i has the largest phenotypic range in the population. Consequently, the rest of the population makes up the $(N-1)/N$-th quantile of the distribution.

Suppose $\omega_{\hat{I}}$ and $\omega_{\hat{O}}$ represent the fitness of i and the average fitness of the rest of the population, respectively. To approximate the time in generations until the entire population has norms of reaction containing Ω_{new}, we compare the population growth rate of the descendants of i with that of the rest of the population. This ignores mutation, and assumes that the i lineage homogeneously sees Ω_{new} while the other lineage does not.

More specifically, let $\hat{I}_t$ be the frequency of individuals at time t that descend from i, and therefore have Ω_{new} within their norm of reaction. Also let $\hat{O}_t = 1 - \hat{I}_t$ be the frequency of individuals that do not see Ω_{new}, and therefore are derived from other lineages.
Then

$$\hat{I}_{t+1} = \frac{\hat{I}_t \omega_{\hat{I}}}{\hat{I}_t \omega_{\hat{I}} + \hat{O}_t \omega_{\hat{O}}}$$

and

$$\hat{O}_{t+1} = \frac{\hat{O}_t \omega_{\hat{O}}}{\hat{I}_t \omega_{\hat{I}} + \hat{O}_t \omega_{\hat{O}}},$$

hence

$$\frac{\hat{I}_{t+1}}{\hat{O}_{t+1}} = \frac{\hat{I}_t \omega_{\hat{I}}}{\hat{O}_t \omega_{\hat{O}}}.$$

If we let $t = 0$ when individual i first encounters Ω_{new}, then the initial ratio $(\hat{I}_0/\hat{O}_0) = (1/(N-1))$. Hence $(\hat{I}_t/\hat{O}_t) = (1/(N-1))(\omega_{\hat{I}}/\omega_{\hat{O}})^t$. The estimated time until all individuals have the optimal phenotype is the minimum t such that $(\hat{I}_t/\hat{O}_t) \geqslant N$. We calculate then,

$$t_{eqm} = \log_{\frac{\omega_{\hat{I}}}{\omega_{\hat{O}}}} N(N-1) = \frac{\log N(N-1)}{\log \frac{\omega_{\hat{I}}}{\omega_{\hat{O}}}}.$$

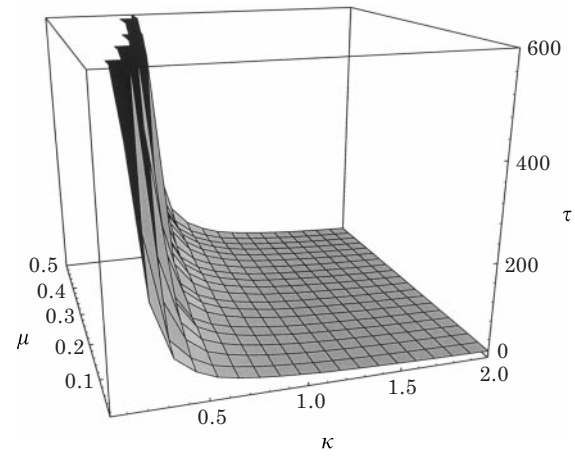

FIG. 9. Time until entire population encounters new optimum: the expected time between the first encounter with Ω_{new} and all i having $\Omega_{new} \in \rho_i$ in terms of the mutation rate μ and the fitness benefit κ.

The rate at which the population moves toward the new optimal phenotype depends on the ratio $\omega_{\hat{I}}/\omega_{\hat{O}}$. This ratio gets larger as κ, the benefit delivered to the i lineage by Ω_{new}, increases. When κ outweighs the cost of the relatively large norm of reaction of i, the i lineage eventually overtakes the population, bringing it towards the next equilibrium. The mutation rate of the population influences the initial distribution of range sizes, hence the relative fitnesses. Figures 9 and 10 illustrate the effect of κ and μ on the expected time to equilibrium for $N = 1000$, $\sigma = 0.25$ and $v = 0.05$. For a given mutation rate, the graph of expected time to the

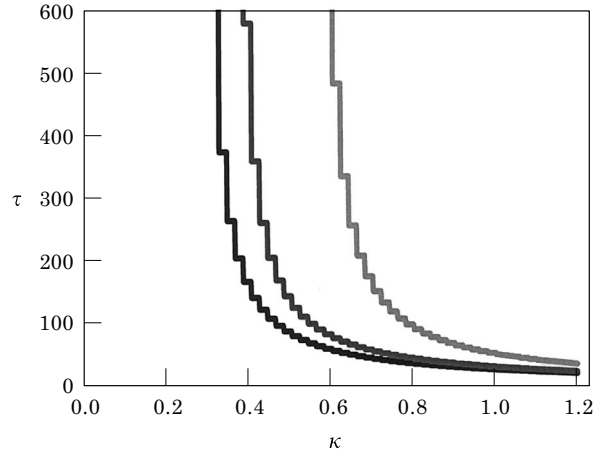

FIG. 10. Time until entire population encounters new optimum for specific μ: the expected time between the first encounter with Ω_{new} and all i having $\Omega_{new} \in \rho_i$ in terms of κ for $\mu = 0.1$ (dark line), $\mu = 0.2$ (medium line) and $\mu = 0.5$ (light line).

new optimum asymptotes vertically at the value κ_α such that $\omega_{\hat{I}} = \omega_{\hat{O}}$. For $\kappa \leqslant \kappa_\alpha$ the advantage to the new optimal phenotype is not strong enough to pull the population to the new domain. As κ increases from κ_α the rate of evolution towards the new equilibrium increases dramatically. Larger μ yields more extended distributions, hence a larger norm of reaction for i. So as μ increases, the magnitude of κ necessary for the success of the i lineage also increases.

When is this approximation reasonable? This analysis rests on two levels of estimation. First the stochasticity of the model limits us to knowing only the expected ratio of growth rate. Further we assume independent lineages, and thereby ignore both mutations that change norm of reaction size, and consequently fitness, and mutations that move an i descendant away from Ω_{new} or an $\hat{O}$ individual onto Ω_{new}. Individuals in i lineage that mutate away from Ω_{new} will be short-lived because they will have long norms of reaction and suboptimal phenotypes. Likewise $\hat{O}$ individuals that increase in length will be unsuccessful. The remaining two effects of mutation would improve the fitness of the affected individuals and consequently hasten the population towards the new equilibrium. Large mutation rates therefore would most likely cause our calculation to over-estimate the time to equilibrium. The above approximations of the time between environmental transitions necessary for the Baldwin–Simpson effect are therefore conservative.

In the final component of the analysis, we estimate the time to equilibrium after the population has moved over to Ω_{new}. We make the very rough assumption that the population uniformly has a range size equal to that of the first individual to encounter Ω_{new}. We again allow this individual (i), and hence the entire population, to have the largest norm of reaction in the distribution given by the leading left eigenvector of M. We then ask how long until the average range size of the population converges to equilibrium conditions. Let $\mathbf{u}_{\Delta\Omega}$ be the vector describing the assumed distribution of range sizes. That is, $\mathbf{u}_{\Delta\Omega}$ has all 0 terms except a 1 at position describing the range size of the original i. Then we solve for the minimum τ such that $\mathbf{M}^\tau \mathbf{u}_{\Delta\Omega}$ attains the equilibrium average range size described by the leading left eigenvector of $\mathbf{M}$. The results are depicted in Fig. 11.

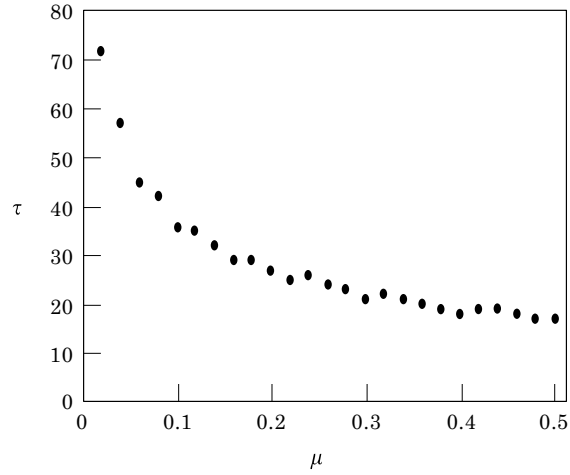

FIG. 11. Time to convergence on equilibrium: this graph depicts the time from the moment at which the entire population has Ω_i within their norm of reaction until the norms narrow to equilibrium as a function of mutation rate μ. We assume $v = 0.05$ and $\sigma = 0.25$.

The Simpson–Baldwin Effect

Together the two processes—the search for the new optimal phenotype and the subsequent convergence on the phenotype—determine the transition time between equilibria in consecutive environments. Immediately following the initial transition and until a single individual encounters Ω_{new}, the population will drift, maintaining a narrow distribution of phenotypic ranges. This is illustrated by the region around Fig. 5(a) mark d, where the average norm of reaction length remains small while no individual encounters the optimal phenotype. The interim spikes in average norm length, seen at Fig. 5(a) mark b, occur after the first encounter with the new optimum. The first individuals to have Ω_{new} inside their norm of reactions typically lie on the very high end of the distribution of range sizes. Strong selection for these optimal types temporarily skews the distribution towards more extended norms of reaction. After the population swings over to the new optimum, most individuals realize the fitness benefit of Ω_i, and selection for decreased norm length narrows the distribution to the new equilibrium.

We can now estimate the minimum time necessary between environmental transitions for

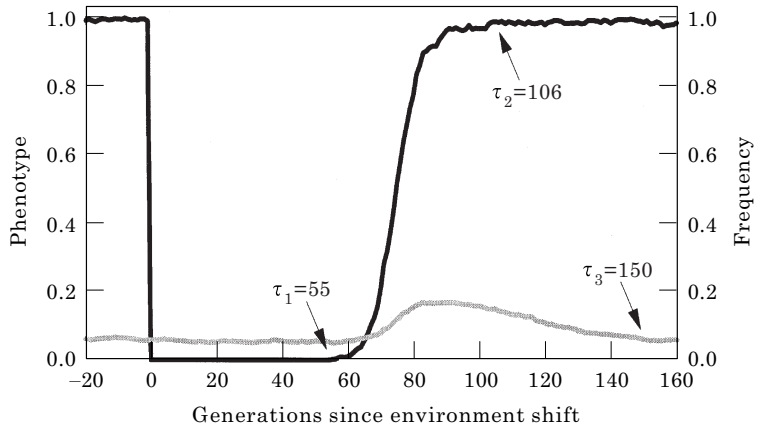

FIG. 12. Simulation results: this data is produced by simulation to and beyond an environmental transition from $\Omega_{old} = 0.3$ to $\Omega_{new} = 0.6$, assuming $N = 1000$, $\mu = 0.1$, $\sigma = 0.25$ and $\kappa = 0.5$. τ_1, τ_2 and τ_3 mark the generations at which the first individual encounters Ω_{new}, almost all norms of reaction contain Ω_{new}, and the new equilibrium is attained, respectively.

a Simpson–Baldwin effect to occur. Suppose $N = 1000$, $\mu = 0.2$, $\sigma = 0.25$, $v = 0.05$ and $\kappa = 0.5$. If a population experiences an environmental shift of magnitude $|\Omega_{new} - \Omega_{old}| = 0.3$, how long will it take to reach a new equilibrium? By the methods described above, we find that the time until one individual is expected to have a norm of reaction containing Ω_{new} is 70 generations (Fig. 8), the estimated time until all individuals see Ω_{new} is 100 generations (Fig. 10), and the approximate number of generations from there to equilibrium is 30 (Fig. 11). The environment must therefore remain constant for roughly 200 generations for a range expansion and subsequent Simpson–Baldwin narrowing to take place.

Figure 12 gives the trajectory for one run of a population size $N = 1000$ with the above-mentioned parameter values. τ_1, τ_2 and τ_3 mark the three episodes estimated above. All three conform fairly well to the predictions.

There are many parametric conditions under which populations will never experience such channeling of phenotypic possibility. For small μ and moderate distances between environmental optima $\Delta\Omega$, the time before a single individual is expected to have a norm of reaction containing the new optimal phenotype is extraordinarily long (assuming still that $N = 1000$, $\sigma = 0.25$ and $v = 0.05$). For example, when the rate of environmental fluctuation ϵ exceeds 0.01, populations with mutation rate $\mu = 0.1$ will only exhibit the Simpson–Baldwin effect for environmental shifts $\Delta\Omega < 0.3$. This limitation explains the differences between the lines in Fig. 5(b). Under those conditions, populations experiencing $\epsilon > 0.05$ never display the minimum average range characteristic of the equilibrium endpoint of a Simpson–Baldwin effect.

The trade-off between mutation rates and the benefit delivered by the optimal phenotypes also restricts the dynamics. The second part of the Simpson–Baldwin effect—convergence on the new optimal state—will never occur when κ is too small to outweigh the mutation pressure towards expanding norms of reaction.

Although we define a Simpson–Baldwin effect as a complete transition from one equilibrium to the next, we are also interested in alternative outcomes. By the methods described above, we can also approximate threshold ϵs such that the population might encounter optimal phenotypes and undergo only partial narrowing of norms of reaction between environmental transitions, and ranges of ϵ under which populations maintain constant large ranges, never having time to happen upon or converge to new optima.

Conclusions

Simpson's version of the Baldwin effect describes how natural selection in a changing environment may produce a seemingly Lamarckian transition. On the surface, it seems that initially a few individuals acquire—but not through genetic inheritance—a new and beneficial trait. Over time, the frequency of individuals that possess the trait increases.

Almost simultaneously, alternative states become less diverse. This continues until almost all individuals have almost exclusively the beneficial trait. What at first seems non-hereditary eventually reveals a genetic component. Our model demonstrates that this phenomenon can be explained by natural selection on phenotypic plasticity. A Simpson–Baldwin narrowing of norms of reaction occurs when the environment shifts to favor a new trait, followed by the mutation of a few norms of reaction to include that trait. Finally the fitness benefit of the new optimal trait outweighs the cost of phenotypic flexibility, causing selection for narrow ranges that contain the optimal trait.

How might this theoretical model apply to the real world? Phenotypic plasticity is widely documented from plant morphological variation in different climates to higher mammalian learning to biochemical modification in bacteria. The norm of reaction intervals may represent a spectrum of leaf size, of foraging techniques, or of enzyme concentrations. Environmental change may be climatic or chemical or may result from niche modification by the organisms themselves. There are costs associated with the flexibility to adapt to such changes. In *E. coli*, for example, there is a tension between the ability to reproduce quickly and the genetically programmed ability to adapt to a rapidly changing environment. The cost of such versatility is extra genetic baggage.

This simple model invites application to biological systems and further theoretical exploration of natural selection on plasticity. What happens in more elaborate environments in which there are multiple optima or other more complicated maps from phenotype space to fitness? How would different cost functions for norm of reaction length change the dynamics? In its current form, the model illustrates that in a rapidly shifting environment, it pays to have phenotypic plasticity. When fluctuations are rare, populations repeatedly experience a Simpson–Baldwin effect.

I extend my utmost gratitude to M. Lachmann and G. Sella for their invaluable technical advice, and to M. W. Feldman for numerous insightful suggestions. This research is supported in part by a U.S. National Defense Science and Engineering Fellowship to L. Ancel, and NIH grant GM28016 to M. W. Feldman.

REFERENCES

Baldwin, J. M. (1896). A new factor in evolution. *Am. Nat.* **30,** 441–451, 536–553.

Hinton, S. J. & Nolan, G. E. (1980). How learning can guide evolution. *Complex Systems* **1,** 495–502.

Morgan, C. L. (1896). *Habit and Instinct*. London: Arnold.

Osborn, H. F. (1896). A mode of evolution requiring neither natural selection nor the inheritance of acquired characteristics. *Trans. New York Acad. Sci.* **15,** 141–142, 148.

Schmalhausen, I. I. (1949). *Factors of Evolution* (Dordick, I., trans.; Dobzhansky, T., ed.). Philadelphia: Blakiston.

Simpson, G. G. (1953). The Baldwin effect. *Evolution* **7,** 110–117.

APPENDIX

Let c_i be the center of the norm of reaction ρ_i. During each time step, c_i may increase by $v/2$ through a range-expanding mutation of the upper bound or through a shrinking mutation of the lower bound. The rate of such changes to the center is therefore $\frac{1}{2}\mu(1-\sigma)+\frac{1}{2}\mu\sigma=\frac{1}{2}\mu$. Likewise the rate of decrease of c_i by $v/2$ is also $\frac{1}{2}\mu$. Each mutation alters the magnitude of a $|\rho_i|$, and therefore affects the fitness of the individual. By symmetry, mutational increments and decrements to c_i have identical fitness consequences.

In calculating the fitness costs of mutation, we divide the population into three categories—those with increased c_i, decreased c_i and constant c_i. Let α_μ be the ratio of the fitness of an individual with an incremented (or decremented) c_i to the average fitness of the population. If we approximate all ρ_i with the average range $|\bar{\rho}|$, then

$$\alpha_\mu=\frac{\frac{1}{2}\mu\sigma(1-(|\bar{\rho}|-v))+\frac{1}{2}\mu(1-\sigma)(1-(|\bar{\rho}+v))}{\mu\sigma(1-(|\bar{\rho}|-v))+\mu(1-\sigma)(1-(|\bar{\rho}|+v))+(1-\mu)(1-|\bar{\rho}|)}$$

$$=\frac{v(2\sigma-1)+1-|\bar{\rho}|}{\mu v(2\sigma-1)+1-|\bar{\rho}|}.$$

Note that we describe α_μ in the context of a population that has not encountered the optimal phenotype. Hence we ignore the environmental component of fitness and consider only range-size effects.

We can therefore approximate the evolution of the midpoints by $c_i' = c_i + b$, where

$$b = \begin{cases} -\frac{v}{2} & \text{with probability } \frac{1}{2}(\mu\alpha_\mu) \\ \frac{v}{2} & \text{with probability } \frac{1}{2}(\mu\alpha_\mu) \\ 0 & \text{with probability } 1 - \mu\alpha_\mu \end{cases}$$