Online Supplemental S1: Model description

for Solving the paradox of stasis: Squashed stabilizing selection and the limits of detection

Benjamin C. Haller[1] and Andrew P. Hendry[1]

1. Department of Biology and Redpath Museum, McGill University, 859 Sherbrooke Street West, Montreal, Quebec, Canada H3A 0C4

The model description presented here is structured according to the guidelines suggested by Grimm et al. (2006). The model is an individual-based non-spatial sexual model with non-overlapping generations. It was implemented in Objective-C using the Cocoa object-oriented framework (Apple Inc., http://www.apple.com, Mac OS X 10.6.8). All model parameters are summarized in Table 1.

Purpose

The purpose of this model is to produce full census records, including genetic and phenotypic trait values and survival records, of a population evolving on an invariant stabilizing fitness function and, optionally, also subject to negative frequency-dependent selection due to intraspecific competition. The model is designed to separate out the regulation of population size from selection involving the focal trait; it is assumed, in other words, that population size is regulated by other factors. Because the generated census records are intended to be analyzed using a mark-recapture protocol based upon common empirical practice, random mortality uncorrelated with the focal trait is also modeled. A neutral (unselected) trait with the same genetics as the focal trait, but physically unlinked with that trait, is included as a control.

Environment and state variables

Each individual in the model is defined by two traits, its selected trait and its neutral trait. For each individual, the selected trait has a genetic trait value a_s and an environmental trait deviation e_s; the individual's phenotype for the selected trait, z_s, is additively generated from these values ($z_s = a_s + e_s$). The neutral trait is similarly defined by genetic, environmental, and phenotypic values a_n, e_n, and z_n ($z_n = a_n + e_n$). An individual's genetic trait values a_s and a_n are themselves defined by one or many values, representing additive alleles, depending upon the underlying genetic architecture (see *Genetic architectures*). The state of the model at any point in time is therefore fully specified by the state (a_s, e_s, a_n, e_n) of all individuals, keeping in mind that a_s and a_n are themselves defined by their allelic values, which are therefore part of the model's state.

The environment in the model is non-spatial and stateless; its properties are governed by parameters discussed in the following subsections.

Genetic architectures

Three genetic architectures were used for the model realizations discussed here, for purposes of comparison and generality: a quantitative-genetics-based architecture ("quantitative"), a diploid 8-locus triallelic architecture ("triallelic"), and a diploid 8-locus continuum-of-alleles architecture ("continuum"). Each architecture will be discussed in turn, including the way in which inheritance and mutations are implemented for that architecture. For convenience this section will refer to any genetic trait values of an individual, whether a_s or a_n, simply as a; for a given realization of the model, the genetic architectures of the selected and neutral traits are always the same.

QUANTITATIVE ARCHITECTURE

The quantitative architecture implements each genetic trait value a according to the so-called infinitesimal model, using a single unbounded continuous value that represents the additive effects of an infinite number of loci of infinitesimal effect size (Fisher 1918; Bulmer 1980). Offspring of parents with trait values a_1 and a_2 receive a trait value a drawn from a normal distribution with its mean equal to the mid-parental value $\frac{1}{2}(a_1 + a_2)$ and with standard deviation $\frac{1}{2}|a_1 - a_2|$ (following, e.g., Heinz et al. 2009). If a mutation occurs (see *Reproduction phase*), the offspring trait value a is then altered by a mutational effect of a magnitude

drawn from a normal distribution with mean 0 and standard deviation α.

TRIALLELIC ARCHITECTURE

The triallelic architecture implements each genetic trait value a using 8 pairs of values, λ_{a1} and λ_{b1} through λ_{a8} and λ_{b8}, where each pair (λ_{an}, λ_{bn}) represents one diploid locus; each allele λ may take one of the three possible allelic values, –1, 0, or +1 (following, e.g., Thibert-Plante and Hendry 2011). The value of a is determined additively from the λ values, scaled by the square root of the number of alleles according to

$$a = \left(1/\sqrt{8\times 2}\right)\sum_{n=1}^{8}\left(\lambda_{an} + \lambda_{bn}\right) \quad \text{(Formula S1.1).}$$

Although many models scale by the number of alleles (e.g. Dieckmann and Doebeli 1999), such scaling has the effect that the maximum genetic variance possible without linkage disequilibrium decreases as the number of loci increases, which is undesirable since it appears to be biologically unrealistic (traits based upon a very large number of loci, such as human height, do not thereby have a variance close to zero). Scaling by the square root of the number of alleles ($\sqrt{8\times 2}$ in this case) ensures that the maximum genetic variance does not depend upon the number of loci upon which the trait is based (Dieckmann, U., and Meszéna, G., pers. comm.). Since versions of the model were tested with anywhere from 2 to 256 loci, this was an important consideration in the model's construction, but it is of less importance to the results we present here, since we here discuss only the 8-locus version of the triallelic architecture.

Offspring in the triallelic architecture inherit alleles from the parents with equal probability and without linkage, according to standard Mendelian inheritance; in other words, each parent will contribute either its λ_{an} or its λ_{bn} value to the offspring, with equal probability. If a mutation occurs for an allele in the triallelic architecture, that allele's value is altered by one step in a random direction; values of –1 and +1 therefore always change to 0, while a value of 0 will change to either –1 or +1 with equal probability.

CONTINUUM ARCHITECTURE

The continuum architecture implements each genetic trait value a using 8 pairs of values, λ_{a1} and λ_{b1} through λ_{a8} and λ_{b8}, where each pair (λ_{an}, λ_{bn}) represents one diploid locus (following, e.g., Yeaman and Guillaume 2009). Each allelic value λ may take any unbounded continuous (thus "continuum") value (Kimura 1965). The genetic trait value a is determined additively from the allelic values, scaled by the square root of the number of alleles of which the trait is composed, as in the triallelic architecture (see Formula S1.1). Offspring inherit allelic values as in the triallelic architecture. If a mutation occurs for an allele, the allelic value is offset by a draw from a normal distribution of mean 0 and standard deviation α, as in the quantitative architecture.

One difference between these genetic architectures is that the triallelic architecture has a bounded genetic range, whereas the quantitative and continuum architectures are unbounded. Care was taken to ensure that the phenotypic range of the population was generally within the bounds of the triallelic architecture, to minimize the importance of genetic boundary effects, but the difference nevertheless does exist.

Another difference is in the magnitude of the mutational variance experienced by the different architectures. The mutational variance, V_M, may be calculated using the standard formula $V_M = pn\mu\sigma_\alpha^2$, where p is the ploidy level (i.e., 2 for diploid), n is the number of loci (so pn is the number of alleles), μ is the mutation rate (the probability of a mutation, per allele), and σ_α^2 is the mutational variance per allele when a mutation does occur (Lynch 1988). For the quantitative genetics architecture, pn is taken to be 1, so $V_M = \mu\sigma_\alpha^2$, and σ_α^2 is the square of the mutational variance parameter α. For the diploid 8-locus continuum-of-alleles architecture, $\sigma_\alpha^2 = \alpha^2/pn$, because of the scaling of allelic effect size described above (see Formula S1.1; note the scaling of the allelic effect by $\sqrt{pn}$ becomes squared in the conversion from standard deviation to variance). Finally, for the diploid 8-locus triallelic architecture, $\sigma_\alpha^2 = 1/pn$, because mutated alleles always change in value by 1, with the allelic effect size then scaled as described above (see Formula S1.1). The realized mutational variances for the

genetic architectures and parameter values used are shown in Table S1.1.

Table S1.1. Mutational variances for the genetic architectures and parameter values used.

Architecture	μ	α	V_M
Quantitative	0.001	0.5	2.5×10^{-4}
		0.05	2.5×10^{-6}
	0.00001	0.5	2.5×10^{-6}
		0.05	2.5×10^{-8}
Triallelic	0.001	1.01	1.0×10^{-3}
	0.00001	1.01	1.0×10^{-5}
Continuum	0.001	0.5	2.5×10^{-4}
		0.05	2.5×10^{-6}
	0.00001	0.5	2.5×10^{-6}
		0.05	2.5×10^{-8}

[1]For the triallelic model, α is not defined as a parameter; but the mutational effect size, in the sense of parameter α, is 1.0, so it is shown as such here.

The mutational heritability $h_M^2 = V_M/V_E$ may easily be calculated from these mutational variances (Lynch 1988; Houle et al. 1996), given the value of V_E for a particular realization of the model. These mutational heritability values range from 1.0 to 2.5×10^{-7} for the parameter values used. This range of values is wider than the empirically observed distribution, which ranges from approximately 3×10^{-2} to 1×10^{-4} (Houle et al. 1996). The wide range of mutational heritabilities in our model is a consequence of our deliberate use of a wide range of mutational variances that span the values used by other models (see *Parameters*). However, more than half of our realizations fall within the range of empirically observed values. Moreover, very few values have been observed for empirical systems thus far, and so the actual range of mutational heritabilities in nature is doubtless wider than has yet been documented.

Initialization

Initially, the model contains N_j individuals, with each individual receiving selected trait and neutral trait genetic values, a_s and a_n, according to the genetic architecture. For the quantitative architecture, initial genetic trait values are drawn from a uniform distribution between –1 and +1. For the triallelic architecture, each allele λ has an equal and independent probability of being either –1, 0, or +1. For the continuum architecture, each allele λ is drawn from a uniform distribution between –1 and +1. These rules result in initial genetic distributions that are reasonably similar between the different architectures, and that in all cases have substantially more genetic variance than is observed in the model following the "burn-in" period (see Methods, *Data collection*). The equilibration of the model during burn-in therefore does not depend principally upon the accumulation of new mutations, which would be slow and very stochastic, but instead principally upon the sorting of the initial standing genetic variation. If there is a bias introduced by the choice of initial state, it is thus towards a stable, consistent equilibrium, and is thus conservative with respect to our hypothesis.

The initial selected trait and neutral trait environmental values, e_s and e_n, are drawn from a normal distribution with mean 0 and variance V_E, just as for new offspring (see *Reproduction phase*).

Interactions

Individuals interact only through negative frequency-dependent selection, here conceptualized as competition although other ecological interactions can also generate this pattern of selection (see Discussion, *Squashed stabilizing selection*). The model of competition used is based upon that of Roughgarden (1972), assuming a Gaussian competition kernel as in many similar models (e.g., Slatkin 1979; Doebeli and Dieckmann 2003), with modifications explained below.

Here z_i will refer to the selected trait value z_s of an individual i, for notational simplicity, since the neutral trait is not involved in interactions. The strength of competition felt by a focal individual i is the mean of the strengths of competition due to every other individual, scaled by the intensity of competition, c, as given by

$$c_{\text{eff}} = \frac{c}{N-1} \sum_{j=1,\, j\neq i}^{N} e^{-(z_i - z_j)^2 / 2\sigma_c^2} \qquad \text{(Formula S1.2).}$$

The mean is used, rather than the sum divided by the carrying capacity as in many other models, because competition is not meant to act as a regulator of

population size in this model, and thus the strength of competition should not depend upon the number of other individuals present; in other words, the competition modeled is frequency-dependent but not density-dependent. From an implementation perspective, this is principally so that when competition is turned off the population size continues to be regulated in essentially the same manner as when it is turned on; the population size will be smaller in the former case than in the latter, because competition causes additional mortality, but the magnitude of that decrease will depend only on the phenotypic distribution of the population, not on the population size, and it is typically not large. From a biological perspective, this choice could reflect, for example, a situation in which competition is exerted by other individuals that are nearby in space, such that (with a constant spatial density of individuals determined, for example, by territorial behavior, or by density regulation based upon traits uncorrelated with the modeled trait) the average strength of competition over a generation depends upon the phenotypic distribution in the population, but not upon the population size.

According to this formula, the effective strength of competition felt by an individual, c_{eff}, varies proportionally with the intensity of competition c, a parameter not present in the model of Roughgarden (1972), which instead has the intrinsic rate of increase r at this position in the equation. This difference is again due to the change in how the population is regulated, as discussed above; in effect, the Roughgarden (1972) model varies the intensity of competition according to the population size compared to the carrying capacity, while our model has a separate parameter for that intensity because the population regulation is separated from the effect of competition.

The effective strength of competition between two individuals of identical phenotype is 1; this decreases with increasing phenotypic dissimilarity between the individuals, as described by a normal function with standard deviation σ_c, the width of the phenotypic competition function. As the effective strength of competition felt by an individual increases, that individual's probability of selective death increases (see *Selection phase*).

Process overview

Generations in the model are non-overlapping, and are divided into three phases which occur sequentially: the mortality phase (in which random deaths occur, unrelated to the trait under selection), the selection phase (in which deaths occur due to selection on individuals), and the reproduction phase (in which the generation of juveniles for the next generation occurs, by random mating of the survivors of random mortality and selection). These phases are described in the subsections that follow.

The separation of random mortality and selective mortality corresponds to a biological life cycle in which the focal trait develops relatively late (i.e., after juvenile mortality); this design allows independent variation and measurement of these two types of mortality. A model combining random and selective mortality into a single phase, or reversing their order, would be expected to produce similar results, since random mortality should not change the genotypic or phenotypic distribution (apart from small, unbiased stochastic effects). The strength of competition would be unaltered by such changes, since competition in the model is not density-dependent.

Mortality phase

In the mortality phase, each individual has the same chance of dying, governed by the probability of random mortality m. Individuals that fail their probability "roll" are immediately removed from the population.

Selection phase

In the selection phase, each individual has a probability of "selective death" that depends upon its absolute fitness $W(z_s)$, which depends, in turn, upon its selected trait phenotype as compared to the optimum phenotype in the environment, θ, as given by

$$W(z_s) = e^{-(z_s-\theta)^2/2\omega^2} \qquad \text{(Formula S1.3).}$$

As this formula shows, absolute fitness is based upon the value of a Gaussian function with optimum θ (the phenotypic optimum) and width parameter ω (analogous to a standard deviation; the "strength" of selection), evaluated at the phenotypic value of the

focal individual and scaled such that an individual with phenotype $z_s = \theta$ has an absolute fitness $W(z_s) = 1$. Absolute fitness decreases from this maximum value as the deviation from the optimum phenotype increases.

Formula S1.3 represents the case of frequency-independent selection, as used in the model without competition. If competition is enabled in the model, the absolute fitness also depends upon the effective strength of competition, c_{eff}, felt by the individual, as

$$W(z_s) = \frac{e^{-(z_s-\theta)^2/2\omega^2}}{1+c_{eff}} \quad \text{(Formula S1.4).}$$

The effective strength of competition exerted by other individuals depends upon their phenotypic similarity to the focal individual (see *Interactions*); it thus gives rise to negative frequency-dependent selection that is combined multiplicatively with the stabilizing fitness function of Formula S1.3. If the effective strength of competition is zero, this formula reduces to Formula S1.3; as the effective strength of competition increases, absolute fitness decreases.

Once the absolute fitness of every individual in the population has been determined, selective deaths then occur, with the probability of death for each individual, $1 - W(z_s)$, depending upon that individual's fitness. Individuals that die due to selection are immediately removed from the population.

Reproduction phase

In the reproduction phase, N_j juveniles (meaning simply individuals prior to random mortality and selection) are generated, after which all of the adults in the population die (are removed from the population). Each juvenile is produced by the mating of two adult parents, randomly chosen from the population with uniform probability. Since individuals in the model are hermaphroditic, no consideration of gender is implemented. After mating the parents remain in the population, and are thus free to mate again, should they be so lucky.

The selected trait and neutral trait genetic values a_s and a_n of the juvenile offspring are calculated according to the genetic architecture, as described in *Genetic architectures* above. After the alleles inherited from the parents have been determined, for each allele in the genetic architecture there is a probability μ of that allele mutating (or for the quantitative architecture, which does not model individual alleles, μ is the probability of a mutation occurring across the genome as a whole). When a mutational event occurs, its consequences are determined according the rules in *Genetic architectures*.

The offspring's selected trait and neutral trait environmental deviances, e_s and e_n, are drawn from a normal distribution with mean 0 and variance V_E. Phenotypic values z_s and z_n are then determined additively, as described in *Environment and state variables*.

At the end of the reproduction phase the generation counter is advanced by one, and the next generation begins with the mortality phase.

Stochasticity

This model has several stochastic elements. The initial state of the model is stochastic (see *Initialization*), although the burn-in period (see Methods, *Data collection*) is intended to minimize the consequences of this. Demographic stochasticity is present due to finite population size, which affects the distribution of survivors in the mortality phase and selection phase, as well as the distribution of individuals that reproduce in the reproduction phase. Offspring genotypes vary stochastically from midparental genetic values due both to the random assortment of parental alleles and to the effects of mutation (see *Genetic architectures* and *Reproduction phase*). Offspring trait environmental deviations vary stochastically, which is intended to model random environmental influences and developmental stochasticity (see *Reproduction phase*). Phenotypic values thus include stochastic effects from both those environmental deviations, and from the underlying genetic values.

Observables

Each realization of the model generates quite a large amount of data. At the beginning of each generation, prior to the mortality phase, the genetic and phenotypic values of both the selected and the neutral trait of every individual are saved; in other words, (a_s, z_s, a_n, z_n) are saved for every individual.

At the beginning of the reproduction phase of every generation, the survival status of every individual that was in the population at the beginning of the generation is saved: whether that individual died during the mortality phase, died during the selection phase, or survived to the reproduction phase. These records together consumed approximately 800 MB of disk space per realization of the model with a population size of 1000 individuals, and 2 GB per realization with a population size of 2500 individuals (compressed). The full raw dataset generated by all realizations of the model is thus well over 1 TB in size, and therefore cannot be made available through an online data repository. However, anyone willing to ship a large hard drive to the corresponding author, B.C. Haller, with return postage paid may obtain a copy of the raw dataset. The dataset made available on Dryad is the (much smaller) result of the first pass of analysis over the raw data (see Methods, *Data analysis*).

The state of all individuals, and various summary statistics, are also observable graphically during model runs, for purposes of testing and experimenting with the model (Grimm 2002).

Parameters

The parameters of the model are given in Table 1, with symbols, units and values. Where multiple values are listed, model realizations were executed for each value listed, across all listed values of all other parameters, except where otherwise noted (see Methods, *Data collection*).

The strength of selection, ω, plays the same role as the parameter σ_K in the model of Dieckmann and Doebeli (1999), while our σ_c is the same as their σ_c. Values were chosen for ω and σ_c to produce both weak competition that does not lead to disruptive selection ($\sigma_c > \omega$), and strong competition that produces disruptive selection ($\sigma_c < \omega$), although evolutionary branching is not possible in this model since there is no mechanism by which assortative mating could evolve (Dieckmann and Doebeli 1999). The effective strengths of stabilizing selection experienced in realizations of the model once a pseudo-equilibrium has been reached appear to be realistic (see Supplemental S2, *The strength of stabilizing selection*).

There has been some debate in the literature regarding the proper mutation rate in models of this type (Doebeli and Dieckmann 2005; Gavrilets 2005; Waxman and Gavrilets 2005a, b). We therefore tested both high and low values for μ, the mutation rate per locus. The mutational effect size α was also varied for the quantitative and continuum genetic architectures. Together with the effects of the genetic architecture itself, this led to a wide range of mutational variances and mutational heritabilities tested (see *Genetic architectures* and Table S1.1), allowing us to assess the robustness of our conclusions to variation in these parameters (see Results and Supplemental S2, *Effects of mutational variance*).

The population size of 1000 used for most realizations was chosen primarily for computational efficiency, but supplemental realizations conducted with a population size of 2500 showed larger population size to be unimportant for our conclusions (see Supplemental S2, *Effects of large population size and sample size*).

The range of environmental variances used produced a wide range of emergent heritabilities that span the observed empirical range (see Supplemental S2, *Effects of heritability*).

The three genetic architectures used in our realizations span a broad range of possible architectures, and comparison among them indicates that genetic architecture is unimportant to our conclusions, although quantitative differences do exist (see *Genetic architectures*; Supplemental S2, *Effects of mutational variance*; Supplemental S2, *Effects of parameters on the selection gradient distribution*).

The random mortality rate was chosen with the aim of conservatism with respect to our hypothesis. A random mortality rate of 0.0, in which the focal trait is the only trait that affects fitness, is clearly unrealistic, but sets a baseline for comparison by showing what selection would be detected without any obscuring influence of uncorrelated mortality. The highest random mortality rate used, 0.5, is likely still lower than would be observed for many species in nature; however, its effects in obscuring the detection of selection were already so strong that larger, perhaps more realistic values would be pointless. The effect of random mortality is made

clear by the values chosen, and the effect of higher values is easy to conjecture.

Finally, the mark-recapture sizes used are intended to reflect the range of sizes typically used in empirical studies (Kingsolver et al. 2001). A few studies have used larger sizes, and these may be compared to our supplemental realizations with larger population size (see Supplemental S2, *Effects of large population size and sample size*).

Model samples

Snapshots of a modeled population adapted to various fitness functions are shown in Fig. S1.1. Videos of the running model, without and with competition, are provided as Movie S1.1 and Movie S1.2.

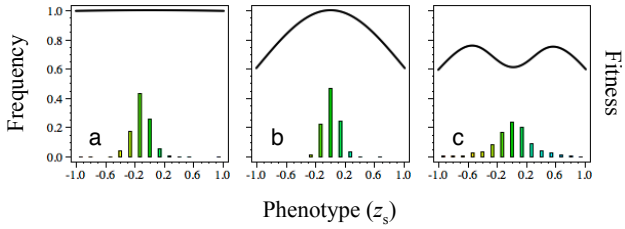

Figure S1.1. Model snapshots of typical phenotypic distributions (bars) and fitness functions (curves): (a) no competition, $\omega = 10$; (b) no competition, $\omega = 1$; (c) competition, $\omega = 1$, $\sigma_c = 0.2$. Other parameter values: $m = 0.0$, $\mu = 0.001$, $\alpha = 0.5$, $V_E = 0.01$, $N_j = 1000$. Despite the breadth of the peak in (a), the population's mean phenotype essentially never leaves the interval (−1, 1). The narrower width of the peak in (b) constrains the population somewhat more, but wander still occurs due to mutational variance. In (c) the effects of strong, narrow competition dimples the top of the fitness function, while the overall stabilizing shape is preserved, an illustration of "squashed stabilizing selection" (Fig. 1e); note that despite the appearance of disruptive selection across most of the phenotypic range of the population, the population mean will stay under the dimple indefinitely due to the larger stabilizing structure of the fitness function. Observed values at these snapshots were: (a) $V_P = 0.024$, $V_G = 0.014$, $h^2 = 0.587$; (b) $V_P = 0.012$, $V_G = 0.002$, $h^2 = 0.194$; (c) $V_P = 0.090$, $V_G = 0.079$, $h^2 = 0.879$. Note the large differences in phenotypic variance, genetic variance, and heritability, despite the same genetic architecture and model parameters; this illustrates how the population is shaped by the selective regime. In standardized units of phenotypic standard deviation, the effective strength of selection at these snapshots was: (a) 645.5; (b) 9.1; (c) 3.3 (see *The strength of stabilizing selection*).

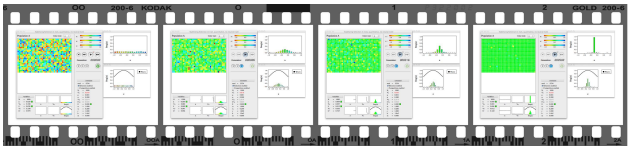

Movie S1.1: A few generations from a typical realization of the model without competition. Parameter values: $N_j = 1000$, $V_E = 0.01$, $\mu = 0.001$, $\alpha = 0.5$, $\theta = 0.0$, $\omega = 1.0$, $\sigma_c = 0.2$, $c = 1.0$, $m = 0.0$. Note the narrow phenotypic distribution compared to the fitness peak breadth, and the wander of the population around the fitness optimum. DOI: DOI HERE!

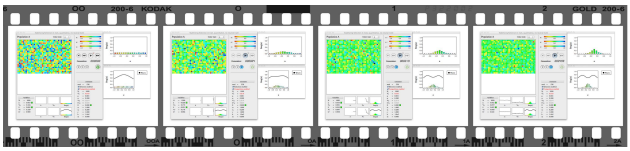

Movie S1.2: A few generations from a typical realization of the model with competition. Parameter values: $N_j = 1000$, $V_E = 0.01$, $\mu = 0.001$, $\alpha = 0.5$, $\theta = 0.0$, $\omega = 1.0$, $\sigma_c = 0.2$, $c = 1.0$, $m = 0.0$. Note the much broader phenotypic distribution relative to Movie S1.1, and the dimpled shape of the fitness function that indicates "squashed stabilizing selection" due to negative frequency-dependent selection. DOI: DOI HERE!

Literature Cited

Bulmer, M. G. 1980. The Mathematical Theory of Quantitative Genetics. Oxford University Press, New York.

Dieckmann, U., and M. Doebeli. 1999. On the origin of species by sympatric speciation. Nature 400:354–357.

Doebeli, M., and U. Dieckmann. 2003. Speciation along environmental gradients. Nature 421:259–264.

Doebeli, M., and U. Dieckmann. 2005. Adaptive dynamics as a mathematical tool for studying the ecology of speciation processes. J. Evol. Biol. 18:1194–1200.

Fisher, R. A. 1918. The correlation between relatives on the supposition of Mendelian inheritance. Trans. R. Soc. Edinb. 52:399–433.

Gavrilets, S. 2005. "Adaptive speciation"—It is not that easy: A reply to Doebeli et al. Evolution 59:696–699.

Grimm, V. 2002. Visual debugging: A way of analyzing, understanding and communicating bottom-up simulation models in ecology. Nat. Resour. Model. 15:23–38.

Grimm, V., U. Berger, F. Bastiansen, S. Eliassen, V. Ginot, J. Giske, J. Goss-Custard, T. Grand, S. K. Heinz, G. Huse, A. Huth, J. U. Jepsen, C. Jorgensen, W. M. Mooij, B. Muller, G. Pe'er, C. Piou, S. F. Railsback, A. M. Robbins, M. M. Robbins, E. Rossmanith, N. Ruger, E. Strand, S. Souissi, R. A. Stillman, R. Vabo, U. Visser, and D. L. DeAngelis. 2006. A standard protocol for describing individual-based and agent-based models. Ecol. Model. 198:115–126.

Heinz, S. K., R. Mazzucco, and U. Dieckmann. 2009. Speciation and the evolution of dispersal along environmental gradients. Evol. Ecol. 23:53–70.
Houle, D., B. Morikawa, and M. Lynch. 1996. Comparing mutational variabilities. Genetics 143:1467–1483.
Kimura, M. 1965. A stochastic model concerning the maintenance of genetic variability in quantitative characters. Proc. Natl. Acad. Sci. U. S. A. 54:731–736.
Kingsolver, J. G., H. E. Hoekstra, J. M. Hoekstra, D. Berrigan, S. N. Vignieri, C. E. Hill, A. Hoang, P. Gibert, and P. Beerli. 2001. The strength of phenotypic selection in natural populations. Am. Nat. 157:245–261.
Lynch, M. 1988. The rate of polygenic mutation. Genet. Res. 51:137–148.
Roughgarden, J. 1972. Evolution of niche width. Am. Nat. 106:683–718.
Slatkin, M. 1979. Frequency- and density-dependent selection on a quantitative character. Genetics 93:755–771.
Thibert-Plante, X., and A. P. Hendry. 2011. The consequences of phenotypic plasticity for ecological speciation. J. Evol. Biol. 24:326–342.
Waxman, D., and S. Gavrilets. 2005a. 20 questions on adaptive dynamics. J. Evol. Biol. 18:1139–1154.
Waxman, D., and S. Gavrilets. 2005b. Issues of terminology, gradient dynamics and the ease of sympatric speciation in Adaptive Dynamics. J. Evol. Biol. 18:1214–1219.
Yeaman, S., and F. Guillaume. 2009. Predicting adaptation under migration load: The role of genetic skew. Evolution 63:2926–2938.