doi:10.1111/j.1558-5646.2007.00303.x

EVOLUTION OF EVOLVABILITY IN A DEVELOPMENTAL MODEL

Jeremy Draghi[1,2] and Günter P. Wagner[1,3]

[1] *Department of Ecology & Evolutionary Biology, Yale University, New Haven, CT 06511*

[2] *E-mail: jeremy.draghi@yale.edu*

[3] *E-mail: gunter.wagner@yale.edu*

Received June 11, 2007
Accepted October 27, 2007

Evolvability, the ability of populations to adapt, can evolve through changes in the mechanisms determining genetic variation and in the processes of development. Here we construct and evolve a simple developmental model in which the pleiotropic effects of genes can evolve. We demonstrate that selection in a changing environment favors a specific pattern of variability, and that this favored pattern maximizes evolvability. Our analysis shows that mutant genotypes with higher evolvability are more likely to increase to fixation. We also show that populations of highly evolvable genotypes are much less likely to be invaded by mutants with lower evolvability, and that this dynamic primarily shapes evolvability. We examine several theoretical objections to the evolution of evolvability in light of this result. We also show that this result is robust to the presence or absence of recombination, and explore how nonrandom environmental change can select for a modular pattern of variability.

KEY WORDS: **Evo-devo, evolvability, modularity, quantitative genetics.**

Evolvability is the ability of populations to adapt through natural selection. This property is succinctly expressed by the observation that, through mutation, recombination and development, organisms can produce offspring that are more fit than themselves (Altenberg 1994): the study of evolvability posits that this observation is surprising and that it demands an explanation. The concept of evolvability promises to integrate ideas about constraint, phenotypic correlations, and mutational biases into a systematic theory of the variational properties that underlie evolution by natural selection. Evolvability may also provide an essential framework for understanding the success of invasive species (Gilchrist and Lee 2007) and the response of populations to anthropogenic change. Central to the idea of evolvability is the genotype–phenotype map: a set of rules that relate genotypes to the range of phenotypes they can produce (Alberch 1991; Wagner and Altenberg 1996). This mapping emphasizes how developmental systems create phenotypic variation from underlying genetic variation, and suggests two levels of processes contributing to evolvability.

Mutation and recombination contribute to evolvability by creating genetic differences between parents and offspring. Several investigations of the evolution of mutation reference the concept of evolvability (Radman et al. 1999; Tenaillon et al. 2001; Bedau and Packard 2003; Earl and Deem 2004; André and Godelle 2006), as do decades of studies on the evolution of sex and recombination (reviewed in Bell 1982; Otto and Barton 1997; Pepper 2003; Goddard et al. 2005). Although these studies have inspired much debate and deepened our understanding of the evolution of variability, they cover only a small fraction of the biological traits that shape evolvability.

A more diverse, and much less-explored, level of influences on evolvability includes the developmental processes that make genetic variation visible to natural selection. Evolvability is increasingly popular as a framework for interpreting a wide range of developmental traits at diverse scales: codon usage in genes (Plotkin and Dushoff 2003; Meyers et al. 2005), RNA structural evolution (Cowperthwaite and Meyers 2007), protein folding and stability (Wagner et al. 1999; Bloom et al. 2006), gene regulatory interactions (Wagner 1996; Tanay et al. 2005; Quayle and Bullock 2006; Tirosh et al. 2006), and angiogenesis and neural outgrowth in animal development (Kirschner and Gerhart 1998).

These two levels of influences on evolvability imply that many of an organism's traits might contribute to its evolvability. Insights from the experimental evolution of proteins (Aharoni et al. 2005; Khersonsky et al. 2006; O'Loughlin et al. 2006; Poelwijk et al. 2007), microorganisms (Burch and Chao 2000), and computer programs (reviewed in Adami 2006; Magg and Philippides 2006) are increasingly described in terms of evolvability, suggesting the exciting unifying potential of this idea. However, we still lack the theoretical tools to make rigorous measurements and comparisons of organismal evolvability, and the literature contains much confusion over the definition and utility of evolvability (Sniegowski and Murphy 2006; Lynch 2007). Central to this ambiguity is the question of whether evolvability can itself evolve by natural selection.

The evolution of evolvability through direct selection is currently a controversial hypothesis for the adaptability of organisms. One problem with this idea is that selection for evolvability seems to conflict with the apparent myopia of natural selection: the benefits to evolvability lie in an unknown future, perhaps beyond the ken of selection acting on contemporary phenotypes (e.g., Kirschner and Gerhart 1998; Poole et al. 2003; Earl and Deem 2004; Sniegowski and Murphy 2006). However, evolutionary biology contains several frameworks for understanding adaptation, such as geometric mean fitness (Stearns 2000) and lifetime reproductive success, in which selection, by integrating information about the past, appears to anticipate the future. Seen in this context, this objection to the evolution of evolvability is simply an empirical question about how well past environments predict future ones, and not a logical paradox.

Another frequently cited objection is that selection for evolvability requires group selection (e.g., Lynch 2007). Because evolution changes only populations, not individuals, it is superficially plausible that the rate of evolution could only evolve through competition among populations. This notion, however, was shown to be faulty over 25 years ago. Wagner (1981) modeled the dynamics of alleles that modify the rate of increase in mean fitness, $\bar{m}$. If p and q are the frequency of genotypes with modifier alleles I and II, respectively, then these modifier alleles evolve during adaptation in an asexual population according to

$$\frac{dp}{dt} = pq(\bar{m}_I - \bar{m}_{II}). \quad (1)$$

This means that if there is a difference between two genotypes in the rate of increase in the mean fitness, that is a difference in evolvability, the resulting difference in mean fitness is exactly the selective advantage caused by differences in evolvability. This result demonstrates that mutations affecting the rate or size of beneficial mutations are subject to individual-level selection. Recombination weakens but does not negate this form of selection for evolvability (Wagner and Bürger 1985), as discussed further below.

Another objection is that an evolvable genotype does not survive its own success: a genotype which produced a mutant that fixed in a population may have been evolvable, but it is now also extinct (Plotkin and Dushoff 2003). This difficulty may also be more apparent than real: just as offspring must merely resemble their parents for selection to cause evolution, evolvability must only be partially heritable to evolve. Again, this is an empirical question about organisms and other evolvable systems (see our discussion below and Plotkin and Dushoff (2003) for examples of how it can be answered).

The last major argument against the efficacy of selection favoring evolvability is that recombination will quickly dissociate an allele that improves variability from any positively selected variants it helps to create (Sniegowski and Murphy 2006). This argument may be damning for alleles that increase the mutation rate, or "mutator" alleles, in populations with any recombination (Tenaillon et al. 2000), implying that the evolutionary relevance of mutators is small at best (Sniegowski et al. 2000; de Visser 2002). Sniegowski and Murphy (2006) suggest that this result argues against all but special cases of evolvability loci with local effects, such as contingency loci in certain bacteria or transposons (de Visser 2002). Because these loci cannot be readily decoupled from any adaptive variants they create, recombination does not directly limit their successful fixation through indirect selection. Although it is unknown whether local mutation-modifying loci are the exception or the rule, it is worth noting that loci that affect development can influence variability through epistatic interactions with other loci. Such epistasis binds an evolvability-modifying allele to the beneficial alleles it facilitates, setting selection in opposition to recombination. Therefore, recombination may restrict, but not preclude, the evolution of evolvability through changes in epistasis.

Despite these apparent theoretical difficulties, several studies have made strong arguments for the evolution of evolvability by the evolution of developmental processes, in both biological systems and models (Ancel and Fontana 2000; Masel and Bergman 2003; Plotkin and Dushoff 2003; Masel 2005; Meyers et al. 2005). A common feature of many of these successful investigations is a tractable and explicit model of the relevant aspects of development: RNA folding in Ancel and Fontana (2000), mRNA translation in Masel and Bergman (2003), and the genetic code in Plotkin and Dushoff (2003) and Meyers et al. (2005). Although the genotype–phenotype maps of even simple organisms are largely unknown, analyzing the evolution of evolvability in a diverse array of simple, well-defined models of development may reveal the mechanisms of selection that shape the variabilities and developmental systems of organisms.

To dissect the influence of natural selection on evolvability, we constructed a model of development focused on the pleiotropic effects of two genes on two quantitative characters. Pleiotropy,

the effects of a single locus on multiple traits, is a ubiquitous feature of real developmental systems, and an elemental component of genotype–phenotype maps (Baatz and Wagner 1997; Hansen 2006). The effects of pleiotropy on adaptation have been studied since Fisher and are central to ideas about complexity and modularity in biology (Wagner and Altenberg 1996; Baatz and Wagner 1997; Orr 2000; Hansen 2003; Welch and Waxman 2003; Griswold 2006). Our model allows the evolution of both phenotypes and pleiotropic effects and therefore enables us to address the important question of how selection on the phenotype alters pleiotropy (Hansen 2006). We find that subjecting simulated populations of these model organisms to fluctuating selection favors pleiotropic relationships that minimize constraints on variability. We demonstrate that this outcome is the result of selection favoring evolvability, and that our results suggest a novel perspective on the influence of selection on the evolvability of developmental systems.

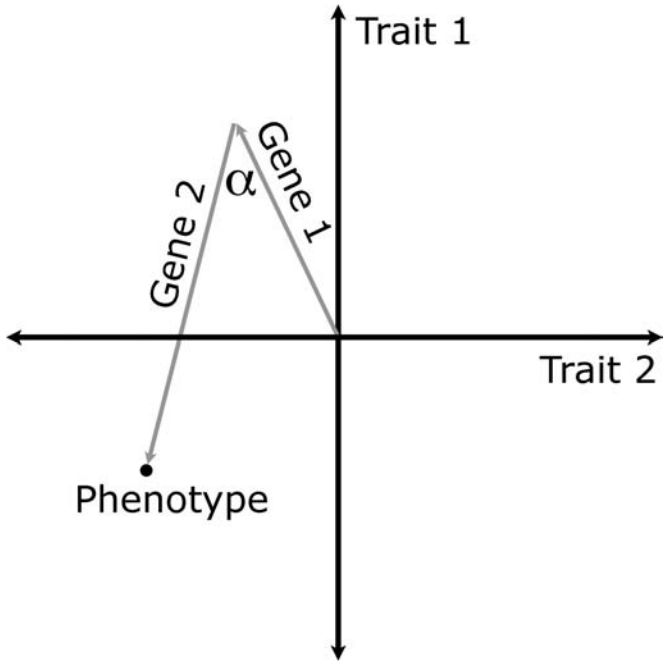

Figure 1. Illustration of the genotype–phenotype map used in the model. Two orthogonal quantitative traits define phenotypic space. Genotypes are pairs of vectors; phenotypes are the endpoints of the sum of those vectors. α is defined as the smallest angle between vectors.

Methods

MODEL DEFINITION AND SIMULATION CONDITIONS

Our model represents development as a geometric relationship between genotypes and phenotypes. Each phenotype is represented as a point on a plane, and can therefore be described by the values of two orthogonal, quantitative traits. Each genotype consists of two genes, each of which is a two-dimensional vector; the genotypic value of an organism is then the endpoint of the sum of these vectors, as depicted in Figure 1. Environmental variation is not modeled, so each genotype specifies a single phenotype. Importantly, this scheme captures the degeneracy of development: although each genotype specifies one phenotype, there are many genotypes that can produce each phenotype. This degeneracy may facilitate the evolution of evolvability by allowing genotypes that have the same phenotype to differ in evolvability.

We simulate evolution by first creating a population of N such genotypes, then defining an optimal phenotype, P. The fitness of each phenotype is inversely proportional to its distance from P: specifically, where d is this distance, each organism i is assigned a fitness

$$w_i = 1/(0.01 + d). \quad (2)$$

The next generation of the population is then formed in one of two ways. In simulations with no recombination, the algorithm picks N organisms, with replacement, from the current population. Each organism is selected with a probability proportional to its fitness, and is copied with mutation into the next generation. In simulations with free recombination, two different individuals are selected according to their fitnesses, and an offspring is generated by randomly choosing between the parental alleles for each of the four numbers defining the genome. As in the no-recombination simulations, parents are selected with replacement and genes are copied with mutation. In both modes of reproduction, two rates define the rate of mutagenesis: μ is the per-locus rate of mutation for the magnitude component of each vector, and μ_a is the corresponding rate for the angle component. The magnitude and angle associated with each vector also differ in the distribution of allowed mutations: magnitudes change by a uniformly distributed value between 0.5 and -0.5, whereas angles mutate to a new angle chosen uniformly from the entire interval [0, 360].

Our study examines cases in which angle mutations are considerably less common than magnitude mutations. If most mutations change the vector magnitude, then the mutant spectrum is dominated by changes along the length of each vector: variation is therefore largely constrained by the directions of the vectors. This pattern of constraint allows a biological interpretation of the characteristics of each gene. Magnitude mutations change the allelic state of a gene, and alter both phenotypic characters in a correlated manner. This correlation is equivalent to the vector angle: mutations in this angle therefore change the pleiotropy of a gene. Vector angles can consequently be viewed as a mutable aspect of development, and therefore as an aspect of variability that can evolve. To observe the evolution of variability, we need a measure of the pleiotropic correlations of both genes. This measure is simply the smallest angle between the two vectors, designated α.

Vector angles measure pleiotropy by the axes of phenotypic space, such as those drawn in Figure 1. However, this measure is only informative if these axes are meaningful. Because the orientation of the axes of a Euclidean space is unconstrained, any mutation can be seen as either affecting a single character (axis) or several, depending on the arbitrary orientation of the axes of the space. This ambiguity is simply a consequence of the many ways in which an organism's phenotype can be parsed into orthogonal traits. Because α measures vectors by their angular difference, it is independent of the choice of axes.

Because of this independence, α cannot measure the amount of pleiotropy. For example, a genotype with an α of 90° could exhibit no pleiotropy, maximum pleiotropy, or anything in between, depending on the choice of characters. What α does measure is the isotropy, or lack of bias, of the phenotypic spectrum of mutations. Vector angles describe the directions in phenotypic space in which mutational changes are most likely, and α captures the correlation between these two directions. α is inversely related to this correlation: genotypes with a small α therefore have a strong correlation between vectors, and can produce only a limited variety of mutants. In contrast, genotypes with high α values can produce a broad spectrum of variants that is unconstrained by the variational biases of the individual genes.

For most simulations presented, populations were initiated with a randomly generated population of N individuals. An optimal phenotype, P, was then uniformly drawn from within the unit disc, and the population evolved toward this optimum for 1000 generations. P was then redrawn from the same uniform distribution to simulate an environmental fluctuation. Evolution proceeded in 1000-generation intervals, punctuated by fluctuations to new optima, for a total of 1000 cycles, or one million generations. Several variants of this basic program were also used to test specific hypotheses, as described below.

All simulations were conducted with custom software written in C and C++, compiled using GNU C and C++ compilers, and executed on a Macintosh G5. The Mersenne Twister library was used for pseudorandom number generation (Matsumoto and Nishimura 1998). Mathematica was used for the Markov model presented below, and Microsoft Excel and R were used for data analysis.

Results

I. EVOLUTION OF α

Figure 2 shows α, or the angle difference between the vectors, for simulations performed over a range of μ_a where $\mu = 0.04$ and $N = 600$. Each point represents the mean of 24 replicates; error bars show one standard error. Each replicate is the median of at least 1000 measurements of the mean population α over a one-million-generation simulation. The horizontal gray line marks 45°, the

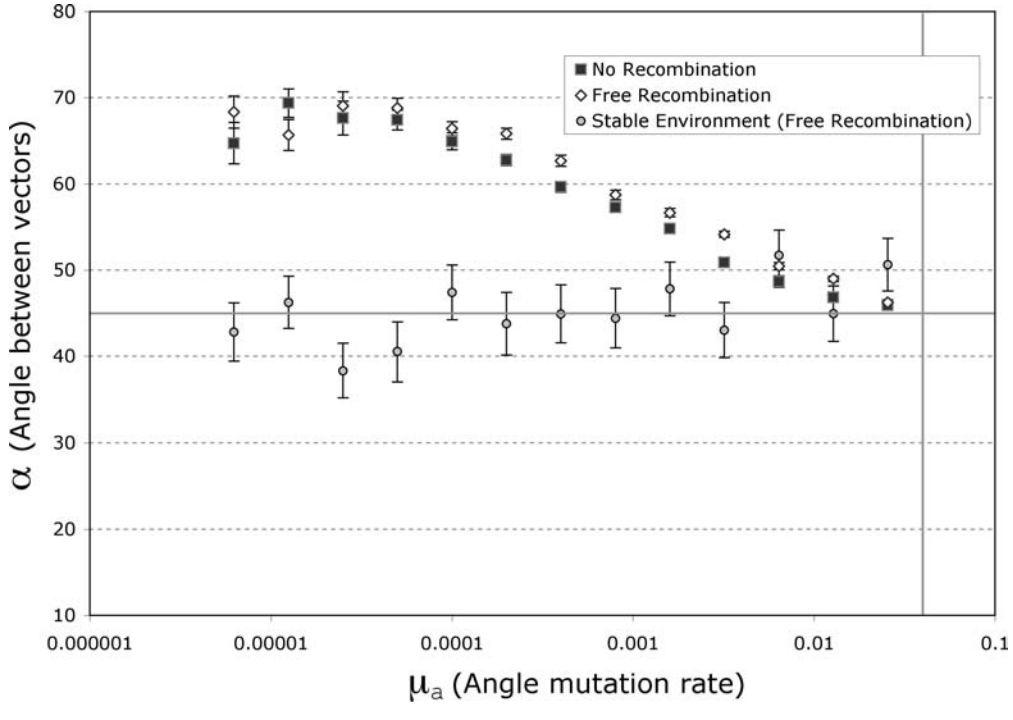

Figure 2. α, the angle difference between the vectors, as the angle mutation rate, μ_a, varies. Each point represents the mean of 24 simulations; each simulation is characterized by the median of 1000 measurements of mean population α. Recombination and its absence are implemented as described in the text. Error bars represent standard errors. The horizontal gray line marks the neutral expectation; the vertical gray line marks the magnitude mutation rate, $\mu = 0.04$. The population size, N, is 600 for all simulations.

null expectation if neither large nor small values of α are consistently favored. The vertical gray line marks $\mu = 0.04$, where μ is equal to μ_a. Open diamonds are trials with free recombination; closed squares represent asexual populations. Large values of α are clearly overrepresented in simulations where μ_a is much smaller than μ. In contrast, there is no trend toward any particular angles for either vector individually (data not shown). Recombination increases the divergence from 45°, but does not qualitatively affect the results. Because this method averages α across members of a population, it may underestimate the occurrence of extreme values of α in diverse populations: additional simulations that avoided averaging α produced nearly identical results (not shown). Appendix 1 demonstrates that this preponderance of high values of α is robust to changes in model parameters.

Figure 3A displays histograms of α, pooled across 24 replicates, for several values of μ_a; only the results for the asexual populations are shown. As Figure 3 shows, these distributions are strongly skewed when $\mu_a \ll \mu$: because of this skew, medians were used to characterize the central tendency of each simulation in Figure 2. These distributions confirm that large values of α are favored when μ_a is small. To understand the source of this preference, we next examined the pattern of substitutions in the angle component of each gene. Two trends could explain the preference for larger α: we may observe genotypes with higher α invading populations more often, and we may see that populations fixed for higher α genotypes are invaded less often.

To study angle substitutions within our simulation, it is helpful to use a slight variation of the model described above. This variation constrains vector angles to the integer degrees, such that α can have only the 91 states in the interval [0, 90]: this minor alteration places a minimum size on angle mutations, focusing our attention on significant changes in α. We used this model to determine the frequency of angle substitutions that lead to new values of α, and the duration for which populations are fixed for particular values of α.

One possible explanation for the pattern seen in Figure 2 is that angle mutants that increase α are more likely to fix in the population. Figure 3B shows the frequency of substitutions as a function of the substituting α, pooled for 12 replicate simulations. Clearly, there is a small bias in the frequency of substitutions, but this trend is similar for each μ_a value and opposes the dominant trend in Figure 3A: large α genotypes are favored in spite of a bias against substitutions that lead to their fixation.

Another possibility is that populations fixed for larger α values persist longer between substitutions. Figure 3C shows a set of histograms for the intervals between substitutions as a function of α. These data, taken from the same set of simulations as those in Figure 3B, clearly show that populations of large α genotypes are more prevalent because they persist for many generations before a new angle mutant invades the population. This differential persistence also varies with μ_a, accounting for the relationship between mean α and angle mutation rate seen in Figure 2.

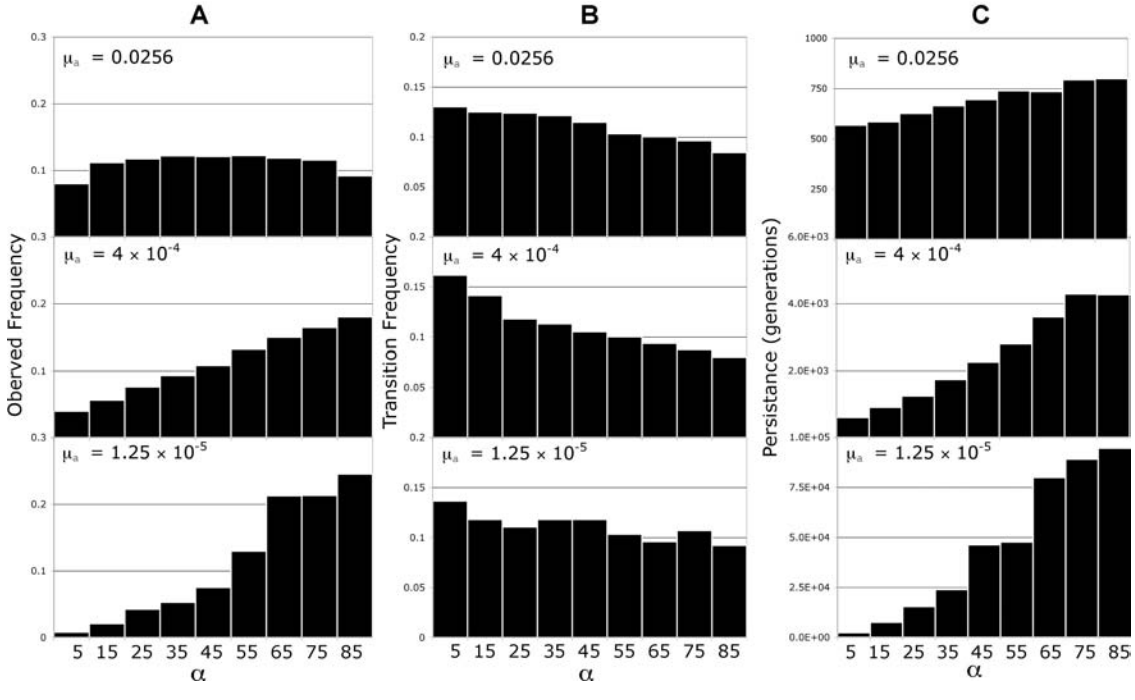

Figure 3. **(A) Frequency of mean population α, pooled across 24 simulations, for three values of μ_a. Other parameters are the same as in Figure 2. (B) Distribution of angle substations grouped by substituting α, pooled across 12 simulations of the variant model described in the text. (C) Distribution of times between angle substitutions grouped by α, pooled across the same 12 simulations.**

II. α AND THE RATE OF EVOLUTION

The results in Figures 2 and 3 suggest that evolution in a randomly changing environment elicits some difference between genotypes based on α, that this difference does not depend strongly on the presence or absence of recombination, and that it is amplified as μ_a decreases. Because α measures the shape of the distribution of mutational effects on the phenotype, we expected that α might capture an aspect of variability that was important for evolution. We hypothesized that the key trait that varies with α is evolvability: specifically, genotypes with higher values of α are able to produce vector magnitude mutations that are adaptive in a greater range of circumstances. Below we develop the basis for this hypothesis and test it with further simulations.

First, it is essential to note that populations in fluctuating environments may experience a mix of directional and stabilizing selection: if periods of stabilizing selection shape the mean value of α, then our hypothesis linking evolvability to the prevalence of high values of α is incomplete. We studied populations subject only to stabilizing selection by initiating simulations with populations at the optimum phenotype, and evolving these populations without shifts in the environment. After examining several combinations of mutation rates in both the presence and absence of recombination, we found no evolution of α in these populations. A subset of these data is graphed in Figure 2. This finding was consistent with our observations that, across all simulations, change in α was always associated with directional adaptation. These observations can be understood by noting that any individual mutation changing α alters the phenotype, and that, at the optimum, any alteration of the phenotype is deleterious. Therefore, stabilizing selection opposes any evolution of α. These results support our focus on the relationship between α and the response to directional selection.

Figure 4 provides a quantitative depiction of the relationship between α and evolutionary constraint on the response to directional selection. At the center of each plot is the phenotype of a model organism: this phenotype is normalized to (0,0) for convenience. We simulated evolution of a clonal population in 360 novel environments, each an equally spaced point on the unit circle depicted in Figure 4. Evolution in each environment revealed the rate of adaptation as a function of the bearing of the new environment to the organisms' axes of variation. Each simulation was replicated 100 times, and model parameters were as above, except that organisms were asexual and only magnitude mutations were permitted. The inner curve in each part of Figure 4 shows the mean position of the population after 20 generations of evolution: (a) $\alpha = 5^\circ$, (b) $\alpha = 45^\circ$, (c) $\alpha = 85^\circ$. Using only magnitude mutations, genotypes with high α can rapidly adapt to any novel environments, whereas genotypes with smaller α values are constrained. When μ_a is small relative to μ, we expect the adaptive utility of magnitude mutations to strongly affect the rate of adaptation: Figure 4 suggests that populations with larger values of α may adapt more quickly in these circumstances.

If α is associated with evolvability, how can this connection explain the preponderance of high α values? We consider two mechanisms relating evolvability to evolutionary success. First, genotypes with higher α values may fix because they are more likely to generate beneficial mutations: this mechanism is often called "indirect selection" in work on mutator alleles and recombination modifiers. Also, populations fixed for higher α values may, because of their greater evolvability, be less likely to be invaded by genotypes with mutant values of α, whether these mutants have a direct benefit or are linked with beneficial magnitude mutations. Although the same factor—the rate of beneficial mutation—drives both of these mechanisms, they are empirically distinguishable: the first mechanism describes competition among rare angle mutants, whereas the second applies to competition between rare angle mutants and a resident population.

We took several approaches to determine if invasion probability and resistance to invasion varied with α. First, we sought to determine if it was possible for selection on evolvability to

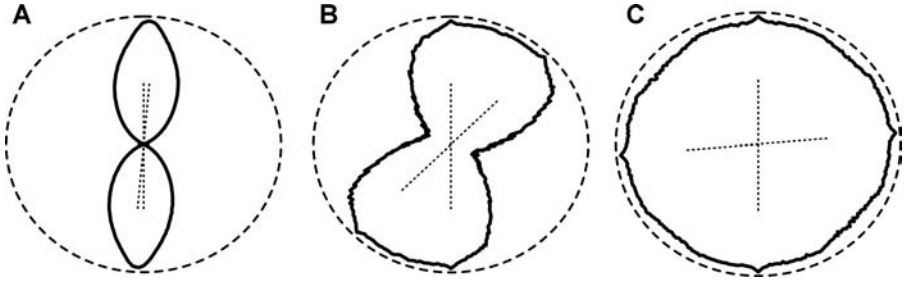

Figure 4. Short-term evolutionary response of populations with prescribed α values to a spectrum of new optima. Each point on the dashed circles represents a new optimum at a unit distance from the phenotype of the initial genotype; the solid line is a radial plot of the mean location of the population after 20 generations of evolution, as a function of the bearing of the new phenotypic optimum. Populations are 600 initially identical individuals; trait angles as indicated by the dotted lines: (A) $\alpha = 5^\circ$, (B) $\alpha = 45^\circ$, (C) $\alpha = 85^\circ$. Evolution occurs normally, except that angle mutations are not permitted; each solid line is the mean of 100 replicates for each of 360 new environments.

Table 1. **Invasion counts for neutral mutants in novel environments. Each entry records the number of replicates in one million trials in which the invader fixed. Invading mutants began at an initial frequency of 1% in a population of 100. Column sums reflect the likelihood of that resident type to be displaced, and so do not include the italicized counts on the main diagonal.**

		Resident						
		0°	15°	30°	45°	60°	75°	90°
Invader	0	*41,759*	32,854	23,516	17,081	12,447	9,977	8,820
	15	48,978	*38,500*	26,705	18,349	13,746	10,287	9,435
	30	54,647	46,271	*31,802*	21,285	15,204	11,654	10,451
	45	58,455	51,268	36,462	*24,130*	17,051	12,779	11,409
	60	60,318	53,726	39,282	26,432	*17,942*	13,179	11,869
	75	60,738	54,891	41,174	27,890	18,862	*13,661*	12,305
	90	60,976	55,544	41,715	28,217	18,987	13,939	*12,340*
	Sums − diagonal	344,112	294,554	208,854	139,254	96,297	71,815	64,289

produce a biased distribution of α values consistent with those in Figure 3. To do this we created model populations consisting of 99 copies of a resident genotype and a single copy of a mutant genotype. The resident and mutant genotypes had specific α values, but indistinguishable fitnesses in the chosen environment. Such populations were then allowed to adapt, with only magnitude mutations permitted, until either the resident or the invader was extinct. The number of times, out of 1 million trials, that the invader was successfully fixed is reported in Table 1 for simulations with a range of α values. Appendix 2 discusses why the results given in Table 1 deviate from the neutral null expectation. Table 1 provides evidence in support of both mechanisms: invasion success is directly associated with the invader's α value, and invasion frequency declines as the resident's α increases.

To confirm that these results were consistent with the histograms in Figure 3, we used a simple Markov model to translate our invasion probabilities into a frequency distribution. Permitting only the seven α states in Table 1, and assuming that (neutral) mutants of each type occurred with equal frequency, we constructed a transition matrix using the probabilities in Table 1. Using Mathematica, we determined the eigenvector corresponding to the dominant eigenvalue of the matrix, which when normalized is the stable evolutionary distribution of α values. This distribution, plotted in Figure 5, is qualitatively similar to the histograms in Figure 3.

These results indicate that both mechanisms acting together can produce a pattern biased in favor of high α values, but they do not quantify the significance of either influence in our evolving populations. To prove unequivocally that α determines the invasion probability of a mutant, we examined the invasion frequencies of pairs of mutants in evolving populations with identical parameters to those simulations used to make Figures 2 and 3. Populations were evolved as before, except that the environment was altered every 500 generations, and each generation was monitored for the occurrence of mutations that changed α. When more than one angle mutation occurred in a single generation, the first two such mutants were screened to see if both mutants had a reasonable chance to fix. Based on preliminary simulations, we established that angle mutations occurring after the first 100 generations of each cycle, and that had a fitness of less than 10% of the mean, were extremely unlikely to fix: such mutants were excluded. If both mutants in a pair passed this test, we followed their evolution, in the context of the population in which they occurred, until both mutants were extinct or until either fixed. Each pair was tracked for 100,000 replicate trials, and pairs for which each mutant fixed at least once were further analyzed. For each relevant pair, we recorded each mutant's α and fitness in the current environment. After analyzing our initial set of data, we established that mutant pairs that occurred near the transition to

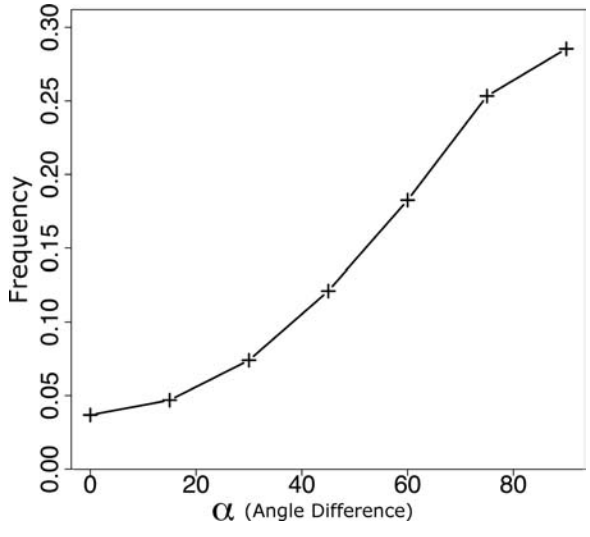

Figure 5. **Frequency of α variants at equilibrium in a discrete Markov model based on the invasion counts is given in Table 1. Each iteration, one of the six possible mutants to a different α, is modeled to occur with equal probability, and the system switches to the mutant state with a probability equal to the corresponding invasion frequency derived from Table 1.**

a new environment were uninformative, because these mutants would evolve in a different, unrecorded fitness landscape for each of their 100,000 replicates: we therefore discarded all pairs that occurred within 100 generations of the transition. After numerous simulations using independent starting populations, we obtained 191 informative pairs for the free-recombination model, and 160 such pairs for the no-recombination model.

This procedure has a significant advantage over examining the fixation probability of individual mutants: because both members of a pair shared the same environment and population of competing alleles, the only factors that determined their relative success were their fitnesses and their propensities to acquire beneficial mutations. Due to our design, the only consistent influence on the latter propensity was α. To determine if α had an effect on invasion success in a large, diverse population, we performed a multiple linear regression of the difference in α between each member of a pair, the natural log of their fitness ratio, and the interaction of these terms, on the natural log of the ratio of fixation odds. The results of these analyses for both the absence and presence of recombination are given in Table 2. In both, there is clear evidence for an effect of α on fixation probability, showing that dependence of invasion probability on α does partially explain the evolution of high values of α. Also, the model coefficients for both significant components, α and the log fitness ratio, are similar between the two cases, suggesting that the effect of invasion probability does not depend on recombination. Using these linear coefficients, we can compare the effects of angle differences and fitness differences on the outcome of competitions between mutants: a difference in α of 10° is equivalent to a fitness difference of about 5.5% in the asexual populations, and an estimated 7% difference in fitness in the sexual populations.

Table 2. Results of a multiple linear regression of difference in α (*Angle*) and the natural logarithm of the ratio of fitnesses (*Ln_Fit*) on the odds ratio of fixation. In both the free recombination and no recombination models, *Angle* and *Ln_Fit* are significant, whereas their interaction is not.

	Estimate	SE	t value	Pr($>\|t\|$)
No recombination				
Intercept	−0.051554	0.255942	−0.201	0.84062
Angle	0.026634	0.008368	3.183	0.00176
Ln_Fit	5.011443	0.351722	14.248	$<2\times10^{-16}$
Angle: Ln_Fit	0.010889	0.011639	0.936	0.35094
Free recombination				
Intercept	−0.339427	0.276155	−1.229	0.22057
Angle	0.029028	0.009066	3.202	0.00160
Ln_Fit	4.311185	0.438099	9.841	$<2\times10^{-16}$
Angle:Ln_Fit	−0.001450	0.013968	−0.104	0.91746

No Recombination: Residual standard error: 3.157 on 156 degrees of freedom.
Multiple R^2: 0.5696, Adjusted R^2: 0.5614.
***F*-statistic: 68.83 on 3 and 156 df, *P*-value: $<2.2\times10^{-16}$.**
Free recombination: Residual standard error: 3.735 on 187 degrees of freedom.
Multiple R^2: 0.352, Adjusted R^2: 0.3416.
***F*-statistic: 32.87 on 3 and 187 df, *P*-value: $<2.2\times10^{-16}$.**

Although this regression analysis shows that α has a detectable effect on invasion probability, it also demonstrates that angle mutations may fix because they have strong and directly beneficial effects on the phenotype. This suggests that the resistance of a rapidly evolving population to invasion by beneficial mutants, particularly beneficial angle mutations, may also be significant. To better understand how evolvability is affected by the dominant α of the population, and how evolvability affects the invasion success of angle mutants, we performed additional simulations with another variant of the model.

This variant allowed populations to adapt normally, with 600 individuals, a μ_a of 0.0016, and a μ of 0.04: however, at each environmental change, the dominant value of α was identified, and the subpopulation with that dominant value was sampled, with replacement, to form a test population. This test population contained much of the variation in vector magnitudes of the original population, but was monomorphic for α. The test population was then allowed to adapt to a new environment, using only magnitude mutations, for 40 generations. Figures 6A and C show mean population distance from the new optimum: (a) depicts no recombination, (b) shows recombination, and open circles represent populations with an α of over 80°; crosses, an α between 40° and 50°; filled squares, an α of less than 10°. Number of replicates vary, but are at least 60 for each line; error bars represent standard errors. Parts (B) and (D) of Figure 6 show the vulnerability to invasion by angle mutants of the same populations. Vulnerability to angle mutants was measured by introducing angle mutations in copies of the test population, then tracking those mutants until they were fixed or were lost. The fraction that fixed is reported as the vulnerability to invasion.

Two patterns are clear from Figure 6: first, as in Figure 4, populations with a high value of α adapt much more quickly using magnitude mutations. Second, rate of adaptation is correlated with vulnerability to invasion by angle mutations. This is consistent with the observation that angle mutations often have dramatic phenotypic effects. When the population is far from the optimum, such large mutations may be strongly favored. However, as the population approaches the optimum, the large phenotypic effects of angle mutations make them increasingly likely to be deleterious. This relationship between distance and vulnerability to substitution provides a mechanism to explain why populations with higher α are more resistant to invasion by new α mutants, the pattern observed in Figure 3C.

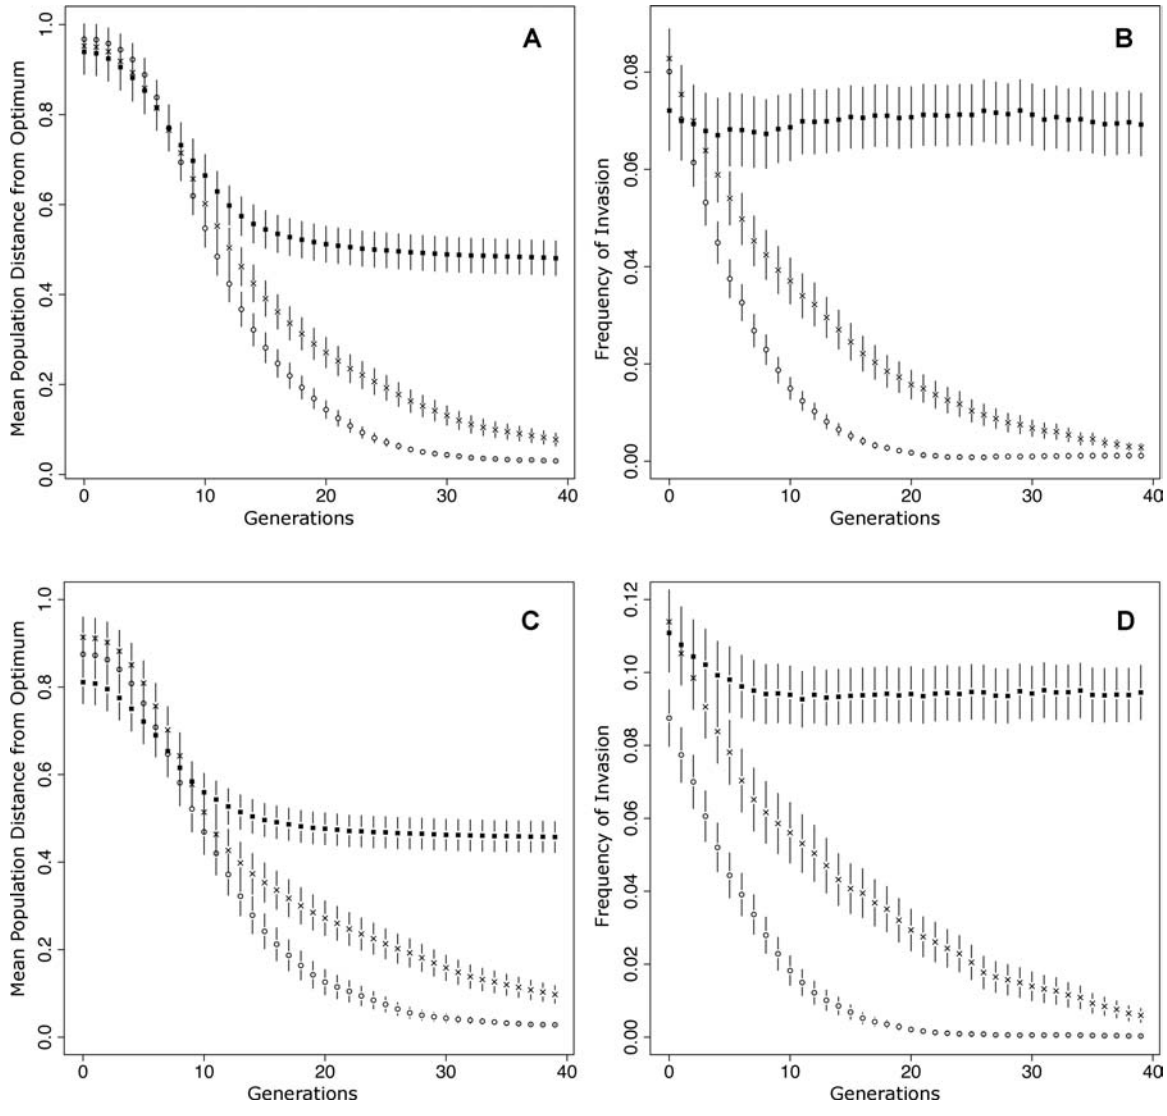

Figure 6. **Evolution and vulnerability to invasion of asexual and sexual populations adapting with only magnitude mutations. Open circles represent populations with an α of over 80°; crosses, an α between 40° and 50°; filled squares, an α of less than 10°. (A) and (B) are asexual populations; (C) and (D) are sexual. Error bars are standard errors.**

III. VARIABILITY AND NONRANDOM ENVIRONMENTAL CHANGE

One natural extension of our model is to nonrandom patterns of environmental change. Because the phenotypic optimum is characterized by two coordinates, we can vary the optimum by changing only one coordinate at a time. Because this pattern of environmental change separates the two components of the phenotype, we will refer to it as modular environmental change. This modular pattern constrains the new optimum to a slice of the unit disc that is parallel to one of the phenotypic plane axes. Although the changed coordinate is chosen randomly, and the new optimum is drawn uniformly from the slice of the unit disc, this scheme may favor genotypes whose axes of variation align with the axes of phenotypic space. Consequently, evolution may favor genotypes with orthogonal vectors that are also aligned with the phenotypic axes.

Figure 7 displays the frequency with which genotypic vectors occur in evolutionary simulations. Individuals are sampled periodically and each vector of their genotypes is measured with respect to the phenotypic axes drawn in Figure 1. Figure 7 contains data from 72 independent simulations, each of 1 million generations of evolution for $N = 600$, $\mu = 0.04$, $\mu_a = 0.0004$, and with recombination. Environmental fluctuations occur every 1000 generations but change only one coordinate of the optimum phenotype at a time.

When the phenotypic optimum is randomly redrawn from the unit disc, each vector displays no bias toward any angle (data not shown). In contrast, Figure 7 demonstrates that, when the

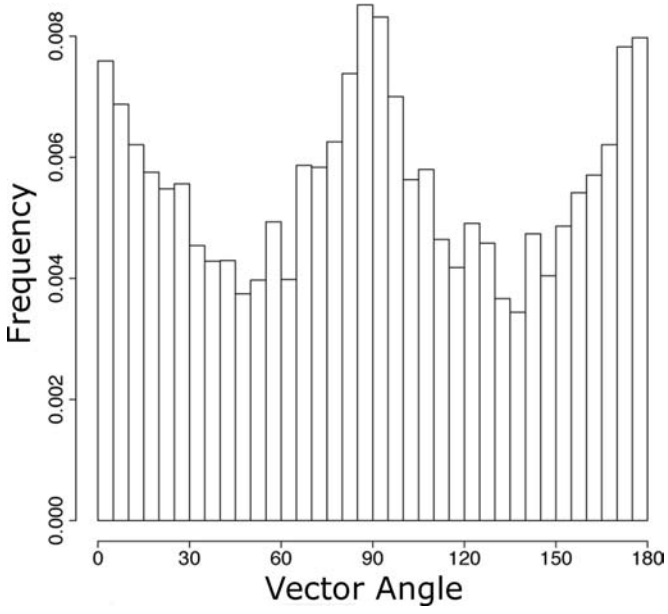

Figure 7. Histogram of vector angles for evolutionary simulations in which environmental change is constrained. Both vectors are considered together for 72 replicate simulations of one million generations of evolution for $N = 600$, $\mu = 0.04$, $\mu_a = 0.004$, and with recombination. As before, a new phenotypic optimum was chosen every 1000 generations: here, only one coordinate of the optimum was changed at a time.

optimum moves along the phenotypic axes, vectors aligned with those axes are favored. Like the bias toward larger values of α discussed above, this results is explained by selection on evolvability. Genotypes with vectors aligned with the axes of phenotypic space are disposed to produce some positively selected variants, and so are overrepresented in evolving populations.

Because evolution in this scenario favors genotypes with vectors aligned with the orthogonal axes of phenotypic space, we expect these evolved genotypes to have especially large values of α. This prediction is also confirmed: the median α in populations subjected to modular environmental change is 66.78° ± 0.36° with recombination and 63.55° ± 0.52° without recombination, whereas the median α for comparable control simulations in which the optimum changes randomly is 58.65° ± 0.37° with recombination and 59.65° ± 0.61° without.

Discussion

EVOLUTION OF EVOLVABILITY IN OUR MODEL

Our results show that indirect selection for α shapes the evolution of this measure of variability, and that α is strongly correlated with evolvability. Because α describes the shape of the spectrum of phenotypic mutations, our results suggest that selection on evolvability favors an isotropic, or unconstrained, pattern of variability. Although we have strong evidence that α influences the invasion probability of mutants, several facts argue that resistance to invasion is also a significant factor in the evolution of α. Figure 3B shows that transitions to low α values are in fact slightly more common in evolution: if α had a strong effect on invasion probabilities, we would expect the opposite of this pattern. Additionally, Figure 2 shows that higher median α values are found in populations with low rates of angle mutations. The influence of invasion probability on the distribution of α is expected to depend on the level of variation in α within the population: α variants must exist for indirect selection to select among them. Because the vast majority of angle mutations are observed to be nonneutral, angle mutation rate strongly determines the standing variation in α. If differing invasion probabilities were entirely responsible for the observed preference for high α values, we would expect this preference to increase with angle mutation rate: instead, we again observe the opposite.

The results shown in Figure 6 suggest that the missing piece of the puzzle is the resistance of a population to invasion by new mutants, particularly beneficial angle mutants. This hypothesis accommodates the distributions of transitions in Figure 3B and explains the distributions of persistence in Figure 3C. It is also consistent with the results of Figure 2: Figure 6 shows that vulnerability to invasion only decreases after populations have had a chance to adapt through the selection of magnitude mutations. If angle mutation rates are high, angle mutants can invade before the differences in evolvability between populations with low or high α are apparent: high angle mutation rates, therefore, are expected to flatten the differences between populations with different α values. Finally, the results of our mutation-invasion simulations, presented in Table 1 and Figure 5, imply that the resistance of high-α populations to invasion by mutants is a major contribution to the distribution of α over evolutionary time.

Our results suggest that two forms of selection on evolvability shape the distributions of α values seen in Figures 2 and 3. Both forms of selection, differential invasion probability and differential resistance to invasion, derive from the same property: genotypes with high α values can create a greater range of possible mutations, and are therefore more likely to be able to produce a positively selected variant in any new environment. The difference between these mechanisms is in the context: invasion probability affects the fitness of rare mutants, whereas resistance to invasion reflects the success of numerically dominant genotypes.

NATURAL SELECTION AND THE EVOLUTION OF EVOLVABILITY

The results of our model have several intriguing implications for the evolution of evolvability in general. First, both mechanisms of selection of evolvability escape several of the problems evident in models of mutation and recombination rate evolution. High α genotypes produce more useful variation without a concomitant

increase in deleterious variation, in contrast to alleles that increase mutation and recombination rates. The strong phenotypic effects of vector angles also ensure that loci conferring greater evolvability cannot be divorced from the beneficial variation they create: this linkage between the phenotypic and variational effects of a mutation explains why recombination does not prevent the evolution of evolvability. This tolerance to recombination is simply a consequence of the epistatic interactions among loci in our model: it may therefore be a general feature of systems that evolve evolvability through changes in epistasis.

Our results also show how evolvability may persist in a changing genome. Magnitude mutations may be more successful because of a large α, but their fixation does not alter α: the determinant of evolvability is therefore passed down largely intact between generations, even in a rapidly evolving lineage. This avoids the potential problem of the inheritance of evolvability noted above. These observations suggest that the evolution of evolvability through changes in developmental systems may occur more readily than inferred from the study of mutator and recombination modifier alleles.

Our model contained two distinct levels of mutational change: frequent magnitude mutations of moderate phenotypic effect, and rare angle mutations of larger effect. Adaptation links evolution in both of these levels: if adaptive change through magnitude mutations is facile, then selection favors stasis in vector angles. If, on the other hand, adaptation via magnitude mutations is frustrated, then change in vector angles, and therefore change in pleiotropy, is strongly favored. Parts of the genotype, by affecting the variability of other traits, can therefore affect their own susceptibility to change. Any trait that facilitates adaptive change in another aspect of the genotype is consequently insulated from the need to change and is under stabilizing selection: this conservation may be a very general mechanism of selection on evolvability.

EVOLUTION OF VARIABILITY AND MODULARITY

Our results indicate that a specific pattern of variability can evolve in response to random environmental change. This pattern, namely orthogonality of the axes of mutational variation or isotropy of the mutant distribution, is reminiscent of the evolutionary concept of modularity. Modularity in evolutionary biology refers to the organization of variability into largely independent units: several authors have suggested that some level of modular organization will maximize the evolvability of complex traits and organisms (Lewontin 1978; Wagner and Altenberg 1996; Hansen 2003; Welch and Waxman 2003; Kashtan and Alon 2005).

Whether the need for evolvability can select for modularity is largely unknown (Wagner et al. 2005), although there is some evidence that, if selection pressures are partitioned into independently changing pieces, selection can favor modularity (Kashtan and Alon 2005). Our results support the idea that selection on evolvability can sort the genotype into largely independent units of variability. Although this orthogonal organization of variability can be called modularity, our randomly varying selection regime does not produce a connection between orthogonal phenotypic traits and corresponding units of variability. This congruence between phenotypic traits and variability only arises when environmental change alters the optimum of one trait at a time, as seen in Figure 7. This confirms the idea that selection can create modularity when selective pressures change in a decomposable way, and so stores information about the nature of environmental change (Wagner and Altenberg 1996; Kashtan and Alon 2005; Kastan et al. 2007).

COMPARISONS WITH QUANTITATIVE GENETICS MODELS

The idea that selection can shape patterns of mutational variability to match patterns of selection is also discussed in quantitative genetics (e.g., Lande 1979, 1980, 1982; Cheverud 1984; Jones et al. 2007). Our results confirm the prediction that a randomly shifting phenotypic optimum will select for a modular, or isotropic, pattern of mutational variation (Jones et al. 2007). However, several differences in emphasis and implementation distinguish our results from those derived from a quantitative genetics framework. Our work focuses on the mutational tendencies inherent in genotypes: these tendencies correspond to the mutational matrix, or **M**-matrix, in quantitative genetics. In contrast, most quantitative genetic studies have focused on the **G**-matrix, the matrix of additive genetic variance and covariance of traits, which is the population-level product of selection, history, and the **M**-matrix (e.g., Lande 1979, 1982; Turelli 1988; Wagner 1989; Schluter 1996; Steppan et al. 2002; Bégin and Roff 2003; Jones et al. 2003, 2004). Although some authors have addressed the evolution of mutational effects (e.g., Wagner et al. 1997; Hermisson et al. 2003; Hansen et al. 2006), so far only one study, Jones et al. (2007), has examined how **M**-matrix evolution affects evolvability.

Jones et al. (2007) address the evolution of evolvability by studying neutral modifiers of mutational patterns in a population under stabilizing selection. Our study departs from this approach by modeling populations with traits that individually affect both the phenotype and variability: our results are therefore not restricted to modifiers with no direct fitness effects. Also, we model populations adapting to a diverse and changing mixture of stabilizing and directional selection. We believe these conditions comprise the most realistic context for the evolution of evolvability, and are the most crucial to understand.

Our work also differs from quantitative genetics approaches by our focus on the dynamics affecting individual mutants, rather than on the phenotypic and genotypic distributions of an entire population. Our emphasis on the detailed dynamics of selection has produced results that directly address current controversies

over the effects of selection, recombination, and heritability on evolvability. We expect insights from quantitative genetics to be invaluable in generalizing our results and extending them to empirical data. Also, the strong influence of mutation on the stability of the **G**-matrix (Steppan et al. 2002; Jones et al. 2003, 2007) highlights how understanding the evolution of evolvability may be crucial to predicting the evolution of the **G**-matrix.

Conclusions

The results of analyzing this simple model suggest that evolvability can readily evolve without changes in mutation or recombination rate. Because evolvability depends on epistatic interactions between loci, it can be maintained by selection in recombining populations. We have shown that evolvability affects the probability of a new mutant to invade a population. This process is familiar from models of mutation rates, but here plays a supporting role to a novel dynamic: the resistance of highly evolvable populations to invasion by mutants with large effects. Both of these processes derive from the epistatic role of some loci in determining evolvability and the phenotype: when a locus facilitates rapid evolution in other areas of the genome, it is less likely to be substituted by an adaptive mutation itself. This sorting of genotypes with evolvability can reshape the pleiotropic relationships connecting the genotype to the phenotype and create an isotropic pattern of variability. This pattern of variability can also be aligned by selection with regularities in how selection pressures change: this alignment of genotypic and phenotypic variability with environmental variability could drive the evolution of modularity.

ACKNOWLEDGMENTS

The authors wish to thank P. Turner and members of the Turner lab for resources and helpful feedback, and J. True, J. Masel, and an anonymous reviewer for constructive and insightful reviews. JAD was supported by a NASA Graduate Students Researchers Program Fellowship.

LITERATURE CITED

Adami, C. 2006. Digital genetics: unraveling the genetic basis of evolution. Nat. Rev. Genet. 7:109–118.

Aharoni, A., L. Gaidukov, O. Khersonsky, S. M. Gould, C. Roodveldt, and D. S. Tawfik. 2005. The 'evolvability' of promiscuous protein functions. Nat. Genet. 37:73–76.

Alberch, P. 1991. From genes to phenotype—dynamic-systems and evolvability. Genetica 84:5–11.

Altenberg, L. 1994. The evolution of evolvability in genetic programming. Pp. 47–74 *in* J. K. E. Kinnear, ed. Advances in genetic programming. MIT Press, Cambridge, MA.

Ancel, L. W., and W. Fontana. 2000. Plasticity, evolvability, and modularity in RNA. J. Exp. Zool. 288:242–283.

Andre, J. B., and B. Godelle. 2006. The evolution of mutation rate in finite asexual populations. Genetics 172:611–626.

Baatz, M., and G. P. Wagner. 1997. Adaptive inertia caused by hidden pleiotropic effects. Theor. Popul. Biol. 51:49–66.

Bedau, M. A., and N. H. Packard. 2003. Evolution of evolvability via adaptation of mutation rates. Biosystems 69:143–162.

Begin, M., and D. A. Roff. 2003. The constancy of the G matrix through species divergence and the effects of quantitative genetic constraints on phenotypic evolution: a case study in crickets. Evolution 57:1107–1120.

Bell, G. 1982. The masterpiece of nature: the evolution and genetics of sexuality. Univ. of California Press, Berkeley, CA.

Bloom, J. D., S. T. Labthavikul, C. R. Otey, and F. H. Arnold. 2006. Protein stability promotes evolvability. Proc. Natl. Acad. Sci. USA 103:5869–5874.

Burch, C. L., and L. Chao. 2000. Evolvability of an RNA virus is determined by its mutational neighbourhood. Nature 406:625–628.

Cheverud, J. M. 1984. Quantitative genetics and developmental constraints on evolution by selection. J. Theor. Biol. 110:155–171.

Cowperthwaite, M., and L. A. Meyers. 2007. How mutational networks shape evolution: lessons from RNA models. Annu. Rev. Ecol. Syst. 38:203–230.

de Visser, J. A. G. M. 2002. The fate of microbial mutators. Microbiol-Sgm 148:1247–1252.

Earl, D. J., and M. W. Deem. 2004. Evolvability is a selectable trait. Proc. Natl. Acad. Sci. USA 101:11531–11536.

Gilchrist, G. W., and C. E. Lee. 2007. All stressed out and nowhere to go: does evolvability limit adaptation in invasive species? Genetica 129:127–132.

Goddard, M. R., H. Charles, J. Godfray, and A. Burt. 2005. Sex increases the efficacy of natural selection in experimental yeast populations. Nature 434:636–640.

Griswold, C. K. 2006. Pleiotropic mutation, modularity and evolvability. Evol. Dev. 8:81–93.

Hansen, T. F. 2003. Is modularity necessary for evolvability? Remarks on the relationship between pleiotropy and evolvability. Biosystems 69:83–94.

———. 2006. The evolution of genetic architecture. Annu. Rev. Ecol. Evol. S 37:123–157.

Hansen, T. F., J. M. Alvarez-Castro, A. J. R. Carter, J. Hermisson, and G. P. Wagner. 2006. Evolution of genetic architecture under directional selection. Evolution 60:1523–1536.

Hermisson, J., T. F. Hansen, and G. P. Wagner. 2003. Epistasis in polygenic traits and the evolution of genetic architecture under stabilizing selection. Am. Nat. 161:708–734.

Jones, A. G., S. J. Arnold, and R. Burger. 2003. Stability of the G-matrix in a population experiencing pleiotropic mutation, stabilizing selection, and genetic drift. Evolution 57:1747–1760.

———. 2004. Evolution and stability of the G-matrix on a landscape with a moving optimum. Evolution 58:1639–1654.

———. 2007. The mutation matrix and the evolution of evolvability. Evolution 61:727–745.

Kashtan, N., and U. Alon. 2005. Spontaneous evolution of modularity and network motifs. Proc. Natl. Acad. Sci. USA 102:13773–13778.

Kastan, N., E. Noor, and U. Alon. 2007. Varying environments can speed up evolution. Proc. Natl. Acad. Sci. USA 104:13711–13716.

Khersonsky, O., C. Roodveldt, and D. S. Tawfik. 2006. Enzyme promiscuity: evolutionary and mechanistic aspects. Curr. Opin. Chem. Biol. 10:498–508.

Kirschner, M., and J. Gerhart. 1998. Evolvability. Proc. Natl. Acad. Sci. USA 95:8420–8427.

Lande, R. 1979. Quantitative genetic-analysis of multivariate evolution, applied to brain-body size allometry. Evolution 33:402–416.

———. 1980. The genetic covariance between characters maintained by pleiotropic mutations. Genetics 94:203–215.

———. 1982. A quantitative genetic theory of life-history evolution. Ecology 63:607–615.

Lewontin, R. C. 1978. Adaptation. Sci. Am. 239:212–228.

Lynch, M. 2007. Colloquium papers: the frailty of adaptive hypotheses for the origins of organismal complexity. Proc. Natl. Acad. Sci. USA 104 (Suppl 1):8597–8604.

Magg, S., and A. Philippides. 2006. GasNets and CTRNNs—a comparison in terms of evolvability. Lect. Notes Comput. Sci. 4095:461–472.

Masel, J. 2005. Evolutionary capacitance may be favored by natural selection. Genetics 170:1359–1371.

Masel, J., and A. Bergman. 2003. The evolution of the evolvability properties of the yeast prion [PSI+]. Evolution 57:1498–1512.

Matsumoto, M., and Nihimura, T. 1998. Mersenne Twister: a 623-dimensionally equidistributed uniform pseudorandom number generator. ACM Trans. Model. Comput. Simul. 8:3–30.

Meyers, L. A., F. D. Ancel, and M. Lachmann. 2005. Evolution of genetic potential. PLos. Comput. Biol. 1:236–243.

O'Loughlin, T. L., W. M. Patrick, and I. Matsumura. 2006. Natural history as a predictor of protein evolvability. Protein Eng. Des. Sel. 19:439–442.

Orr, H. A. 2000. Adaptation and the cost of complexity. Evolution 54:13–20.

Otto, S. P., and N. H. Barton. 1997. The evolution of recombination: removing the limits to natural selection. Genetics 147:879–906.

Pepper, J. W. 2003. The evolution of evolvability in genetic linkage patterns. Biosystems 69:115–126.

Plotkin, J. B., and J. Dushoff. 2003. Codon bias and frequency-dependent selection on the hemagglutinin epitopes of influenza A virus. Proc. Natl. Acad. Sci. USA 100:7152–7157.

Poelwijk, F. J., D. J. Kiviet, D. M. Weinreich, and S. J. Tans. 2007. Empirical fitness landscapes reveal accessible evolutionary paths. Nature 445:383–386.

Poole, A. M., M. J. Phillips, and D. Penny. 2003. Prokaryote and eukaryote evolvability. Biosystems 69:163–185.

Quayle, A. P., and S. Bullock. 2006. Modelling the evolution of genetic regulatory networks. J. Theor. Biol. 238:737–753.

Radman, M., I. Matic, and F. Taddei. 1999. Evolution of evolvability. Ann. NY Acad. Sci. 870:146–155.

Schluter, D. 1996. Adaptive radiation along genetic lines of least resistance. Evolution 50:1766–1774.

Sniegowski, P. D., and H. A. Murphy. 2006. Evolvability. Curr. Biol. 16:R831–R834.

Sniegowski, P. D., P. J. Gerrish, T. Johnson, and A. Shaver. 2000. The evolution of mutation rates: separating causes from consequences. Bioessays 22:1057–1066.

Stearns, S. C. 2000. Daniel Bernoulli (1738): evolution and economics under risk. J. Biosci. 25:221–228.

Steppan, S. J., P. C. Phillips, and D. Houle. 2002. Comparative quantitative genetics: evolution of the G matrix. Trends Ecol. Evol. 17:320–327.

Tanay, A., A. Regev, and R. Shamir. 2005. Conservation and evolvability in regulatory networks: the evolution of ribosomal regulation in yeast. Proc. Natl. Acad. Sci. USA 102:7203–7208.

Tenaillon, O., F. Taddei, M. Radman, and I. Matic. 2001. Second-order selection in bacterial evolution: selection acting on mutation and recombination rates in the course of adaptation. Res. Microbiol. 152:11–16.

Tirosh, I., A. Weinberger, M. Carmi, and N. Barkai. 2006. A genetic signature of interspecies variations in gene expression. Nat. Genet. 38:830–834.

Turelli, M. 1988. Phenotypic evolution, constant covariances, and the maintenance of additive variance. Evolution 42:1342–1347.

Wagner, A. 1996. Does evolutionary plasticity evolve? Evolution 50:1008–1023.

Wagner, G. P. 1981. Feedback selection and the evolution of modifiers. Acta Biotheor. 30:79–102.

———. 1989. Multivariate mutation-selection balance with constrained pleiotropic effects. Genetics 122:223–234.

Wagner, G. P., and R. Burger. 1985. On the evolution of dominance modifiers.2. A nonequilibrium approach to the evolution of genetic systems. J. Theor. Biol. 113:475–500.

Wagner, G. P., and L. Altenberg. 1996. Perspective: complex adaptations and the evolution of evolvability. Evolution 50:967–976.

Wagner, G. P., G. Booth, and H. C. Bagheri. 1997. A population genetic theory of canalization. Evolution 51:329–347.

Wagner, G. P., C. H. Chiu, and T. F. Hansen. 1999. Is Hsp90 a regulator of evolvability? J. Exp. Zool. 285:116–118.

Wagner, G. P., Mezey, J., and Calabretta, R. 2005. Natural selection and the origin of modules. Pp. 33–50 *in* W. Callebaut, and Rasskin-Gutman, D., ed. Modularity: understanding the development and evolution of natural complex systems. MIT Press, Cambridge, MA.

Welch, J. J., and D. Waxman. 2003. Modularity and the cost of complexity. Evolution 57:1723–1734.

Associate Editor: J. True

Appendix 1

EXPLORATION OF MODEL PARAMETERS

This section briefly examines how the evolution of evolvability depends on the parameters and functions used to construct our model. Below we consider how changes in magnitude mutation rate, population size, frequency of environmental change, and the fitness function alter the main results plotted in Figure 2.

Online Supplementary Table S1 displays the mean and standard errors of the median α value for sets of replicate populations. Each set contains 24 replicate simulations of one million generations each, with 600 individuals, environmental shifts every 1000 generations, and recombination and mutation rates as specified in the table. These data demonstrate that the evolution of evolvability is not dependent on the somewhat high rate of magnitude mutation, 0.04 per locus per generation, used in the bulk of our simulations. Results are roughly constant for a given ratio of μ:μ_a, although the absolute magnitudes of these rates clearly influence the strength of selection on evolvability. These results are similar with and without recombination.

Online Supplementary Table S2 shows analogous results for a range of population sizes and periods between environmental shifts. In these simulations $\mu = 0.04$ and $\mu_a = 0.0004$. These results show that increasing the population size may slightly reduce selection for high values of α, whereas changing the frequency of environmental shifts has no significant effect. King and Masel (2007) predict that the product of population size and the frequency of environmental change is an important predictor of the efficacy of selection on traits like evolvability, particularly when that product is near one. In contrast, in our model changes to the period between shifts seem to have little effect. This difference is explained by the cost to evolvability in King and Masel (2007): in our model, α is maintained by strong stabilizing selection in well-adapted populations, and no cost to evolvability is apparent.

The data in Online Supplementary Table S2 also suggest a small decrease in mean α with increasing population size, again in contradiction to the predictions of King and Masel (2007). We suggest that the dependence of α derives from the disparity between the rates of magnitude mutation and angle mutation. In Part II of the Results we show the mean α is primarily the result of competition for fixation between magnitude and angle mutations: rapid adaptation through magnitude mutations prevents beneficial angle mutants from arising and fixing. As population size increases, the rate at which potentially beneficial angle mutants occur will clearly increase: if the rate of adaptation through magnitude mutations fails to increase at the same rate, however, then mean α will decrease. To test this, we measured the rate of adaptation in clonal populations with $\alpha = 45^\circ$, no recombination, $\mu = 0.04$, and $\mu_a = 0$. These populations evolved in new environments for 20 generations, as in the simulations depicted in Figure 4. With $N = 200$, the mean distance from the optimum after 20 generations was 0.767 ± 0.001, whereas with $N = 2000$, it was 0.657 ± 0.001. Figure 6 demonstrates that distance from the optimum is a strong indicator of the probability of angle mutants to invade: clearly, vulnerability to invasion does not decrease linearly with increasing population size. These data illustrate how increasing population size favors the substitution of angle mutants, decreasing α.

Online Supplementary Figure S2 shows how variation in the fitness function changes the outcome of simulations in which $N = 600$, environmental shifts occur every 1000 generations, $\mu = 0.04$ and $\mu_a = 0.0004$. These results measure the α of a randomly chosen individual, not the population average, because these two measures may differ for some fitness functions. Online Supplementary Figure S2 demonstrates that the evolution of evolvability occurs with negative exponential fitness functions over a range of slopes, but that more gently sloping functions enhance selection for high α. We suggest that very steep fitness functions, such as those in parts (E) and (F) in Online Supplementary Figure S2, may favor large-effect angle mutations over small-effect magnitude mutations, and may therefore level differences in evolvability. The fairly small population size of our simulations will preclude weakly beneficial mutants from fixing: Online Supplementary Figure S2 shows that most magnitude mutations will have a negligible effect on fitness in poorly adapted populations with steep fitness functions. Because large-effect angle mutations can access the steeply increasing area of the fitness landscape, they are more likely to fix in such populations. As we have established above, circumstances that favor the fixation of angle mutations will reduce the efficacy of selection on evolvability.

LITERATURE CITED

King, O., and J. Masel. 2007. The evolution of bet-hedging adaptations to rare scenarios (unpubl. ms.).

Appendix 2

UNDERSTANDING THE INVASION SIMULATION RESULTS

Figure 5, Table 1, and Part II of the Results section reference a set of simulations in which a mutant genotype competes against a clonal, resident population in a novel environment. In these simulations, both the mutant and the resident genotypes have equal fitness, but differences in the α value of each genotype may predispose one type to acquire beneficial mutations and so displace the other. When the mutant and resident also share the same value of α, we might expect the mutant to follow the dynamics of a neutral allele: we would therefore predict that the mutant would invade in 1 of N cases, or about 10,000 times in the one million replicates performed for each scenario. This expectation is contradicted by the data in Table 1.

The faulty assumption behind this expectation of neutral behavior lies in averaging evolvability. Although two genotypes with the same α have the same evolvability on average, in any specific environment one will have an advantage over the other. This advantage derives from the orientation of a genotype's vectors relative to the direction of steepest increase in fitness: whichever genotype is best aligned with this fitness gradient will have access to superior adaptive mutations. When this chance difference in evolvability favors the resident genotype, the mutant's probability of invasion will be between 0 and $1/N$: when the reverse occurs, the mutant's chance of success will be between $1/N$ and 1. The average probability over these two scenarios may therefore exceed $1/N$ if differences in evolvability strongly influence invasion success.

To measure the effects of chance differences in evolvability on fixation, we focus on simulations in which both types have $\alpha = 0$. In this scenario, we can measure the difference between the solitary vector angle of a genotype and the angle of quickest increase in fitness in a given environment: we refer to this difference as the bearing on the optimum, and a value of 0 indicates that the variational axis is directly aligned with the fitness gradient. Online Supplementary Figure S1 parses the outcomes of simulation replicates for this scenario according to the difference in bearings between the invader and resident: positive values of this difference indicate a chance difference in evolvability favoring the invading mutant.

As the gray bars in Online Supplementary Figure S1 show, there is symmetrical variation in the difference in bearing on the optimum between the mutant and resident genotypes. The striped bars show the fraction of successful invasions in each class of relative bearings. When the difference in evolvability favors the resident, the mutant rarely invades; when the difference favors the mutant, it often invades. These data suggest that pairs of equally fit alleles that differ, even idiosyncratically, in evolvability will behave differently than expected from neutral theory, and will substitute one another more often than commonly predicted.

Supplementary Material

The following supplementary material is available for this article:

Figure S1. Measured distributions of the relative bearings of genotype pairs for the simulations described in Table 1 and Results Part II.

Figure S2. Effects of alternative fitness functions on the evolution of evolvability.

Table S1. Means and standard errors of sets of median values of α.

Table S2. Means and standard errors of sets of median values of α.

This material is available as part of the online article from:
http://www.blackwell-synergy.com/doi/abs/10.1111/j.1558-5646.2007.00303.x
(This link will take you to the article abstract).

Please note: Blackwell Publishing is not responsible for the content or functionality of any supplementary materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.