

The effect of gene interactions on the long-term response to selection

Tiago Paixão^{a,1} and Nicholas H. Barton^a

^aInstitute of Science and Technology Austria, A3400 Klosterneuburg, Austria

Edited by Joseph Felsenstein, University of Washington, Seattle, WA, and approved March 7, 2016 (received for review September 22, 2015)

The role of gene interactions in the evolutionary process has long been controversial. Although some argue that they are not of importance, because most variation is additive, others claim that their effect in the long term can be substantial. Here, we focus on the long-term effects of genetic interactions under directional selection assuming no mutation or dominance, and that epistasis is symmetrical overall. We ask by how much the mean of a complex trait can be increased by selection and analyze two extreme regimes, in which either drift or selection dominate the dynamics of allele frequencies. In both scenarios, epistatic interactions affect the long-term response to selection by modulating the additive genetic variance. When drift dominates, we extend Robertson's [Robertson A (1960) *Proc R Soc Lond B Biol Sci* 153(951):234–249] argument to show that, for any form of epistasis, the total response of a haploid population is proportional to the initial total genotypic variance. In contrast, the total response of a diploid population is increased by epistasis, for a given initial genotypic variance. When selection dominates, we show that the total selection response can only be increased by epistasis when some initially deleterious alleles become favored as the genetic background changes. We find a simple approximation for this effect and show that, in this regime, it is the structure of the genotype–phenotype map that matters and not the variance components of the population.

epistasis | genetic variance | adaptation | genetic drift

The relation between an organism's genotype and its phenotype is immensely complicated, yet quantitative genetics predicts the correlations between relatives, and the response to selection over tens of generations, based primarily on an additive model. How do interactions between genes affect the response to selection? This question is not easy to answer: Although the additive model can be represented by a few parameters, there are an enormous number of possible relationships between genotype and phenotype. Some insight may come from studying specific well-understood systems, for example, regulation of gene expression by binding of transcription factors, or folding of RNA molecules. Here, we take the alternative approach, seeking statistical regularities, while making minimal assumptions about the nature of epistasis.

There has been a long-standing debate about the effect of genetic interactions on adaptation (1–7). Although some claim that they are unimportant because they contribute very little to the total genetic variance of a population, and consequently to its short-term response, others claim that their long-term effects can be substantial. At the heart of the debate is the distinction between what has been called physiological and statistical epistasis. The former is independent of the state of the population, whereas the latter is the statistical contribution of gene interactions to the trait variance, which depends on the allele frequencies (8, 9).

In the absence of new mutations, the total response to selection is limited by the initial standing variation. Selection acts on the variation present in the population, and will typically act to reduce it; the extent to which it does so depends on the allele frequencies, and on the relation between genotype and phenotype (i.e., the “genetic architecture”). Interactions between genes (epistasis) are a key component of this genetic architecture, and their effect on the response to selection has been controversial.

On the one hand, artificially selected populations are usually well approximated by the infinitesimal model, which, in its simplest form, assumes infinitely many loci of small and additive effect (10, 11). This suggests that gene interactions do not play an important role. On the other hand, biology can hardly be additive, and genomes are finite. If the trait depended on a small number of additive alleles, these would quickly fix—contradicting the robust observation of sustained response to selection. One possibility is that epistasis sustains additive genetic variance for longer: Alleles that were initially deleterious or near-neutral may acquire favorable effects as the genetic background changes, “converting” epistatic variance into additive, and so prolonging the response to selection.

Note that, if the map between genotype and phenotype were arbitrary, no predictions could be made: The fittest genotype could have any value, and genotypes can be organized in any way over the fitness landscape. We will make a statistical argument, which assumes that epistasis primarily involves low-order interactions, and that these can be treated as, to some degree, independent of each other. In particular, we assume that there is no systematic tendency for alleles with a positive marginal effect to interact positively (or negatively), on average; if there were, the rate of evolution would clearly accelerate (or decelerate).

Although much effort has been spent measuring the sign of epistasis among new mutations (12, 13), or among segregating variants (14–16), it is not clear that this would tell us about the natural pattern over a wider range. A population under selection can reach trait values far beyond the initial trait variance it displays, so the epistasis among variants in the existing population may not predict longer-term evolution. Here, we analyze two extreme limits: the limit of very weak selection, in which drift dominates the dynamics of the genetic variances, and the deterministic limit, where selection is the only force. We investigate

Significance

The role of gene interactions in the response to selection has long been a controversial subject; whereas some have claimed they are not relevant for adaptation, others have argued that their long-term effects are of high significance. In this manuscript, we derive simple and general predictions for the effect of gene interactions on the long-term response to selection in two extreme regimes. We show that, when the dynamics of allele frequencies are dominated by genetic drift, the long-term response is surprisingly simple, depending only on the initial components of the trait variance, regardless of the detailed genetic architecture. In the opposite regime, when selection dominates the dynamics of allele frequencies, the long-term response depends only on the genotype–phenotype.

Author contributions: T.P. and N.H.B. designed research, performed research, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. Email: tiago.paixao@ist.ac.at.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1518830113/-DCSupplemental.

the effect of epistasis on the long-term response of a population, and the mechanistic basis for these effects.

Results

The Interaction Between Epistasis and Random Drift. Robertson (17) showed that, in the weak selection limit ($N_e s \ll 1$), where N_e is the effective population size and s is the selection coefficient at a single locus, the total response of an additive trait in a finite diploid population is simply proportional to the additive variance initially present, provided there is no dominance. He used the fact that the probability of fixation of an allele can be written as $u(p_0) = (1 - \exp(-4N_e s p_0)) / (1 - \exp(-4N_e s)) = p_0 + 2p_0(1 - p_0)N_e s + O((N_e s)^2)$. The total expected response, summed over all generations until every allele is lost or fixed, is $R^\infty = \sum_i \alpha_i (u(p_0) - p_0)$, where α_i is the additive effect of each allele on the trait. When selection is weak relative to random drift, this can be approximated as $R^\infty \approx 2 \sum_i \alpha_i p_0 (1 - p_0) N_e s_i$. If the selection gradient on the trait is β , the selection favoring an allele with effect α_i is $s_i = \beta \alpha_i$, and so $R^\infty \approx 2 \sum_i \alpha_i^2 p_0 (1 - p_0) \beta N_e = 2\beta N_e V_A = 2N_e R^0$, where $V_A^0 = \sum_i \alpha_i^2 p_0 (1 - p_0)$ is the initial additive genetic variance, and $R^0 = \beta V_A^0$ is the response in the first generation.

This calculation is valid in the limit of weak selection at every locus, reflected in the fact that the probability of fixation is approximated by the first-order perturbation to neutrality. It is also only strictly valid when the trait is additive, excluding any form of interaction between genes. Indeed, it seems very hard to generalize to allow for epistasis, because, then, the calculation of the fixation probability of an allele depends on the trajectories of the alleles at all other loci it interacts with.

An alternative derivation is possible, which focuses on the change in variance components due to drift. Consider an additive trait that is determined by very many loci, each under negligibly weak selection. Drift will disperse allele frequencies, decreasing the additive variance by a factor $(1 - 1/2N_e)$ per generation. Therefore, the total selection response is just $R^\infty = \sum_{t=0}^{\infty} \beta V_A^0 (1 - 1/2N_e)^t = \beta 2N_e V_A^0$, as before. Both arguments show that the total response is proportional to the initial additive variance in the population (which, for an additive genetic architecture with no dominance, where heterozygotes display the mean phenotype of the homozygotes, corresponds to the total initial genetic variance). This is to be expected, because both assume weak selection at individual loci, so that the variance declines primarily through random drift.

What is the effect of epistasis on the total response of the population? Classical quantitative genetics (e.g., ref. 18) uses statistical arguments to derive expressions for the expected components of genetic variance in an inbred population. Barton and Turelli (19) derive these using an explicit population genetic model, assuming two alleles per locus; Hill et al. (20) show how these are related to the classical expressions, and that they apply without any constraint on the distribution of allelic effects at each locus. In the following, we make use of these results to derive expressions for the expected long-term response of populations with arbitrary genetic architectures. We first focus on haploid populations and then extend our results to diploid populations.

The Long-Term Response for Haploids. Suppose that a haploid population initially has additive genetic variance V_A , pairwise epistatic variance V_{AA} , and so on for higher-order epistatic components; it is convenient to denote the k th order variance component as $V_{A(k)}$. Throughout, we will assume that the population is always at linkage equilibrium (such that $\langle p_i^{t+1} p_j^{t+1} \rangle = \langle p_i^{t+1} \rangle \langle p_j^{t+1} \rangle$, where the expectation $\langle \cdot \rangle$ is taken over multiple realizations of the process). If the expected change in allele frequency due to selection at any locus ($\langle \Delta p_i \rangle = \beta A_i p_i^t q_i^t$, where β is the selection gradient and A_i is the effect of locus i on the trait) is small compared with its variance ($\text{Var}(p_i^{t+1}) = \langle (p_i^{t+1})^2 \rangle - \langle p_i^{t+1} \rangle^2 = p_i q_i / N$; according to the Wright–Fisher model), we can assume that the dynamics of the variance components are dominated by drift. This does not necessarily mean that selection is weak on the trait; if the trait is affected by many loci, these conditions can be met at any individual

locus even with strong selection on the trait (the infinitesimal model). Under these conditions, after an arbitrary period of inbreeding, the expected additive variance of a haploid population is

$$\begin{aligned} \langle V_A^t \rangle &= (1 - F_t) (V_A^0 + 2F_t V_{AA}^0 + 3F_t^2 V_{AAA}^0 \dots) \\ &= (1 - F_t) \sum_{k=1}^{\infty} k F_t^{k-1} V_{A(k)}^0 \end{aligned}$$

where F_t is the probability of identity between two genes at a locus (equation 29 of ref. 19, but see also *Dynamics of Variance Components Under Genetic Drift*). Assuming a uniform rate of inbreeding, so that $1 - F_t = (1 - 1/N_e)^t$, the total response to directional selection β is

$$\begin{aligned} \langle R^\infty \rangle &= \beta \sum_{t=0}^{\infty} \langle V_A^t \rangle = \beta \sum_{t=0}^{\infty} (1 - F_t) \sum_{k=1}^{\infty} k F_t^{k-1} V_{A(k)}^0 \\ &= \beta N_e \sum_{k=1}^{\infty} V_{A(k)}^0 = \beta N_e V_G^0. \end{aligned} \quad [1]$$

Remarkably, the total response to directional selection is proportional to the initial total genotypic variance, $V_G^0 = \sum_{k=1}^{\infty} V_{A(k)}^0$ —a simple extension of Robertson’s (17) result. The change in mean can thus be very much greater than the change in the first generation, if N_e is large (Fig. 1). For a given initial total genetic variance, V_G^0 , the initial response is slower with epistasis, because it is proportional only to the additive component, V_A^0 . However, as allele frequencies fluctuate randomly, epistasis generates additional additive variance, such that the total selection response is the same. For a given initial genotypic variance, epistasis slows the initial response, but does not affect its final value, whereas, for a given initial additive variance, epistasis increases the final response by V_G^0/V_A^0 , compared with an additive genetic architecture.

Our result is based on the dynamics of genetic variances, but Robertson’s (17) result (reproduced at the beginning of the section) suggests an alternative view: As alleles fluctuate in frequency, they change the average effects of alleles at other loci, thereby changing their probability of fixation. In particular, as we show in




Fig. 1. Long-term trait value reached by finite populations with a pairwise genetic architecture as a function of the initial fraction of epistatic variance V_{AA}^0/V_G^0 for different population sizes as indicated. For each population size, we increase the variance in interaction strength, keeping the variance of main effects constant at $\sigma_a = 1/10$, from $\sigma_e = 10^{-3}$ to $\sigma_e = 1/2$, resulting in an initial genetic variance ranging from $V_G = 2$ to $V_G \approx 3$. Initial epistatic variance saturates with the strength of interactions (see *Calculating Expectations for the Initial Components of Variance*). Lines represent the theoretical prediction from Eq. 1. Selection strength was set to $\beta = 10^{-3}$. Circles represent means of 10^4 simulations, and error bars (too small to be seen) represent the SE.

The Interaction Between Strong Directional Selection and Epistasis. Alleles that are initially deleterious but become beneficial as the response unfolds should contribute to this increase in long-term response. This view predicts that an increase in long-term response should be correlated with an increase in the number of alleles that are beneficial at the end of the response, which we confirmed through simulations (Fig. S1).

The Long-Term Response for Diploids. For a diploid population, with no dominance (meaning that the phenotype of the heterozygote is simply the mean of the homozygous phenotypes), epistasis has a stronger effect. From equation 45 of Barton and Turelli (19),

$$\begin{aligned} \langle V_A^t \rangle &= (1 - F_t) \left(V_A^0 + 2(2F_t)V_{AA}^0 + 3(2F_t)^2 V_{AAA}^0 \dots \right) \\ &= (1 - F_t) \sum_{k=1}^{\infty} k(2F_t)^{k-1} V_{A(k)}^0. \end{aligned}$$

Assuming that $1 - F_t = (1 - 1/2N_e)^t$, the total response to directional selection β is

$$\begin{aligned} \langle R^\infty \rangle &= \beta \sum_{i=0}^{\infty} V_A = \beta \sum_{i=0}^{\infty} (1 - F_t)^i \sum_{k=1}^{\infty} k(2F_t)^{k-1} V_{A(k)}^0 \\ &= 2N_e \beta \sum_{k=1}^{\infty} 2^{k-1} V_{A(k)}^0 > 2N_e \beta V_G^0. \end{aligned}$$

Now, epistasis can increase the total response disproportionately. This is because we are calculating the contribution of epistatic components to the variance of a population that will ultimately be completely inbred—which amplifies the genetic variance contributed by each locus by a factor of 2 per order of interaction. (To see this, compare the variance of the trait in a diploid population at Hardy–Weinberg proportions, $q^2 : 2pq : p^2$, with that after complete inbreeding, $q : 0 : p$). However, it is unlikely that higher-order components will be large enough for this effect to be substantial (6). With two alleles per locus, the k th order variance is proportional to $(pq)^k$, the product of allele frequencies across the interacting loci. Because $pq < 1/4$, this decreases at least as 2^{-2k} .

We have assumed constant directional selection, β , on the trait. However, the result readily extends to more complex forms of selection. Assuming that fitness increases monotonically with the trait value, we can transform to a scale on which log fitness depends linearly on the trait. On this new scale, the variance components will be different, but our general results show how the total response depends on these components, when measured on the appropriate scale.

The Interaction Between Strong Directional Selection and Epistasis.

We now turn to the deterministic case, where selection is the predominant process. Imagine that, in the initial haploid population, alleles at n loci have marginal effects α_i , so that, with a selection gradient β , they have selective advantage $\beta\alpha_i$. We label the alleles by their initial effect, so that $\alpha_i > 0$. If the trait is additive, and the population infinitely large, then, if selection remains constant, the fittest genotype $\{1, 1, \dots\}$ will eventually fix, and the mean will increase by $R^\infty = \sum_i \alpha_i q_i^0$, where $q_i = 1 - p_i$. However, with epistasis, the marginal effect of alleles will change through time. If these effects retain the same sign, the same final state will be reached, but it will have a different trait value. If some effects change sign, then a different genotype may be reached. There may still be a single global fitness peak, or there may be multiple peaks, so that the outcome is sensitive to the initial allele frequencies. However, we avoid considering this complication by simply comparing the final state of populations with and without epistasis. More precisely, we compare the total change in mean of an epistatic trait with the total change in mean of an additive trait, starting with the same allele frequencies and marginal effects.

Typically, the initial distribution of allele frequencies will be U-shaped, as expected for a population in a stationary distribution with low mutation rates per site ($N_e\mu < 1$ for haploids). Thus, an enormous range of variation can be released as initially rare alleles increase, and are brought together by recombination. We will assume that the population is so large that alleles that are initially deleterious, but will ultimately be favored, will not be lost from the population. This is not realistic, but it is equivalent to assuming a low rate of recurrent mutation, such that alleles are continually reintroduced, so that, if they are favored, they will become common enough to increase deterministically. We assume that epistasis is not so strong that the marginal effects of alleles fluctuate more rapidly than the timescale needed for them to establish.

If most alleles are very rare or very common, they will spend a relatively short time passing from low to high frequency. Thus, we can caricature the process as a series of separate substitutions; at each substitution, the marginal effects of all of the other alleles change. This will change the time at which favorable alleles substitute, and may cause them to lose their advantage entirely. This caricature is not quite the same as the Strong Selection Weak Mutation approximation (21), in which populations are assumed to be fixed for a single genotype, and to evolve by fixation of new mutations, because rare alleles still follow a complex dynamic. In reality, substitutions may overlap, and their effects change continuously. This limiting scenario is helpful in understanding the following argument, but, because we focus on the final outcome, we do not need to assume that alleles substitute one by one. All we need to know is what is the difference between the final and initial effects of an allele.

Assuming only pairwise interactions, two alleles per locus, and linkage equilibrium throughout (see *Materials and Methods*), the ultimate change in mean is

$$R^\infty = \sum_i \alpha_i (p_i^\infty - p_i^0) + \frac{1}{2} \sum_{i \neq j} \epsilon_{ij} (p_i^\infty p_j^\infty - p_i^0 p_j^0).$$

At any time, allele i has selective advantage $s_i = \beta A_i$, where $A_i = \alpha_i + \sum_j \epsilon_{ij} p_j$ is the average effect of locus i on the trait mean that changes linearly with the frequencies of all of the other alleles (see *Materials and Methods*). Recall that we labeled alleles such that, initially, $A_i > 0$ for all alleles. For simplicity, assume that the set S_0 of alleles starts at low frequency, and a set S_1 starts at high frequency. Therefore, n_0 alleles are expected to sweep from low to high frequency, thereby perturbing the selection on all other alleles. Provided that epistasis is sufficiently weak that no selection coefficients change sign, this set will sweep to fixation, and the difference between the response of an additive trait with the same initial effects (and therefore necessarily with the same initial additive variance) will be $\Delta R = \sum_{ij \in S_0} \epsilon_{ij}$. Therefore, as long as epistasis is not directional, epistasis will, on average, make no difference to the total response. More precisely, provided that the total pairwise epistasis, summed over all pairs of alleles that sweep from low to high frequency, is zero, epistasis will have no net effect on the selection response.

This calculation is based on the assumption that additive effects cannot change sign; if interactions are strong enough to turn an initially beneficial allele, k , into a deleterious one as frequencies at other loci change, the long-term response can be different, because the difference between the two architectures is now $\Delta R = -\alpha_k + \sum_{j \neq k} \epsilon_{kj}$. By assuming that the epistatic coefficients are independent random variables, we can approximate the probability that epistasis will change the ultimate genotype, and the expected change in trait mean that results.

The probability that an allele switches sign as the remaining favorable alleles increase is simply the probability that the sum of all interactions is less than the negative of its main effect: $p_{\text{flip}} = J_\epsilon(-\alpha_k) = \int_{-\alpha_k}^{\infty} P_\epsilon(x) dx$, where J_ϵ is the cumulative probability distribution and P_ϵ is the probability density of the sum of all epistatic interactions with the i th locus, and the n_0 loci that are

sweeping from low to high frequency. We assume that the number of interactions is large enough, and that they are sufficiently independent, that this is approximately Gaussian: $\epsilon = \sum_{i \neq j \in S_0} \epsilon_{ij} \sim \mathcal{N}(0, \sqrt{n_0} \sigma_\epsilon)$, where σ_ϵ is the standard deviation (SD) of the ϵ_{ij} . Note that we only need to assume that the ϵ_{ij} are drawn independently from a distribution with finite variance, so that their sum approaches a Gaussian. Note also that one of the n_1 alleles that is at high frequency will interact with all of the n_0 alleles that sweep from low to high frequency, whereas one of the n_0 sweeping alleles will only interact with the $n_0 - 1$ other alleles that are also sweeping. However, we assume that $n_0 \gg 1$, and so approximate $n_0 - 1$ by n_0 throughout.

We can now calculate the average probability of an allelic reversal by integrating over the distribution of the main effects, $P_\alpha[\alpha]$,

$$\langle p_{\text{flip}} \rangle = \int_0^\infty J_\epsilon(-\alpha) P_\alpha(\alpha) d\alpha. \quad [2]$$

This can be calculated explicitly for specific choices for P_α ,

$$\langle p_{\text{flip}} \rangle = \frac{1}{\pi} \tan^{-1}(\mathcal{E}) \quad (\text{half-Gaussian})$$

$$\langle p_{\text{flip}} \rangle = \frac{1}{2} \left(1 - e^{-\mathcal{E}^2} \operatorname{erfc}[\mathcal{E}] \right) \quad (\text{Exponential})$$

where $\mathcal{E} = \sqrt{n_0} \sigma_\epsilon / \sigma_\alpha$ is the ratio of the SDs of the strength of genetic interactions versus the main effects, the relative “strength” of genetic interactions. These two functions necessarily increase linearly from zero for small \mathcal{E} , and saturate at 1/2 for large \mathcal{E} ; they differ by, at most, 4.4%. Comparing with numerical simulations, we see that this calculation captures the expected number of alleles that would be deleterious, compared with the original ultimate genotype (in which alleles remain positive throughout the evolutionary response) (Fig. 2). However, it assumes that all loci respond independently to the changing genetic background. In reality, as alleles switch sign and go extinct, other loci will be affected as well. When comparing to the actual global peak of the fitness landscape, we see that the actual number of alleles that change sign is less than predicted by our approximation of independence (Eq. 2). This is because, for each allele that reverses its sign, an interaction is effectively removed, thereby reducing p_{flip} for the remaining loci. Numerically, we found that a good approximation to the expected number of allelic reversals is $1 - p_{\text{flip}}(1 - p_{\text{flip}})$ of that assuming independence, independently of number of loci (Fig. 2A). This expression should converge to the one assuming independence at low values of p_{flip} , which we confirmed through simulations (Fig. 2A). Each of these allelic reversals (or “flips”) will translate into some change in trait value compared with an additive population. For each allele, this increase will be $\Delta z_i = \sum_{j \neq i} \epsilon_{ij}$, and so the average contribution of each allele that switches sign will be

$$\langle \Delta z_i \rangle = \frac{1}{J_\epsilon(-\alpha)} \int_0^\infty \int_0^{-\alpha} \epsilon P_\epsilon(\epsilon) P_\alpha(\alpha) d\epsilon d\alpha. \quad [3]$$

For specific distributions of main effects,

$$\langle \Delta z_i \rangle = \frac{1}{\langle p_{\text{flip}} \rangle} \frac{\mathcal{E}^2 \sigma_\alpha}{\sqrt{2\pi} \sqrt{1 + \mathcal{E}^2}} \quad (\text{half-Gaussian}) \quad [4]$$

$$\langle \Delta z_i \rangle = \frac{1}{\langle p_{\text{flip}} \rangle} \frac{e^{\mathcal{E}^2} \operatorname{erfc}[\mathcal{E}] \mathcal{E}^2 \sigma_\alpha}{\sqrt{2}} \quad (\text{Exponential}). \quad [5]$$

These functions are again very similar, tending to zero for weak epistasis, and to $c \sigma_\alpha + \sqrt{2/\pi} \sqrt{n_0} \sigma_\epsilon$ for large \mathcal{E} , where $c = 0.505$ for the half-Gaussian, and 0.444 for the exponential. Finally, the

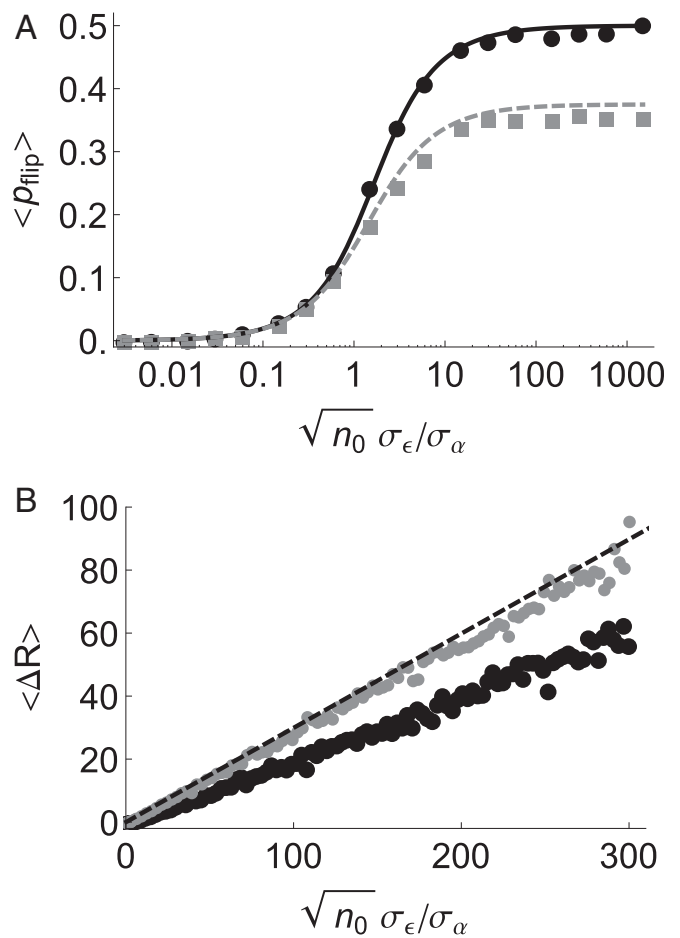


Fig. 2. (A) Mean fraction of alleles that reverse their sign when all initially beneficial alleles go to fixation (black circles) or in the genotype of highest trait value (gray circles), as a function of the relative strength of genetic interactions $\mathcal{E} = \sqrt{n_0} \sigma_\epsilon / \sigma_\alpha$. The black line denotes the prediction assuming independence of alleles, and the gray dashed line is the same prediction multiplied by $1 - p_{\text{flip}}(1 - p_{\text{flip}})$. (B) Expected increase in trait value if all of the initially beneficial alleles fix (black circles) or at the actual global peak (gray circles). Dashed line corresponds to the prediction obtained based on the gray line in A. All points correspond to the mean of 100 trials, randomizing the coefficients for the indicated variances for ϵ_{ij} and α_i for $n = 10$ loci.

total trait increase due to allelic reversals is simply the expected number of alleles that reverse times the increase due to each of those,

$$\langle \Delta R \rangle = n \langle p_{\text{flip}} \rangle \langle \Delta z_i \rangle. \quad [6]$$

Substituting from Eqs. 2 and 4, we see that $\langle \Delta R \rangle$ increases close to linearly with σ_ϵ , as $\sim n \sqrt{2/\pi} \sqrt{n_0} \sigma_\epsilon$ (Fig. 2B). Note that, whereas only the n_0 loci that sweep from low to high frequency contribute (via epistasis) to changes in marginal effects, all $n = n_0 + n_1$ loci may be affected by these changes. An initially favorable allele that is still at low frequency when it loses its advantage will never appear, whereas one that is already at high frequency will sweep back down to low frequency (recall that we assume either an extremely large population or recurrent mutation, so that deleterious alleles never get fixed). Either way, any of the n initially favored alleles that lose their advantage will be eliminated. The previous calculations are upper bounds for the effect of epistasis on the long-term response of a population, because we considered that n_0 loci have low initial allele frequencies. If initial allele frequencies are appreciable, part of the interaction

effect is absorbed by the initial average effect of the loci, and alleles may not reverse their sign, but this can be seen as a reduction in n_0 .

When drift dominates, the ultimate response depends only on initial variance components. When selection dominates, the effect of epistasis on the ultimate response depends on reversals in selection on individual alleles, so that a different genotype is reached; this depends on the ratio of SDs of epistatic versus main effects, multiplied by the square root of the number of interacting alleles that sweep from low to high frequency ($\mathcal{E} = \sqrt{n_0}\sigma_\epsilon/\sigma_a$), as we showed above. The transition between the two regimes depends only on the selection strength each locus experiences (Fig. 3); the trait can be under high selection strength, but, as long as it is highly polygenic, the prediction derived for the drift-dominated regime will hold (Fig. 3) because, for the same initial genetic variance, the selection coefficients scale as $n^{-1/2}$ with the number of loci. In the limit of infinite numbers of genes, this prediction would hold regardless of the selection gradient on the trait. In the selection-dominated regime, the distribution of allele frequencies is mostly irrelevant because the ultimate response depends on alleles that are initially vanishingly rare (or absent if we allow mutation). The key parameter is $\mathcal{E} = \sqrt{n_0}\sigma_\epsilon/\sigma_a$, which is not constrained by the initial variance components: If alleles are at extreme frequency, the epistatic variance must be small, even if \mathcal{E} is large. However, we can relate this parameter to variance components in the F2 population formed by crossing the ancestral with the derived population (see *Measuring Strength of Interactions from F2 Crosses* and Fig. S2). This implies that the initial epistatic variance does not predict the long-term response of the population, because the population may contain strong interactions (large \mathcal{E}) that are not manifested as epistatic variance (which depends on allele frequencies) but can contribute to the long-term response as allele frequencies reach appreciable levels.

Discussion

The role of epistasis in evolution has long been controversial. Wright (1) argued that epistasis would cause populations to become trapped at local “adaptive peaks” and proposed that a “shifting balance” between selection and random drift could allow them to explore alternative peaks, so as to move toward the global optimum. This theory motivated much work on the structure of natural populations, yet it remains unclear whether adaptation is significantly slowed by trapping on local peaks (1, 22). Mayr (23) criticized the supposed neglect of epistasis by “bean-bag genetics,” provoking a robust defense by Haldane (24). More recently, it has been proposed that epistatic variance can be “converted” into additive variance following a bottleneck, aiding adaptation (25). The failure of large genome-wide association studies to assign much heritable variance to specific loci (the so-called “missing heritability”) has been attributed to epistasis (26, 27), although this explanation is unnecessary (28). Overall, the practical success of the additive model in quantitative genetics appears hard to reconcile with the strong molecular interactions between genes.

We investigate how epistasis affects the response to selection, by asking a simple and clearly defined question: By how much does epistasis influence the ultimate change in the mean of a selected trait? We compare the effects of directional selection on two populations that initially have the same genetic variance for a trait; in one, inheritance is strictly additive, whereas, in the other, there can be strong gene interactions. We find simple results in the two extreme cases, where either drift or selection dominate.

In the first case, where drift is stronger than selection on individual alleles, the outcome can be predicted from the initial variance components. This seems remarkable but can be understood as a perturbation to neutrality: When selection is spread over very many loci, its effect on any one locus is weak relative to drift, and so the variance components are hardly perturbed by selection. This is an extension of the infinitesimal model to non-additive inheritance (29). For haploids, the total selection response is proportional to the initial genotypic variance (including both




Fig. 3. The transition from the drift-dominated to the selection-dominated regimes. Stochastic simulations with increasing scaled selection ($N_e\beta$) showing the total trait increase in units of initial SD of trait value for $n = 50$ (squares) and $n = 1,000$ (circles) loci. Dotted line corresponds to the drift-dominated prediction ($R = \beta N_e V_G^0$), and dashed lines correspond to the deterministic limit taken from numerical analysis, for an epistatic architecture (black) and an additive architecture (gray) of the same initial genetic variance. Initial allele frequencies were sampled from a beta distribution with mean $1/4$ and variance $1/16$, $\sigma_a = 1/10$, and $\sigma_\epsilon = 0.5$. Simulations were performed either with $\beta = 10^{-2}$ or $\beta = 10^{-5}$ and varying N_e to yield the reported product $N_e\beta$.

additive and nonadditive components). For diploids, k th-order components of variance have effect multiplied by 2^k . Nevertheless, it is extremely difficult to find plausible models in which such higher-order variance components are significant (10). One way to see this is to imagine a population in which all additive effects are zero—as would be the case at an equilibrium under balancing selection. However, any change in allele frequencies will necessarily generate nonzero additive effects, and, consequently, substantial additive variance. Higher-order epistatic variance is also likely to be small if alleles are at extreme frequencies: the k th-order epistasis is proportional to the product of $\langle pq \rangle$ across k loci, and so cannot be large (10).

In contrast, when selection is strong relative to drift ($Ns \gg 1$), the population will fix at an adaptive peak. This requires strong selection on each allele ($Ns \gg 1$), but also recurrent mutation, so any allele that is favored will eventually succeed. Thus, the issue is primarily about the genotype–phenotype map rather than any population genetic process. We show that, unless epistasis is systematically biased, the ultimate response is increased, on average, only if the fittest peak is changed, because epistasis changes some alleles from being favored to being deleterious. The expected response can be predicted by a simple argument, which assumes that the net effect of the changing background is a normally distributed perturbation, and that the chance that selection on an allele changes sign is simply the chance that this random epistatic perturbation exceeds the main effect of the allele.

When drift dominates, our conclusions follow simply from the variance components, without further assumptions. When selection dominates, we assume that epistasis is not systematically biased toward (or against) interactions between favorable alleles. If a systematic bias is allowed, then epistasis can have an arbitrarily large effect. To see this, imagine a trait that is some function $g[z]$ of an additive trait, z . If $g[z]$ curves upward, then the selection response will accelerate, and can become much larger than the corresponding additive model with the same initial additive variance. Indeed, if we are allowed to assign trait values to genotypes arbitrarily, we can construct paths that follow any pattern of change through time. Our conclusions about the ultimate effects of selection depend on the assumption that epistasis adds a random perturbation, without any definite bias. Empirically, this seems to be the consensus: Epistasis is pervasive between genetic loci but is generally nondirectional in the sense

that there is no prevalent pattern of either positive or negative interactions (30).

Throughout this paper, we assumed that the populations remained at linkage equilibrium. However, linkage disequilibrium (LD) will be generated in a number of ways. Drift alone will produce some LD, although this should be symmetric around no disequilibrium and so should not, on average, have strong effects. In the selection-dominated regime, however, LD should be generated consistently in a directional fashion. Nagylaki's theorem (31) states that, under weak selection, as we assume here, the population is guaranteed to approach and remain close to linkage equilibrium. Furthermore, the fact that we assume that interactions have no preferred direction should further hinder the buildup of LD. Nevertheless, LD will affect the response: A truly infinite population with no recombination (full linkage) is guaranteed to find the global peak, and linkage can only affect the rate at which this is approached. In this sense, linkage can help because recombining populations can get "trapped" at local peaks.

Do gene interactions and epistasis affect the long-term response? In the drift-dominated regime, which can apply even when selection on the trait is strong for polygenic traits, the long-term response is merely proportional to the initial genetic variance. Epistatic architectures can reach higher trait values compared with additive populations of the same additive variance, because the former necessarily harbor more genetic variance. However, this effect will be small and the response slower because epistatic variance typically represents a small fraction of the total genetic variance. In the selection-dominated regime, it is the specifics of the interaction structure that matter. Substantial increases of the long-term response can be reached when interactions are strong and induce allelic reversals, but the initial epistatic variance of the population is not predictive of this (see *Deterministic Limits as a Function of Allele Frequency Distribution* and Figs. S3 and S4). These results set expectations for the effect of epistasis on the long-term response under directional selection and help reconcile the success of the infinitesimal additive model (11) with the biological fact that genetic interactions are pervasive in nature.

Materials and Methods

Simulations. Except where noted, we assume a haploid population of effective size N_e ; similar results can be obtained for a diploid population with $2N_e$ copies of each gene, provided there is no dominance. We also assume

that the trait is under weak directional selection, with selection gradient β , and that the population is close to linkage equilibrium. We ignore the environmental component of variance, because this does not affect the response to directional selection. Our results for finite populations under weak selection are independent of the trait architecture. However, for simulations, we assume that the trait of a diallelic haploid genotype of n loci $\underline{X} = (X_1, \dots, X_n)$ is given by $z(\underline{X}) = \sum_{i=1}^n \alpha_i X_i + 1/2 \sum_{i \neq j} \epsilon_{ij} X_i X_j$. Assuming that the population is at linkage equilibrium, one can write, for the mean trait value, $\bar{z} = \sum_i \alpha_i p_i + 1/2 \sum_{i \neq j} \epsilon_{ij} p_i p_j$. The effect of a particular locus on the trait mean is then $A_i = \partial \bar{z} / \partial p_i = \alpha_i + \sum_{j \neq i} \epsilon_{ij} p_j$. We typically assume that main effects α_i and epistatic coefficients ϵ_{ij} are randomly sampled from Gaussian distributions $N(0, \sigma_\alpha)$ and $N(0, \sigma_\epsilon)$, respectively.

We assume directional selection for increased trait values such that the fitness of a particular genotype \underline{X} is $w(\underline{X}) = 1 + \beta z(\underline{X})$. At linkage equilibrium, the mean fitness of the population can then be written as $\bar{w} = 1 + \beta \bar{z}$. The change in allele frequency at any locus is then $\Delta p_i = p_i q_i (d \ln \bar{w}) / dp_i = \frac{\beta}{\bar{w}} A_i p_i q_i$. The dynamics of macroscopic variables, such as trait mean, variance, etc., can be derived from this expression.

Finite Population Simulations. The simulations were performed by keeping track only of allele frequencies, thereby abolishing LD. Every generation, allele frequencies are updated according to the deterministic expectation, $\Delta p_i = \beta A_i p_i (1 - p_i) / \bar{w}$, where $A_i = \alpha_i + \sum_{j \neq i} \epsilon_{ij} p_j$. A binomial sample of N copies is then sampled with this probability of success for each locus, and the allele frequencies are updated according to this sample. We repeat this procedure until no genetic variation exists at any locus, i.e., the population is fixed for one genotype.

Deterministic Numerical Simulations. Every iteration, we iterate the deterministic recursion $\Delta p_i = \beta A_i p_i (1 - p_i) / \bar{w}$, where $A_i = \alpha_i + \sum_{j \neq i} \epsilon_{ij} p_j$. The recursion is iterated for a set number of generations chosen so that the remaining variance in the population is vanishingly small.

Fitness Landscape Analysis. To find the absolute maximum trait increase, we associate a trait value $z = \sum_{i=1}^n \alpha_i X_i + 1/2 \sum_{i \neq j} \epsilon_{ij} X_i X_j$ to every genotype (X_1, \dots, X_n) and exhaustively search for the maximum of the landscape.

ACKNOWLEDGMENTS. The authors thank Jitka Polechová and Michael Turelli for helpful comments. This work was supported by European Research Council Advanced Grant ERC-2009-AdG-250152. This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under Grant Agreement 618091 Speed of Adaptation in Population Genetics and Evolutionary Computation (SAGE).

- Wright S (1932) The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the VI International Congress of Genetics* (Int Congr Genetics, Ithaca, NY), pp 356–366.
- Hayman BI, Mather K (1955) The description of genic interactions in continuous variation. *Biometrics* 11(1):69–82.
- Whitlock MC, Phillips PC, Moore FB, Tonsor SJ (1995) Multiple fitness peaks and epistasis. *Annu Rev Ecol Syst* 26:601–629.
- Phillips PC (2008) Epistasis—The essential role of gene interactions in the structure and evolution of genetic systems. *Nat Rev Genet* 9(11):855–867.
- Hansen TF (2013) Why epistasis is important for selection and adaptation. *Evolution* 67(12):3501–3511.
- Mäki-Tanila A, Hill WG (2014) Influence of gene interaction on complex trait variation with multilocus models. *Genetics* 198(1):355–367.
- Ávila V, et al. (2014) The action of stabilizing selection, mutation, and drift on epistatic quantitative traits. *Evolution* 68(7):1974–1987.
- Lynch M, Walsh B (1998) *Genetics and Analysis of Quantitative Traits* (Sinauer, Sunderland, MA).
- Cheverud JM, Routman EJ (1995) Epistasis and its contribution to genetic variance components. *Genetics* 139(3):1455–1461.
- Hill WG, Goddard ME, Visscher PM (2008) Data and theory point to mainly additive genetic variance for complex traits. *PLoS Genet* 4(2):e1000008.
- Weber KE, Diggins LT (1990) Increased selection response in larger populations. II. Selection for ethanol vapor resistance in *Drosophila melanogaster* at two population sizes. *Genetics* 125(3):585–597.
- Costanzo M, et al. (2010) The genetic landscape of a cell. *Science* 327(5964):425–431.
- Jasnos L, Korona R (2007) Epistatic buffering of fitness loss in yeast double deletion strains. *Nat Genet* 39(4):550–554.
- Kelly JK (2005) Epistasis in monkeyflowers. *Genetics* 171(4):1917–1931.
- Vissler JAGMD, Hoekstra RF, Ende HVD (1996) The effect of sex and deleterious mutations on fitness in *Chlamydomonas*. *Proc R Soc Lond B Biol Sci* 263(1367):193–200.
- Whitlock MC, Bourguet D (2000) Factors affecting the genetic load in *Drosophila*: Synergistic epistasis and correlations among fitness components. *Evolution* 54(5):1654–1660.
- Robertson A (1960) A theory of limits in artificial selection. *Proc R Soc Lond B Biol Sci* 153(951):234–249.
- Kempthorne O (1954) The correlation between relatives in a random mating population. *Proc R Soc Lond B Biol Sci* 143(910):102–113.
- Barton NH, Turelli M (2004) Effects of genetic drift on variance components under a general model of epistasis. *Evolution* 58(10):2111–2132.
- Hill WG, Barton NH, Turelli M (2006) Prediction of effects of genetic drift on variance components under a general model of epistasis. *Theor Popul Biol* 70(1):56–62.
- Gillespie JH (1983) A simple stochastic gene substitution model. *Theor Popul Biol* 23(2):202–215.
- Coyne JA, Barton NH, Turelli M (1997) Perspective: A critique of Sewall Wright's shifting balance theory of evolution. *Evolution* 51(3):643–671.
- Mayr E (1963) *Animal Species and Evolution* (Belknap, Cambridge, MA).
- Haldane JBS (2008) A defense of beanbag genetics. *Int J Epidemiol* 37(3):435–442.
- Cheverud JM, Routman EJ (1996) Epistasis as a source of increased additive genetic variance at population bottlenecks. *Evolution* 50(3):1042–1051.
- Zuk O, Hechter E, Sunyaev SR, Lander ES (2012) The mystery of missing heritability: Genetic interactions create phantom heritability. *Proc Natl Acad Sci USA* 109(4):1193–1198.
- Manolio TA, et al. (2009) Finding the missing heritability of complex diseases. *Nature* 461(7265):747–753.
- Eichler EE, et al. (2010) Missing heritability and strategies for finding the underlying causes of complex disease. *Nat Rev Genet* 11(6):446–450.
- Fisher RA (1918) The correlation between relatives on the supposition of Mendelian inheritance. *Philos Trans R Soc Edinburgh* 52(2):399–433.
- Mackay TFC (2014) Epistasis and quantitative traits: Using model organisms to study gene-gene interactions. *Nat Rev Genet* 15(1):22–33.
- Nagylaki T (1993) The evolution of multilocus systems under weak selection. *Genetics* 134(2):627–647.