The relative strength of selection on modifiers of genetic architecture under migration load

Stephen R. Proulx [1] Henrique Teotónio [2]

1 Department of Ecology, Evolution, and Marine Biology, UC Santa Barbara, Santa Barbara, CA 93106, USA. **2** Institut de Biologie, École Normale Supérieure, CNRS 8197, PSL Research University, F-75005 Paris, France.

Running title: Selection on modifiers under migration load.

Keywords: mutation selection balance, dominance, epistasis, migration, local adaptation, canalization, modifier theory

Corresponding author: Stephen R. Proulx, Department of Ecology, Evolution, and Marine Biology, UC Santa Barbara, Santa Barbara, CA 93106, USA. Telephone: 805-280-8116. Email: sproulx@ucsb.edu

Statement of authorship: SRP conceived of the study and did the analysis. SRP and HT wrote the paper.

Abstract

Gene flow between populations adapting to differing local environmental conditions creates a "migration load" because individuals might disperse to habitats where their survival is low or because they might reproduce with locally maladapted individuals. The amount by which the mean relative population fitness is kept below one creates an opportunity for modifiers of the genetic architecture to spread due to selection. Prior work that separately considered modifiers changing dispersal or recombination rates, or altering dominance or epistasis, has typically focused on the direction of selection rather than its absolute magnitude. We here develop methods to determine the strength of selection on modifiers of the genetic architecture, including modifiers of the dispersal rate, after populations evolved local adaptation. We consider scenarios with up to five loci contributing to local adaptation and derive a matrix model for the deterministic spread of modifiers. We find that selection for modifiers of epistasis and dominance is stronger than selection for decreased recombination, and that selection for partial reductions in recombination are extremely weak, regardless of the number of loci contributing to local adaptation. The spread of modifiers for a reduction in dispersal depends on the number of loci, pre-existing epistasis and extent of migration load. We identify a novel effect, that modifiers of dominance are more strongly selected when they are unlinked to the locus that they modify. Overall, these results help explain population differentiation and reproductive isolation and provide a benchmark to compare selection on genetic architecture modifiers in finite population sizes and under demographic stochasticity.

Introduction

Gene flow is a fundamental process that introduces, maintains, or homogenizes genetic variation between demes adapted to local environmental conditions. Along with this variability comes a fitness cost, the "migration load", that decreases the population mean fitness below that of a deme lacking immigration (Fisher, 1930; Haldane, 1937; Crow, 1958, 1970; Kondrashov, 1984). It has long been known that selection in natural populations can act to reduce the dispersal or recombination rates between loci experiencing recurrent migration, dubbed the "reduction principle", cf. (Altenberg et al., 2017). Much of the study on the reduction principle has been on proving that the direction of selection is towards reducing mixing of genotypes, rather than specifying the relative magnitude of selection on modifiers of the genetic architecture that can, in principle, diminish migration loads. Heritability for these modifiers has been described in the empirical literature including those that change dispersal rates (Saastamoinen et al., 2018), recombination rates (Rockman and Kruglyak, 2009; Peñalba and Wolf, 2020), and those that alter dominance (Durand et al., 2014; Billiard et al., 2021) and epistasis (Noble et al., 2017; Phillips, 2008) within and between locally adapted loci. It is likely that many of these types of modifiers segregate concurrently within and between demes during and after the process of local adaptation. Hence, it is important to quantify the relative magnitude of potential selection on modifiers of genetic architecture to understand differentiation between populations, stability of hybrid zones and speciation.

Interest about selection on modifiers of genetic architecture under migration load has been spurred by the observation that genomic regions contributing to phenotypic differentiation and reproductive isolation between populations or incipient species do

not often recombine (Turner et al., 2005; Nosil et al., 2009; Carneiro et al., 2013; Cruickshank and Hahn, 2014). How these "genomic islands" of differentiation and isolation originate and are maintained is intriguing. One possibility is that selection causes a build up of linked alleles that contribute to local adaptation, either by favoring mutations, such as chromosomal inversions, that eliminate recombination between pre-existing locally-adapted alleles (Kirkpatrick and Barton, 2006; Bürger and Akerman, 2011; Charlesworth and Barton, 2018), or by favoring mutations that just happen to be linked to locally-adapted alleles when they appear during ongoing local adaptation (Feder and Nosil, 2010; Feder et al., 2012; Flaxman et al., 2013; Yeaman, 2013; Aeschbacher and Burger, 2014; Yeaman et al., 2016). Another possibility is that genomic islands of differentiation and isolation are created by genetic drift and the accumulation of mutations in genome locations of reduced recombination rates under allopatry, in a manner similar to how Dobzhansky-Muller incompatibilities arise and are maintained (Bank et al., 2012; Blanckaert and Hermisson, 2018).

Most theoretical work so far has described the expected direction and strength of selection on recombination modifiers under several simplifying assumptions (Kirkpatrick and Barton, 2006; Charlesworth and Barton, 2018; Bürger and Akerman, 2011; Akerman and Bürger, 2014). First, the recombination modifier is usually modeled as a chromosomal inversion so that its spread can be calculated without taking into account recombination of the modifier into alternative genotypes. Second, the fitness functions are generally assumed to be additive, meaning there is no dominance between locally-adapted alleles or epistasis between locally-adapted loci. Third, migration is usually assumed to be "continent to island" and thus dispersal is uni-directional. Under weak migration and no epistasis between loci, for example, Charlesworth and Barton (2018) show that the strength of selection on an inversion depends on the migration load, and that selection increases as the number of locally-adapted loci increase and is higher when the ancestral genotype has free recombination. In general, when the modifier creates chromosomal inversions, the strength of selection on it is completely determined by the fitness cost incurred by the overall population (Proulx and Phillips, 2005).

When modifiers of genetic architecture are heritable and can be selected, and there is multi-directional dispersal between demes, calculation of relative selection strength on them depends on the full distribution of genotype densities in the demes. Much of the discussion centers on the distribution of offspring genotypes of a parent, but this view obscures the fact that modifiers alter their own associations with alternative genotypes, and therefore that the strength of selection depends both on the frequency of genotype or environment backgrounds in the absence of the modifier as well as on the reproductive value of those states. For example, Proulx and Servedio (2009) explore this interaction when alleles affecting mating preferences are selected under migration-selection balance. They find that selection on changes in the transition probability from a particular genotype or environmental background will be weak if the frequency of that state in the stable stage distribution (absent the modifier) is low, even if the potential fitness benefit of a transmission modification is high. Conversely, selection on changes in the transition probability from a genotype or environment background into a different genotype or environment background will be weak if the absolute reproductive value of the new background state is low.

Selection on modifiers of the genetic architecture under migration load will be most strong for changes in the transition probabilities that affect relatively frequent states

and cause a large change in the absolute reproductive value of their descendant states, possibly leading to local extinction (Szep et al., 2021). Because modifiers affect these transition probabilities without altering the total number of offspring produced, any increase in the probability of transition between two states must be associated with an equal reduction of transmission to other states. Therefore, one must calculate the cumulative effect of these changes in transmission in order to quantify the magnitude of selection on modifiers of genetic architecture under migration load.

We here investigate selection in a scenario with bi-directional dispersal between two demes that face two habitats with differing environmental conditions, where a diploid genotype is determined by a linear chromosome with arbitrary recombination between several adjacent loci (up to five), and where fitness can include dominance and epistasis. Although we assume a linear chromosomal recombination map, if the recombination probability is 0.5, then we recover a scenario with freely segregating chromosomes. Mating between individuals, precisely syngamy, is assumed to be random. We numerically solve for equilibrium – a steady-state between migration, selection, segregation, recombination and mating – of genotype frequencies in each deme and calculate the migration load. We then introduce a rare mutant modifier of the genetic architecture and estimate the selection strength on the invading modifier as the eigenvalue of the transition matrix. The modifier can alter the the overall magnitude of the migration load (in this case being defined as a "canalization" modifier, cf. (Proulx and Phillips, 2005)), but it can also alter the dominance of single or multiple loci, the epistasis between loci, the recombination rate between loci, and the dispersal rate between demes.

The model

Definitions

We consider two demes with reciprocal migration, each adapting to habitats with different environmental conditions. We start the model off by considering an established population that has k loci in a linear chromosome contributing to local adaptation, and that dominance within a locus and epistasis between the loci determine fitness. The population of "resident" haplotypes is followed until it comes to an equilibrium between standing variation, dispersal, recombination, mating and selection. We assume that the genetic architecture is fixed in the resident population, which means, for example, that the recombination rates between loci do not depend on the allele values at those loci. Generally, local adaptation can be maintained so long as dispersal is not too high and when selection for locally adapted alleles is relatively balanced among the habitats. In all of the analyses presented here we assume symmetric migration rates and selection coefficients, and this ensures that the fully polymorphic equilibrium is stable (Akerman and Bürger, 2014). Our method works under asymmetry as well, but it then requires additional checking to ensure that a polymorphism in local adaptation is maintained.

The life-cycle is explicitly defined for the resident population, as the parameters for the modifier mutants are defined relative to the resident population in a transition "invasion matrix". A schematic of the life-cycle for the resident population is shown in Figure 1, and notation is given in Table 1. The resident life-cycle proceeds by joining gametes from different individuals through mass syngamy (analogous to random mating) to generate diploid individuals. Each diploid individual is thus constituted by two

haplotypes. We follow (Kirkpatrick et al., 2002) and index haplotypes where each locus has a binary allele value of 0 or 1 and the index of each haplotype is the binary numerical value of the bit string converted to decimal notation with 1 added. For example, if there are 4 loci then the haplotypes are indexed 1 through 16 and haplotype $0, 1, 0, 1$ has an index value of 5.

Figure 1: Schematic representation of model definitions (left) and the resident population life-cycle (right). During the meiosis stage, allelic state at the different loci are swapped between parental haplotypes according a recombination rate vector parameterizing the probability of swapping between locally-adapted loci. Mendelian segregation is assumed. Parameter notation can be found in Table 1, explanation of the algebra and calculation of selection strength on the modifiers of genetic architecture ($\hat{S}$, $\hat{r}_i$, $\hat{\phi}$, $\hat{\theta}_i$, $\hat{m}$) of the resident population when they are invading can be found in the main text and in the supplemental Mathematica file.

Diploids migrate between the two demes and are then subject to viability selection within each deme (see next section for the individual fitness function). After selection, individuals undergo meiosis with Mendelian segregation and cross-over recombination between homologous chromosomes. Recombination is modeled as if there were a genetic linkage map in a single chromosome. The location of the k loci is given by their index and a recombination vector describes the rate of crossing-over between successive loci, giving $k-1$ recombination rates. After recombination, gametes are produced by defining a table of the frequency of all possible gamete types given each possible pair of parental haplotypes. Following (Kirkpatrick et al., 2002), a square matrix G with dimensions 2^k is defined such that $G_{i,j}$ is itself a vector of length 2^k (technically a tensor of 3 dimensions). Each position l in the vector represents the probability that a parent with haplotypes i and j would produce a gamete with index l. The gamete production table is computed for situations with up to $k = 5$ local adaptation loci and 1 modifier loci. This model has a total of 6 loci with a 64 haplotypes. Thus, the gamete production table

Table 1: Model parameters' notation

Variable	Definition
	General parameters
i	Generally allele values, 1 is adapted to deme 1 and 2 is adapted to deme 2
j	Generally deme value, taking on either 1 or 2
k	The number of genes contributing to local adaptation
h_i	Frequency of a haplotype
$l_i(h_1, h_2)$	The effective number of locally adapted alleles
	Resident genetic architecture parameters
S	Is the total fitness cost associated with the least-adapted genotype
m_j	The migration rate from patch j
r_m	The recombination rate for the modifier
r_i	The recombination frequency between locus i and $i + 1$
ϕ_j	The epistasis parameter in deme j in favor of allele 1
$\theta_{i,j}$	The dominance parameter for locus i in deme j in favor of allele 1
	Modifier genetic architecture parameters
$\hat{x}$	The parameters associated with the modifier allele are indicated with ˆ
	Calculated variables
λ	The dominant eigenvalue associated with the invasion matrix for a modifier allele
L_i	The fitness load in deme i

has $64^3 = 262,144$ entries. These gamete tables are defined generically without inserting numerical values and stored as an object. Values can be efficiently substituted into these stored objects. This means that the gamete table does not need to be re-calculated during simulations, but it does need to be evaluated with the recombination parameter values once per simulation of the modifier invasion.

The life-cycle of the resident population is described using matrix and tensor operations whenever possible. Within each deme, diploid frequencies are given by the multiplicative outer product of the gamete frequency vector with itself. We represent the diploid genotype frequencies after dispersal as the dispersal weighted average of the deme specific frequency matrix. Selection and reproduction is represented by element wise multiplication of the fitness matrix and the allele frequency matrix, followed by normalization of allele frequencies so that they sum to 1. Gamete formation is represented by element wise multiplication of the genotype frequency with the gamete production table and then reduced in rank by summing the frequency of each haplotype produced across all diploid parental genotypes (see supplemental Mathematica file).

Resident population equilibrium

Given that there are k loci and random mating, the 2^k possible haplotypes are followed. Haploid population density in deme j is presented as:

$$\vec{h_j} = \begin{bmatrix} p_{j,1} \\ p_{j,2} \\ \vdots \\ p_{j,2^k} \end{bmatrix}$$

where the haplotype index maps onto the binary bit string of allele values:

i	L1	L2	$\cdots$	L($2^k - 1$)	L(2^k)
$1 \leftrightarrow \{0$	0	$\cdots$		0	$0\}$
$2 \leftrightarrow \{0$	0	$\cdots$		0	$1\}$
$3 \leftrightarrow \{0$	0	$\cdots$		1	$0\}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$2^k \leftrightarrow \{1$	1	$\cdots$		1	$1\}$

We initialize the recursion with the locally favored haplotype being common (frequency of 0.9999) and with the other haplotype frequencies assigned at random. This avoids the pitfall of having the simulation maintain a symmetric but unstable equilibrium. The diploid population density matrix is then formed by taking the outer product of $\vec{h}$ with itself, i.e. multiplying the frequency of all haplotype pairs in all possible orders:

$$\mathbf{d_j} = \vec{h_j} \otimes \vec{h_j},$$

with migration then mixing the densities between demes:

$$\mathbf{dm_j} = \mathbf{d_j}(1-m) + \mathbf{d_{\sim j}}(m),$$

where $\sim j$ inverts the binary value of j, and therefore represents the index for the other deme.

Diploid individual fitness is a function of habitat and genotype and we model it by assuming that each locus has a parameter determining dominance, and that epistasis for all loci are determined by a single parameter. Our approach is to use generalized logistic functions to describe the effective number of locally adapted alleles. Specifically, for a population with k loci, each individual has $2k$ alleles that are each either locally adapted (allele state of 1 in deme 1) or not (allele state of 0 in deme 1). Fitness in habitat i for an individual that inherited haplotypes h_1 and h_2 is then:

$$W_i(h_1, h_2) = e^{\log(1-S)l_i(h_1,h_2)}, \tag{1}$$

where S represents the total fitness cost of having the most maladapted genotype for the resident haplotypes (the "migration load"), and $l_i(h_1, h_2)$ represents the effective fraction of alleles that are maladapted carried by an (h_1, h_2) individual in habitat i. In our parameterization, when the dominance and epistasis parameters are set to 0, $l_i(h_1, h_2)$ is equal to the actual fraction of alleles that are maladapted. In other words, when there is no epistasis or dominance the effective number of locally adapted alleles is the same as the actual number of locally-adapted alleles. If $l_i = 1$, then fitness is $1 - S$, and if $l_i = 0$ then fitness is 1. The effective number of locally-adapted alleles is thus:

$$l_i(h_1, h_2) = \frac{1}{1 + \exp\left(\phi - \log\left(\frac{dom(h_1,h_2)}{1-dom(h_1,h_2)}\right)\right)}, \tag{2}$$

where ϕ is the resident epistasis parameter, and

$$dom(h_1, h_2) = \frac{1}{k} \sum_{i=1}^{k} \left(\frac{\text{het}_i}{1 + \exp(-\theta_i)} + \text{hom}_i \right) \tag{3}$$

with θ_i being the dominance parameter, het_i an indicator variable that takes on 1 if the individual is heterozygous at locus i and 0 otherwise, hom_i an indicator variable that takes on the value 0 if the individual is homozygous for the locally favored allele at locus i and 1 otherwise. This formulation for the fitness function implies that there is complete additivity when $\phi = 0$ and $\theta_i = 0$ for all i. Further, changes in the parameters from 0 are symmetric on a log scale. As ϕ and θ move away from 0, the effect maxes out by values of around 5 (see, for example, Figure 2).

Figure 2: Example of the fitness function showing the effects of dominance and epistasis. In this example the total fitness cost is set at $S = 0.1$, the figures shows a genotype with 20 loci, 5 of which are homozygous for the adapted allele, and 10 of which are homozygous for the maladapted allele. The parameter ϕ or θ is varied from (-3,3). The black line is the fitness value with no dominance or epistasis shown for reference. The blue curve shows the effect when ϕ is varied, while the green curve shows the effect when θ is varied. Note that in this example θ is the same for all loci.

The fitness matrix $\mathbf{W}_j$ in environment j has elements:

$$\mathbf{W}_{j[i_1, i_2]} = w_j(i_1, i_2),$$

dimension $[2^k, 2^k]$, and hence takes the form:

$$\mathbf{W_j} = \begin{bmatrix} w_j(1,1) & w_j(1,2) & \cdots & w_j(1,2^k) \\ w_j(2,1) & w_j(2,2) & \cdots & w_j(2,2^k) \\ \vdots & \vdots & \ddots & \vdots \\ w_j(2^k,1) & w_j(2^k,2) & \cdots & w_j(m,n) \end{bmatrix}$$

After selection, we re-normalize the vector of genotype frequencies by dividing by population mean fitness which ensures that the genotype frequencies sum to 1.

The diploid density after migration and selection is

$$\mathbf{dms_j} = \frac{\mathbf{dm_j} \circ \mathbf{W_j}}{\bar{w}_j},$$

where we use $\circ$ to represent element-by-element matrix multiplication (i.e. the Hadamard product); the mean fitness in deme j being $\bar{w}_j$. This means that the elements of $\mathbf{dms_j}$ are:

$$\mathbf{dms_{j}}_{(i1,i2)} = \frac{\mathbf{dm_{j}}_{(i1,i2)} w_j(i1,i2)}{\bar{w}_j}.$$

After selection, the next stage of the life-cycle is gamete production where segregation and recombination of parental haplotypes takes place. Gamete production is accomplished by Hadamard tensor multiplication, which means that each diploid genotype frequency is multiplied by the vector of recombination probabilities of producing each gamete haplotype. For example, in a one locus scenario the gamete production vector is $v = [1/2, 1/2]$, so if the frequency of a heterozygote is x then the product is $x * v = [x/2, x/2]$. The gamete production tensor is a tensor $\mathbf{G} \in \Re^{2^k x 2^k x 2^k}$ and thus:

$$(d \to h)_j = \mathbf{G} \circ \mathbf{dms_j},$$

where now $(d \to h)_j \in \Re^{2^k x 2^k x 2^k}$ where $(d \to h)_{j(i1,i2,i3)}$ represents the density of type $i3$ haplotypes produced by parents with diploid genotype $(i1, i2)$. While $(i1, i2)$ and $(i2, i1)$ parents are identical in all ways, they appear as separate entries for computational convenience. And the total density of haplotypes is found by summing up the terms for a specific haplotype, that is, $\sum_{i3}(d \to h)_{j(i1,i2,i3)}$. This can be represented in tensor operations by first unfolding, or flattening, $(d \to h)_j$ by a single level and then by taking the inner product with a one's vector.

Having detailed the life-cycle, we find equilibrium haplotype frequencies for the resident population by iterating the recursion equations just presented above, while censusing the population at the haploid gametic stage. Iteration occurs for a maximum of 5000 generations or until the haplotype frequencies change by less than 10^{-8} (see supplemental Mathematica file for further information).

Selection on the modifier

The primary goal of our study is to calculate the magnitude of the leading eigenvalue of a transition invasion matrix describing modifiers of the genetic architecture of a resident (ancestral) populations that has reached a population genetic equilibrium between migration, selection, segregation, recombination and syngamy. We define this transition

invasion matrix for the modifier, the **A** matrix, based on the background genotype and deme that the modifier is associated with. The vector of mutant genotype densities represents the density of individuals that bear one haplotype containing the mutant modifer allele and one haplotype containing a resident (no modifier) haplotype. The modifier is always placed in position 1 of the genetic linkage map. Each pair of adjacent loci experiences an independent chance of crossing-over given by the recombination vector. The modifier may alter the parameters which are indicated by $\hat{S}$ (canalizing modifier), $\hat{r}_i$ (recombination modifier), $\hat{\phi}$ (epistasis modifier), $\hat{\theta}_i$ (dominance modifier), and $\hat{m}$ (dispersal modifier). Here, the modifier alters only one parameter at a time, although in principle it may alter several or all the parameters at once.

Because the modifier is rare, matings between two individuals carrying the modifier are vanishingly rare. It also means that the modifier only affects the parameters of diploid adults that carry it. The recombination modifier in particular acts in "trans" and is dominant, with the mutant gamete frequency table being determined by the gamete frequencies of diploids carrying one copy of the modifier. An important consideration will be to describe how modifier alleles alter the dominance and epistasis parameters in the two habitats. This involves considering how the genotype relates to an implicit phenotype and how this implicit phenotype determines fitness. If the modifier alters phenotype in an environment-independent way, and modifications in phenotype have opposite effects in the two environments, then we expect the modifier allele to have opposite effects on the dominance and epistasis parameters in the two habitats. This would occur if for example there was Gaussian selection in the two habitats with differing optima and the modifier affects phenotypic dominance in the same way in each habitat. On the other hand, if phenoypic dominance is environment specific then we might expect that a modifier would increase the effective number of locally adapted alleles in both habitats simultaneously.

The invasion modifier **A** matrix defines the number and genotype of mutant offspring produced per mutant individual in the prior generation. The construction of this matrix follows the same life-cycle as for the resident dynamic equations, but differs in that the frequency of potential mates is set by the frequency of resident genotypes at the resident equilibrium, and in that the mutant parameters are used. Thus the matrix elements $A_{i,j}$ represent the number of diploid adults in the following generation that carry the modifier and the mutant haplotype and resident haplotype denoted by the index j that descend from a single individual with the mutant haplotype and resident haplotype denoted by the index i.

To obtain **A**, we start by creating the vector describing the density of mutants as ν, which has dimension 2^{2k+1}. This dimensionality is due to enumerating all combinations of mutant and resident haplotypes in both demes ($2^k * 2^k * 2$). We define index mapping functions that relate the index of the vector ν onto the mutant haplotype, the resident haplotype, and the deme:

i		μ_i	R_i	d_i
1	$\leftrightarrow$	$\{1$	1	$1\}$
2	$\leftrightarrow$	$\{1$	2	$1\}$
$\vdots$		$\vdots$		
$2^k + 1$	$\leftrightarrow$	$\{1$	1	$2\}$
$2^k + 2$	$\leftrightarrow$	$\{1$	2	$2\}$
$\vdots$		$\vdots$		
$2 * 2^k$	$\leftrightarrow$	$\{2^k$	2^k	$2\}$

The matrix $\mathbf{A}$ is constructed in three stages, first with a matrix $\mathbf{D}$ combining the segregation and recombination matrix from the two demes and then applying migration, and second applying the matrix $\mathbf{F}$ to adults in the deme they end up in. The matrices $\mathbf{R}_j$, one for each deme, have dimension 2^k each, and detail how segregation, recombination and syngamy produce offspring in the next generation. These matrices only differ in the resident genotype frequencies and its elements are:

$$\mathbf{R}_j[i_1, i_2] = \mathrm{Rp}(\mu_{i_1} \to \mu_{i_2}) * v_{R_{i_2}} * 2,$$

where Rp is a function that returns the probability that a mutant modifier found in a diploid bearing mutant haplotype μ_{i_1} and resident haplotype R_{i_1} would produce a mutant gamete with haplotype μ_{i_2}. $v_{R_{i_2}}$ represents the probability that the gamete fuses with a resident gamete bearing haplotype R_{i_2}. The factor of 2 conditions the probability that an offspring gamete contains the modifier (1/2 of offspring because of Mendelian segregation).

Migration then moves mutant genotypes between the two demes. For this, we create a block matrix composed of the recombination/syngamy matrices:

$$\mathbf{D} = \begin{bmatrix} \mathbf{R_1} * (1 - \hat{m}_1) & \mathbf{R_1} * (\hat{m}_1) \\ \mathbf{R_2} * (\hat{m}_2) & \mathbf{R_2} * (1 - \hat{m}_2) \end{bmatrix}$$

which has dimension 2^{1+2k}.

After migration, we apply selection by creating a vector f of relative fitness values, for each diploid genotype. This vector has length 2^{2k+1} and its elements are:

$$f_i = \begin{cases} \frac{w_1(\mu_i, R_i)(\hat{S}, \hat{\phi}, \hat{\theta})}{\bar{w}_1} & i \leq 2^k \\ \frac{w_2(\mu_i, R_i)(\hat{S}, \hat{\phi}, \hat{\theta})}{\bar{w}_2} & i > 2^k \end{cases},$$

where the fitness function uses the mutant parameters for S, ϕ and θ.

Because the fitness function is applied at the adult stage, the transmission of genotypes in $\mathbf{D}$ to the next stage depends only on their column index, not their row index. So, we can write:

$$\mathbf{F} = \begin{bmatrix} f \\ f \\ \vdots \\ f \end{bmatrix}$$

and hence obtain the transition matrix for the invasion modifier:

$$\mathbf{A} = \mathbf{D} \circ \mathbf{F}$$

It is important to note that the $\mathbf{A}$ matrix involves terms that represent the equilibrium haplotype frequencies of the resident population and the mean fitness of the resident population, but does not require any additional information about the resident dynamics.

Finally, the invasion recursion for the modifier across generations t can then be described by:

$$\nu_{t+1} = \nu_t . \mathbf{A}$$

The long-term growth rate multiplier of the modifier, that is, the relative selection "strength" or selection "coefficient", is the dominant eigenvalue of $\mathbf{A}$ (Otto and Day, 2011). With our approach, the potential loss by genetic drift of the mutant modifier is not modelled and thus we assume that once the modifier appears it will be established in the resident population and follow deterministic dynamics, cf. (Chelo et al., 2013). However, the invasion matrix generated here can be used to define a multi type branching process to study stochastic establishment and will be considered in a future work.

Results and Discussion

Invasion of a modifier of the genetic architecture of a population with no dominance or epistasis, and with free recombination

We first examine selection on modifiers of the genetic architecture of an ancestral population characterized by multiplicative fitness and free recombination, i.e. $r_i = 0.5$, $\phi_j = 0$, and $\theta_{i,j} = 0$. We also define that the modifier has symmetric effects in the two demes, such that $\phi_2 = -\phi_1$ and $\theta_{i,2} = -\theta_{i,1}$. Biologically this amounts to assuming that the effect of genes on phenotype is environment independent and the effect of phenotype on fitness is opposite but symmetric between loci (Rice, 2002; Grieshop et al., 2021). We set relatively high levels of prior local adaptation, with the total fitness cost of the most maladapted genotype being 0.1 and migration rates of 0.05. We further modeled tight linkage between the modifier and the first local adaptation locus ($r_m = 0$) and identical values for each dominance or recombination modifier ($\hat{\theta}_{i,1} = \hat{\theta}_{.,1}$ and $\hat{r}_i = \hat{r}_.$). Modifier parameters are modelled either over the available range (for $\hat{S}, \hat{r}, \hat{m}$ this is between 0 and the ancestral values), or in the case of dominance and epistasis between 0 and 3. For dominance and epistasis the effect of negative values is completely symmetric to that of positive values, representing the opposite polarity of which allele is dominant. The values of ϕ and θ are limited to a maximum of 3 because the effect of the modifier on the effective number of locally adapted alleles is saturated by this point.

Results are shown in Figure 3. Modifiers of recombination have the smallest selection coefficients, and in particular there is very weak selection on recombination modifiers unless they almost completely eliminate recombination between locally adapted alleles. While selection for reduced recombination is slightly higher when more loci contribute to local adaptation, it is still the case that selection only becomes noticeably strong as recombination gets close to zero. Selection for reduced dispersal into unfavorable habitats is generally stronger than selection for reduced recombination. The strength of

Figure 3: Strength of selection on different modifiers as measured by $\lambda - 1$. Parameter values for the ancestral population are $S = 0.1, M = 0.05, R = 0.5; \theta = 0; \phi = 0$. The dashed horizontal line represents the theoretical maximum for the spread of any modifier, cf. (Proulx and Phillips, 2005). To allow direct comparison of the different types of modifiers, we plot the fraction of change in the parameter on the x-axis, so that a value of 0 for a recombination modifier, for instance, represents $\hat{r} = r = 0.5$ while a value of 1 represents $\hat{r} = 0$. Top panels left to right show results for two local adaptation loci, while bottom panels left to right show results for five local adaptation loci. For dominance modifiers, the modifier is assumed to affect dominance at all local adaptation loci simultaneously. Likewise, the recombination modifier is assumed to reduce recombination between all loci simultaneously. For both epistasis and dominance, we allowed the parameter range is $(0, 3)$, and display the fraction of the range changed. Note that negative values of the epistasis/dominance parameters are allowed, and simply polarize phenotype toward the optimum in deme 2, and thus have symmetric effects on the strength of selection. For recombination and migration, the change in the parameter represents the percentage reduction in the rate.

selection on migration modifiers decreases substantially as the number of loci increase, even though the magnitude of the migration load is slightly higher. This is because the migration modifier is directly linked to only a single local adaptation locus and can recombine into a more favorable genetic background after migrating into an unfavorable habitat.

Selection is similarly strong on both epistasis and dominance, and saturates as the modifier parameter increases. Even small increases in the dominance or epistasis parameter results in an appreciable selection coefficient. Selection on epistasis is stronger than selection on dominance, for the same level of phenotypic bias. The difference between selection on epistasis and dominance is higher when there are more loci contributing to local adaptation. Generally, selection for modifiers of epistasis and dominance is stronger than selection for modifiers of recombination and migration. As the number of loci increases, migration modifiers become relatively weakly selected. In all scenarios, selection on a partial increase in dominance or epistasis is much more strongly selected than a partial reduction in recombination.

Effect of the number of local adaptation loci

We next varied the number of genes contributing to local adaptation while keeping the magnitude of dispersal and the migration load constant. More than 5 loci contributing to local adaptation becomes computationally straining, though asymptotic patterns can be observed as we go up to 5 loci. We determine the maximum possible strength of selection when one of the parameters is set to an extreme value. For recombination and migration modifiers this involves setting the parameter to 0. For epistasis and dominance modifiers it involves setting parameters to 10 because this causes complete dominance or epistasis.

Figure 4: Strength of selection on different modifiers as measured by $\lambda - 1$. Parameter values for the ancestral population are $S = 0.1, M = 0.05, R = 0.5; \theta = 0; \phi = 0$. Panel (A) shows the strength of selection of the maximum effect for each parameter (for dominance and epistasis we use a parameter value of 10, for migration and recombination we use a value of 0). Panel (B) shows the effect when the parameters have a modest value (for dominance and epistasis we use a parameter value of 1.5, for migration and recombination we use a value of half the resident population value).

Results show that as the number of local adaptation loci increases from 2 to 5 the total migration load goes up by about 10%, which increases the maximum theoretical strength of selection (Figure 4A). The strength of selection on epistasis and recombination show increases of similar amounts, indicating that the fraction of the migration load that they can ameliorate remains nearly constant. The strength of selection on dominance decreases slightly with increase number of loci. Most different is the response of migration, where the strength of selection decreases by about 50% with number of loci.

We further evaluated the effect of the number of local adaptation loci on the strength of selection for moderate effect modifiers. For migration and recombination modifiers we reduce the parameter value by 50% of the ancestral/resident population rate. For epistasis and dominance we set the parameter to 1.1, because this leads to a 50% reduction in the fitness cost when half the alleles are locally adapted. We find that pattern of change is largely the same, with the main difference being that the strength of selection on recombination is relatively smaller than on the other modifier types (Figure 4B).

Dominance modifiers are more strongly selected when unlinked to the loci they affect

In Figure 3 results are presented for modifiers of dominance that affect all loci simultaneously (global dominance). It is thus interesting to consider modifiers of dominance that affect only one of the local adaptation loci. For this, the modifier is assumed to be at one end of a linear chromosome and tightly linked (r=0) with the first local adaptation locus (locus 1). In Figure 5 we show results from a scenario where the ancestral population has free recombination (r=0.5) between successive local adaptation loci. When the modifier affects locus 1, the strength of selection is lower than when the modifier affects the second locus 2 or a higher index locus in scenarios with more than two loci.

As shown before in Figure 4, selection on global dominance modifiers decreases slightly as the number of local adaptation loci increases. When considering modifiers that affect dominance at single local adaptation loci, we find that the strength of selection is substantially lower, and that the effect of modifying multiple loci is more than additive. Regardless of the number of loci, the strength of selection is stronger when the modifier affects unlinked genes. When only two loci contribute to local adaptation the affect of position is much larger than when five genes contribute to local adaptation.

Figure 5: Strength of selection on dominance modifiers is dependant on modifier linkage. Parameter values for the ancestral population are $S = 0.1, M = 0.05, R = 0.5; \theta = 0; \phi = 0$. In Panel (A) 2 loci contribute to local adaptation. The green curves show selection on a modifier that simultaneously affects dominance of all loci, or a modifier that affects dominance of the locus tightly linked to the modifier (locus 1) or a locus that freely recombines with the modifier (locus 2). In panel (B) 5 loci contribute to local adaptation. Loci 2-5 all freely recombine with the modifier, and so selection on modifiers that affect dominance at any of these loci have the same strength of selection.

Changing the levels of background epistasis

As the level of epistasis changes in the population, selection on modifiers will also change. A full analysis of the evolutionary dynamics would follow successive changes in the genetic architecture which would require more simulation effort and be less generalizable. We vary background epistasis levels as a way to assess how selection would change as the genetic architecture evolves.

Figure 6: Strength of selection on different modifiers depends on the level of pre-existing epistasis. Parameter values for the ancestral population are $S = 0.1, M = 0.05, R = 0.5; \theta = 0$. ϕ is varied between 0 and 10. Strength of selection is shown for modifiers with extreme effects.

Because we consider a modifier of epistasis that shifts phenotypes towards one habitat type, the fitness effect of the modifier is positive in one habitat and negative in the other. This means that the maximum rate of spread of a full canalizing modifier (i.e. a modifier that causes its bearer to have relative fitness of 1) is given by a migration rate dependent eigenvalue, $\lambda(A)$ where

$$A = \begin{vmatrix} \frac{(1-m_1)}{L_1} & \frac{m_1}{L_2} \\ \frac{m_2}{L_1} & \frac{(1-m_2)}{uL_2} \end{vmatrix}. \tag{4}$$

The theoretical maximum strength of selection for an epistasis modifier is given by the eigenvalue of this matrix minus 1: $\lambda(A) - 1$ (see Appendix).

As the amount of epistasis increases, $\lambda(A) - 1$ decreases and approaches a lower asymptote when extreme epistasis effectively creates a single partially dominant locus (Figure 6). Selection for migration reduction increases with ϕ, because the migration modifier can reduce a larger fraction of the migration load when there are fewer loci contributing to adaptation (as in Figure 3).

Selection for additional increases in epistasis shows a complex pattern. First, as ϕ increases from 0, there is an increase in selection for further increases in ϕ. Selection for further increases in ϕ then reaches a local maximum, and after ϕ reaches about 3 then selection for further increases in ϕ substantially decreases. Eventually, further selection for increased ϕ is reduced to zero. When $\phi = 0$, selection for a complete reduction in recombination is about $1/4$ as large as selection for increased ϕ. Like selection for epistasis, recombination is initially more strongly selected as ϕ increases, and does not substantially decrease until ϕ is around 4. Selection decreases to about 4%

of the theoretical maximum, but does not decay towards 0. Once ϕ is about about 4.5, selection for tight linkage supersedes selection for increased epistasis.

In contrast to the other modifier types, selection for dominance decreases as ϕ increases from 0. Once moderate levels of ϕ are reached, selection for tight linkage supersedes selection for recombination.

Weak selection and canalization

For completeness, we explored how the magnitude of selection for amelioration of the migration load affects the strength of selection on modifiers. We did this by varying the parameter S which determines the fitness cost that the most maladapted genotype experiences. Modifiers affecting S can be considered to "canalize" the relevant phenotypes (Proulx and Phillips, 2005). Several (heritable) modifiers that behave in this fashion in mutant or natural populations have been described (Queitsch et al., 2002; Felix and Barkoulas, 2015).

Figure 7 shows that selection on each type of modifier of the genetic architecture of the resident population increases as the fitness benefit of local adaptation goes up. While the response of selection for migration modifiers appears linear, the other types of modifier show a non-linear response. To better appreciate how evolution of genetic architecture would proceed where several modifiers are selected concurrently during local adaptation, we calculated the ratio of the strength of selection on epistasis modifiers to the strength of selection on recombination modifiers, and the ratio of the strength of selection on dominance modifiers to the strength of selection on recombination modifiers. Panel B of Figure 7 demonstrates that even as the benefit of local adaptation becomes very weak, the relative advantage of epistasis/dominance modifiers, as compared with recombination modifiers, remains large and constant.

Figure 7: Selection for modifiers depends on the strength of selection for local adaptation. As the value of S is increased, the strength of selection on all types of modifier goes up (panel A). The relative advantage of epistasis and dominance modifiers, as compared with recombination modifiers, is shown in panel B. Epistasis modifiers have about a 6 fold higher invasion rate than a modifier that eliminates recombination (i.e. an inversion). Parameter values for the ancestral population are $M = 0.05, R = 0.5; \phi = 0; \theta = 0$.

Conclusions

We have presented a novel framework to compare the magnitude of selection between modifiers of the genetic architecture of a structured population undergoing local adaptation in two habitats with different environmental conditions. Unlike standard theoretical approaches in population genetics, our framework allows one to understand selection in the context of the full distributions of genotype densities and their respective reproductive values. We identify four previously unreported features of selection on genetic architecture modifiers. First, selection for reduced recombination, even when the modifier is a chromosomal inversion that reduces recombination to zero, is weaker than for the other types of modifier. Second, when there is a constant total benefit of local adaptation but varying number of loci, increasing the number of loci increases selection strength for modifiers of epistasis and recombination, but decreases selection strength for dominance and migration modifiers. Third, for modifiers of dominance, the strength of selection is higher when they alter dominance at loci that the modifier itself is not linked to. Fourth and finally, even at moderate levels of epistasis in the resident population, selection for further increases in epistasis is larger than selection for tighter linkage of the locally adapted loci. Selection on modifiers to reduce dispersal rates are complex, depending on the number of loci, pre-existing epistasis and initial levels of local adaptation. Further investigation on dispersal rates using our approach is needed to reach more general conclusions.

In general, under migration selection balance when local populations are well mixed, then selection should act to reduce migration (Altenberg et al., 2017). Our results for selection to decrease dispersal are in line with this in that we see consistent directional selection towards lower dispersal rates. However, we find that selection on dispersal is weaker than on other types of modifiers unless the number of local adaptation loci is low and epistasis is strong. Adding complication, it is known that other factors, such as kin selection, might be more important than levels of migration load. When local resource competition between kin is possible, relatively high levels of dispersal are often favored (Hamilton and M., 1977; Ronce, 2007; Billiard and Lenormand, 2005). There is ample selection for dispersal to find a suitable habitat when local population density approaches the carrying capacity, and dispersal to find suitable habitat in these circumstances may lead to by-product dispersal at different distance scales. While dispersal itself is responsible for the entire fitness cost that drives selection for all other modifiers here investigated, selection acts more strongly on other modifiers of genetic architecture than on modifiers that reduce migration load. This seems paradoxical, because if selection acted to reduce migration then selection for other genetic architecture modifiers would disappear. How can we reconcile these ideas? Part of the answer is that modifiers of migration only alter the rate at which individuals carrying the modifier move between habitats, and do not change the impact or rate of interaction with other genotypes that do disperse between habitats. Modifiers that decrease dispersal protect the individual from arriving in a habitat that they are not adapted to, but it does not protect them from the effects of mating with others who themselves dispersed. Selection for reduced migration weakens as the number of loci increase because the immediate cost to a modifier of mating with maladapted individuals decreases.

Charlesworth and Barton (2018) point out that the strength of selection for inversions under continent-island selection is mathematically related to the amount of migration

load generated in the absence of an inversion and the migration load generated once an inversion is fixed. Selection on an inversion, however, is determined by the fraction of the migration load that the inversion is able to ameliorate (Proulx and Phillips, 2005). Charlesworth and Barton (2018) consider a scenario where there is one-way migration from a continent to an island, fitness is multiplicative (i.e. additive on log scale), and locally adapted alleles are segregating on the island at migration-selection balance. An inversion that causes locally adapted alleles to become tightly linked can spread because it causes the marginal fitness of the haplotype carrying the modifier and a full complement of locally adapted alleles to be pegged at 1, and therefore is positively selected with a magnitude of $L - m$, or the load minus the migration rate (see Appendix). Under bi-directional reciprocal dispersal we consistently find that selection to reduce the recombination rate is weaker than other changes to the genetic architecture. Because all of the modifiers we consider act as effective "canalizers" of the migration load, the strength of selection on each is determined by the fraction of the migration load that they can ameliorate. This can further be analyzed by considering how the modifier alters the reproductive value classes of the offspring, cf. (Proulx and Servedio, 2009). We observe that when the modifier only partially reduces recombination then selection on the modifier is quite weak.

Perhaps the most significant result is that the strength of selection for dominance and epistasis modifiers is relatively stronger than that for recombination modifiers. We further found that modifiers of dominance spread fastest when they are unlinked to the locus they modify, but are linked to another locus that also contributes to local adaptation. To better appreciate this, consider two loci, A and B, both varying for alleles that confer local adaptation between habitats. If a modifier of dominance that causes the habitat 1 allele to be dominant at locus B/b and is in tight linkage with B, then it will confer a direct advantage whenever the modifier and B are in habitat 1 but a cost when in habitat 2. If the modifier is linked to the A allele at the A/a locus, then in habitat one it is beneficial when the carrier is heterozygous at the B/b locus. If the modifier migrates into the habitat two, it is only costly if one of the B alleles is not locally adapted. We thus expect that genomic islands of differentiation between populations are more likely caused by selection on partially-dominant mutations that just happen to appear next to already existing locally-adapted alleles. A possible example of this process in natural populations might have occurred at the self-incompatibility locus of Brassica, a non-recombining gene cluster showing high levels of DNA sequence divergence within and between species, where a hierarchy of partial dominance has been described (Durand et al., 2014). More empirical effort will be needed to describe the dominance effect distributions of local adaptation to test the results presented here.

Many prior analyses of selection on modifiers assume that the fitness differences between alternative genotypes are small. There are two main reasons for this, one is mathematical practicality and the other is the belief that the bulk of naturally segregating alleles have small effects on fitness. For the cases we study here, both of these reasons are less applicable. Under bi-directional migration we do not have good approximations for either the amount of migration load or the equilibrium genotype frequencies, so no reduction of mathematical complexity is achieved by assuming weak selection. We can still approximate selection on modifiers using a weak load approximation, and while these approximations perform well (when we numerically solve for the genotype frequencies), they do not provide much additional insight (see Appendix). Second, large differences

in fitness can be maintained with only a few locally-adapted loci and high dispersal rates, patterns which have been observed in natural populations, e.g. (Hoekstra et al., 2005; Hall and Willis, 2006).

In summary, we find little evidence for selection for decreased recombination rates under migration load. Instead, selection should diminish migration loads by favoring modifiers of dominance and epistasis within and between locally-adapted loci. Selection on dispersal modifiers is more complex. Our modelling assumes deterministic dynamics so a venue for future study is to find out if genetic drift and demographic stochasticity qualitatively change our conclusions.

Acknowledgements

We thank Reinhard Burger, Joachim Hermisson, Thomas Lenormand and Denis Roze for discussion, and special thanks to Nick Barton for suggesting the fitness function. HT is supported by the Agence Nationale pour la Recherche (ANR-17-CE02-0017-01, ANR-18-CE02-0017-01).

Appendix - Maximum spread of inversions and epistasis modifiers

We first write several definitions to simplify the notation. The frequency of haplotypes carrying i locally maladapted alleles is in deme d defined as $x_{(d)(i)}$. Define the average fitness among haplotypes that do not disperse (residents) from deme d as $\overline{W}_{HR(d)} = \sum_i x_{(d)(i)}(1-S)^{i/(2k)}$, and the average fitness among haplotypes that do disperse (migrants) as $\overline{W}_{HM(d)} = \sum_i x_{(d)(i)}(1-S)^{(k-i)/(2k)}$. The diploid average fitness is in deme d is $\overline{W}_{D(d)} = \overline{W}_{HR(d)}^2(1-m_d) + \overline{W}_{HM(\sim d)}^2 m_{\sim d}$, where $\sim d$ represents "not deme d", i.e. the other deme. We then define the fitness load among non-dispersing haplotypes as $L_{R(d)} = 1 - \overline{W}_{HR(d)}$, and the fitness load among dispersing haplotypes as $L_{M(d)} = 1 - \overline{W}_{HM(d)}$. The total haploid load is defined as $L_d = L_{R(d)}(1-m_d) + L_{M(\sim d)}m_{\sim d}$.

The invasion matrix for the inversion is given by

$$A_{\text{inv}} = \begin{vmatrix} \dfrac{\overline{W}_{HR(1)}(1-m_1)}{\overline{W}_{D(1)}} & \dfrac{\overline{W}_{HM(1)}m_1(1-S)^{1/2}}{\overline{W}_{D(2)}} \\ \dfrac{\overline{W}_{HM(2)}m_2}{\overline{W}_{D(1)}} & \dfrac{\overline{W}_{HR(2)}(1-m_2)(1-S)^{1/2}}{\overline{W}_{D(2)}} \end{vmatrix} \tag{5}$$

Define λ_{inv} as the dominant eigenvalue of A_{inv} which is the discrete time growth factor, i.e. R_0, for the inversion.

With unidirectional migration from deme 2 to 1, the matrix reduces to a lower diagonal matrix with an eigenvalue of $\lambda = \dfrac{\overline{W}_{D(1)}}{\overline{W}_{D(1)}^2 + m_2}$. We define the selection coefficient for the inversion as $s_{\text{inv}} = \lambda_{\text{inv}} - 1$, which is approximately $s_{\text{inv}} = L - m$, which is equation 4 from Charlesworth and Barton (2018). Note that this equation is quite general, it does not depend on any assumptions about the number of loci. It stems from the simple fact that the inversion haplotype has constant marginal fitness set to 1 and that the frequency of the inversion haplotype is reduced by m each generations (see also Proulx and Phillips (2005) for a general derivation of this result).

Under bi-directional migration with symmetry between the demes, we can find the eigenvalues and then approximate around weak selection. Under symmetry, we suppress the deme indices. We first Taylor expand around weak load, and then Taylor expand around small S, including second order terms (supplementary Mathematica file). We found that our approximation was both qualitatively and quantitatively accurate when we include the second order terms in S, but fails qualitatively when we only include the first order terms. Defining $s_{\text{inv}} = \lambda - 1$, we have

$$s_{\text{inv}} = L - \frac{S}{4}(L+1) - \frac{S^2}{8}\left(L + 1 - \frac{1}{4m}(L_M + 1)\right) \tag{6}$$

It is difficult to gain much intuition directly from this expression.

We develop a similar approximation for a complete modifier of epistasis, using the simplifying assumption that the modifier is always in a haplotype that contains at least one allele that is locally adapted in deme 1. This gives the matrix

$$A_{\text{epi}} = \frac{1}{\overline{W_D}} \begin{vmatrix} (1-m) & m(1-S) \\ m & (1-m)(1-S) \end{vmatrix} \tag{7}$$

Approximating around weak load and weak S, and dropping terms of order $S^2 * L$ we come to

$$s_{\text{epi}} = 2L - \frac{S}{2}(2L+1) - \frac{S^2}{4}\left(1 - \frac{1}{2m}\right) \tag{8}$$

References

Aeschbacher, S. and Burger, R. (2014). The Effect of Linkage on Establishment and Survival of Locally Beneficial Mutations. *Genetics*, 197(May):317–336.

Akerman, A. and Bürger, R. (2014). The consequences of gene flow for local adaptation and differentiation : a two-locus two-deme model. *Mathematical Biology*, 68:1135–1198.

Altenberg, L., Liberman, U., and Feldman, M. W. (2017). Unified reduction principle for the evolution of mutation, migration, and recombination. *PNAS*, 114(12):E2392–E2400.

Bank, C., Burger, R., and Hermisson, J. (2012). The Limits to Parapatric Speciation: Dobzhansky–Muller Incompatibilities in a Continent–Island Model. *Genetics*, 191(3):845–863.

Billiard, S., Castric, V., and Llaurens, V. (2021). The integrative biology of genetic dominance. *Biological Reviews*, 96(6):2925–2942.

Billiard, S. and Lenormand, T. (2005). Evolution of migration under kin selection and local adaptation. *Evolution*, 59(1):13–23.

Blanckaert, A. and Hermisson, J. (2018). The limits to parapatric speciation ii: Strengthening a preexisting genetic barrier to gene flow in parapatry. *Genetics*, 209(1):241–254.

Bürger, R. and Akerman, A. (2011). The effects of linkage and gene flow on local adaptation : A two-locus continent – island model. *TPB*, 80(4):272–288.

Carneiro, M., Baird, S. J., Afonso, S., Ramirez, E., Tarroso, P., Teotonio, H., Villafuerte, R., Nachman, M. W., and Ferrand, N. (2013). Steep clines within a highly permeable genome across a hybrid zone between two subspecies of the european rabbit. *Mol Ecol*, 22(9):2511–25.

Charlesworth, B. and Barton, N. H. (2018). The Spread of an Inversion with Migration. *Genetics*, 208(January):377–382.

Chelo, I. M., Nédli, J., Gordo, I., and Teotónio, H. (2013). An experimental test on the probability of extinction of new genetic variants. *Nature communications*, 4(1):1–8.

Crow, J. E. (1958). Some possibilities for measuring selection intensity in man. *Human Biology*, 30:1–13.

Crow, J. E. (1970). Genetic loads and the cost of natural selection. In Kojima, K. I., editor, *Mathematical models in population genetics*, pages 128–177. Springer, Berlin.

Cruickshank, T. E. and Hahn, M. W. (2014). Reanalysis suggests that genomic islands of speciation are due to reduced diversity, not reduced gene flow. *Molecular Ecology*, 23(13):3133–3157.

Durand, E., Meheust, R., Soucaze, M., Goubet, P. M., Gallina, S., Poux, C., Fobis-Loisy, I., Guillon, E., Gaude, T., Sarazin, A., Figeac, M., Prat, E., Marande, W., Berges, H., Vekemans, X., Billiard, S., and Castric, V. (2014). Dominance hierarchy arising from the evolution of a complex small rna regulatory network. *Science*, 346(6214):1200–1205.

Feder, J. L., Gejji, R., Yeaman, S., and Nosil, P. (2012). Establishment of new mutations under divergence and genome hitchhiking. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367:461–474.

Feder, J. L. and Nosil, P. (2010). The efficacy of divergence hitchhiking in generating genomic islands during ecological speciation. *Evolution*, 64:1729–1748.

Felix, M. A. and Barkoulas, M. (2015). Pervasive robustness in biological systems. *Nat Rev Genet*, 16(8):483–96.

Fisher, R. A. (1930). *The genetical theory of natural selection: a complete variorum edition*. Oxford University Press.

Flaxman, S. M., Feder, J. L., and Nosil, P. (2013). Genetic hitchhiking and the dynamic buildup of genomic divergence during speciation with gene flow. *Evolution*, 67:2577–2591.

Grieshop, K., Ho, E. K. H., and Kasimatis, K. R. (2021). arXiv preprint: Dominance reversals, antagonistic pleiotropy, and the maintenance of genetic variation.

Haldane, J. B. S. (1937). The effect of variation of fitness. *The American Naturalist*, 71(735):337–349.

Hall, M. C. and Willis, J. H. (2006). Divergent selection on flowering time contributes to local adaptation in mimulus guttatus populations. *Evolution*, 60(12):2466–2477.

Hamilton, W. D. and M., M. R. (1977). Dispersal in stable habitats. *Nature*, 269:578–581.

Hoekstra, H. E., Krenz, J. G., and Nachman, M. W. (2005). Local adaptation in the rock pocket mouse (Chaetodipus intermedius): natural selection and phylogenetic history of populations. *Heredity*, 94(2):217–228.

Kirkpatrick, M. and Barton, N. (2006). Chromosome Inversions, Local Adaptation and Speciation. *Genetics*, 434(May):419–434.

Kirkpatrick, M., Johnson, T., and Barton, N. (2002). General models of multilocus evolution. *Genetics*, 161(4):1727–1750.

Kondrashov, A. S. (1984). Deleterious mutations as an evolutionary factor. I. The advantage of recombination. *Genetical Research*, 44:199–217.

Noble, L. M., Chelo, I., Guzella, T., Afonso, B., Riccardi, D. D., Ammerman, P., Dayarian, A., Carvalho, S., Crist, A., Pino-Querido, A., Shraiman, B., Rockman, M. V., and Teotonio, H. (2017). Polygenicity and Epistasis Underlie Fitness-Proximal Traits in the Caenorhabditis elegans Multiparental Experimental Evolution (CeMEE) Panel. *Genetics*, 207(4):1663–1685.

Nosil, P., Funk, D. J., and Ortiz-Barrientos, D. (2009). Divergent selection and heterogeneous genomic divergence. *Molecular Ecology*, 18(3):375–402.

Otto, S. P. and Day, T. (2011). *A biologist's guide to mathematical modeling in ecology and evolution.* Princeton University Press.

Peñalba, J. V. and Wolf, J. B. W. (2020). From molecules to populations: appreciating and estimating recombination rate variation. *Nature Reviews Genetics*, 21(8):476–492.

Phillips, P. C. (2008). Epistasis–the essential role of gene interactions in the structure and evolution of genetic systems. *Nat Rev Genet*, 9(11):855–67.

Proulx, S. R. and Phillips, P. C. (2005). The opportunity for canalization and the evolution of genetic networks. *Am Nat*, 165(2):147–62.

Proulx, S. R. and Servedio, M. R. (2009). Dissecting selection on female mating preferences during secondary contact. *Evolution*, 63:2031–2046.

Queitsch, C., Sangster, T. A., and Lindquist, S. (2002). Hsp90 as a capacitor of phenotypic variation. *Nature*, 417(6889):618–24.

Rice, S. H. (2002). A general population genetic theory for the evolution of developmental interactions. *Proceedings of the National Academy of Sciences*, 99(24):15518–15523.

Rockman, M. V. and Kruglyak, L. (2009). Recombinational landscape and population genomics of caenorhabditis elegans. *PLOS Genetics*, 5(3):1–16.

Ronce, O. (2007). How does it feel to be like a rolling stone? ten questions about dispersal evolution. *Annual review of ecology, evolution, and systematics*, 38:231–253.

Saastamoinen, M., Bocedi, G., Cote, J., Legrand, D., Guillaume, F., Wheat, C. W.,
Fronhofer, E. A., Garcia, C., Henry, R., Husby, A., Baguette, M., Bonte, D., Coulon,
A., Kokko, H., Matthysen, E., Niitepõld, K., Nonaka, E., Stevens, V. M., Travis,
J. M. J., Donohue, K., Bullock, J. M., and Delgado, M. D. M. (2018). Genetics of
dispersal. *Biol. Rev*, 93:574–599.

Szep, E., Sachdeva, H., and Barton, N. H. (2021). Polygenic local adaptation in
metapopulations: A stochastic eco-evolutionary model. *Evolution*, 75(5):1030–1045.

Turner, T. L., Hahn, M. W., and Nuzhdin, S. V. (2005). Genomic islands of speciation
in anopheles gambiae. *PLoS biology*, 3(9):e285.

Yeaman, S. (2013). Genomic rearrangements and the evolution of clusters of locally
adaptive loci. *PNAS*, 110:E1743–E1751.

Yeaman, S., Aeschbacher, S., and Burger, R. (2016). The evolution of genomic islands by
increased establishment probability of linked alleles. *Molecular Ecology*, 25:2542–2558.