

Phenotypic lag and population extinction in the moving-optimum model: insights from a small-jumps limit

Michael Kopp · Elma Nassar · Etienne Pardoux

Received: date / Accepted: date

Abstract Continuous environmental change – such as slowly rising temperatures – may create permanent maladaptation of natural populations: Even if a population adapts evolutionarily, its mean phenotype will usually lag behind the phenotype favored in the current environment, and if the resulting phenotypic lag becomes too large, the population risks extinction. We analyze this scenario using a moving-optimum model, in which one or more quantitative traits are under stabilizing selection towards an optimal value that increases at a constant rate. We have recently shown that, in the limit of infinitely small mutations and high mutation rate, the evolution of the phenotypic lag converges to an Ornstein-Uhlenbeck process around a long-term equilibrium value. Both the mean and the variance of this equilibrium lag have simple analytical formulas. Here, we study the properties of this limit and compare it to simulations of an evolving population with finite mutational effects. We find that the “small-jumps limit” provides a reasonable approximation, provided the mean lag is so large that the optimum cannot be reached by a single mutation. This is the case for fast environmental change and/or weak selection. Our analysis also provides insights into population extinction: Even if the mean lag is small enough to allow a positive growth rate, stochastic fluctuations of the lag will eventually cause extinction. We show that the time until this event follows an exponential distribution,

This study benefited from a PHC Amadeus exchange grant (project number 31642SJ) to Michael Kopp and Joachim Hermisson (University of Vienna, Austria)

M. Kopp

Aix-Marseille Université, CNRS, Centrale Marseille, I2M, 3 Place Victor Hugo, 13331 Marseille Cedex 3, France

Tel.: +33-413-551201

E-mail: michael.kopp@univ-amu.fr

E. Nassar

Aix-Marseille Université, CNRS, Centrale Marseille, I2M, 3 Place Victor Hugo, 13331 Marseille Cedex 3, France

E. Pardoux

Aix-Marseille Université, CNRS, Centrale Marseille, I2M, 3 Place Victor Hugo, 13331 Marseille Cedex 3, France

whose mean depends strongly on a composite parameter that relates the speed of environmental change to the adaptive potential of the population.

Keywords adaptation · evolutionary rescue · global change · Ornstein-Uhlenbeck process · phenotypic lag · population extinction

Mathematics Subject Classification (2000) 92D15

1 Introduction

With global change threatening the survival of many species, an increasing number of theoretical and empirical studies focuses on the potential role of “rapid evolution” and “evolutionary rescue” in preventing extinction (reviewed in, e.g., Hairston et al, 2005; Gonzalez et al, 2013; Kopp and Matuszewski, 2014). An important parameter in these studies is the mode of environmental change, which a given population may experience as either sudden or gradual (e.g., Gomulkiewicz and Holt, 1995). Here, we focus on gradual change, as experienced, for example, by oceanic phytoplankton exposed to increases in surface temperature and acidity (e.g., Collins et al, 2014). This scenario can be addressed in the so-called moving-optimum model, which assumes that one or more quantitative characters are subject to stabilizing selection towards an optimal value that increases (or decreases) over time. If the optimum moves at constant speed, an evolving population will follow at a certain phenotypic distance or “lag”, whose size depends on the speed of environmental change, the strength of selection and the available genetic variation (reviewed in Kopp and Matuszewski, 2014). In a seminal study, Bürger and Lynch (1995, see also Lynch and Lande, 1993) derived a “critical rate of environmental change”, beyond which the lag becomes too large for the population to tolerate and extinction is inevitable. Their model (and subsequent extensions, e.g., Gomulkiewicz and Holt, 1995; Jones et al, 2004; Gomulkiewicz and Houle, 2009; Chevin et al, 2010) uses a quantitative-genetics approach, that is, it assumes that adaptation occurs from standing genetic variation, the trait under selection is determined by many loci with small individual effects, and genetic variance is constant (although the latter assumption can be relaxed in simulations).

In contrast, Kopp and Hermisson (2007, 2009a) focused on adaptation from new mutations, and investigated how the moving optimum affects the probability and time of fixation for alleles of both small and large effect. Following this work, Kopp and Hermisson (2009b) formulated an “adaptive-walk approximation”, which neglects fixation time and assumes that the population evolves via a stochastic jump process, where an “adaptive jump” (or “step”) occurs whenever a beneficial mutation arrives and escapes loss due to genetic drift (see the strong-selection-weak-mutation model introduced by Gillespie, 1983a,b). They showed that the characteristics of the adaptive walk (in particular, its step-size distribution) depend crucially on a composite parameter γ , which can be interpreted as the ratio of the speed of environmental change to the “adaptive potential” of the population. Small γ (slow change/high adaptive potential) corresponds to an “environmentally-limited regime”, in which the population stays close to the optimum and only small-effect mutations contribute to adaptation. In contrast, large γ defines a “genetically-limited regime”, in which the phenotypic

lag is large and the adaptive walk is shaped primarily by the distribution of incoming mutations.

Matuszewski et al (2014) extended this model to account for a multidimensional phenotype with universal pleiotropy (meaning that each mutation affects every trait). This amounts to a moving-optimum version of Fisher’s well-known “geometric model” (Fisher, 1930), which was originally introduced to argue for the pre-eminence of small mutations in adaptive evolution, but later has been shown to provide empirically accurate predictions under a wide range of conditions (e.g., Orr, 1998, 2005; Martin and Lenormand, 2006; Tenaillon, 2014). A key feature of Fisher’s model is a “cost of complexity”, since an increase in the number of phenotypic dimensions decreases the proportion of beneficial mutations and, hence, the rate of adaptation (Orr, 2000). Under a moving optimum, this translates into a larger phenotypic lag and an adaptive walk that proceeds via rarer (but larger) steps (Matuszewski et al, 2014).

The adaptive-walk approximation to the moving-optimum model also poses some interesting mathematical problems, which have been covered in depth by Nassar (2016) and Nassar and Pardoux (2017). Recently, Nassar and Pardoux (2018) developed a “small-jumps limit” to the adaptive-walk approximation, by assuming that infinitely small mutations arrive at an infinitely high rate (similar to the “canonical equation” of adaptive dynamics, e.g. Geritz et al, 1998; Champagnat et al, 2002; Champagnat, 2006; Boettiger et al, 2010). In this limit, the evolution of the phenotypic lag converges to an Ornstein-Uhlenbeck process around a long-term equilibrium, providing simple analytical predictions for its mean and variance. The aim of the present paper is to use this small-jumps limit as an approximation to adaptive walks with finite step sizes, as well as to more genetically explicit models of polygenic adaptation under a moving optimum. We show that the approximation works well in part of the genetically-limited regime. For the environmentally-limited regime, we obtain some scaling relations based on a different approximation from Matuszewski et al (2014). Finally, we exploit the above-mentioned Ornstein-Uhlenbeck process to gain some insight into the long-term extinction risk of populations in slowly-changing environments.

2 The model

Following Matuszewski et al (2014), our model is set up as follows: A population of constant size N is subject to Gaussian stabilizing selection towards a (d -dimensional) moving optimum that changes linearly with speed vector \mathbf{v} . That is, at time t , the phenotypic lag displayed by an individual with trait vector \mathbf{z} equals $\mathbf{x} = \mathbf{v}t - \mathbf{z}$, and the corresponding fitness is

$$W(\mathbf{x}) = \exp(-\mathbf{x}'\Sigma^{-1}\mathbf{x}), \quad (1)$$

where the positive definite matrix Σ describes the shape of the fitness landscape (and $'$ denotes transposition). Without loss of generality (Matuszewski et al, 2014), we set

$$\mathbf{v} = v\mathbf{e}_1 \text{ with } v > 0, \text{ and} \quad (2)$$

$$\Sigma = \sigma^2\mathbf{I}, \quad (3)$$

where e_1 is the unit vector $(1, 0, \dots)'$ and \mathbf{I} the identity matrix in \mathbb{R}^d . In other words, \mathbf{v} is a horizontal vector and Σ is isotropic. Note that, in the following, we will retain boldface notation for the lag \mathbf{x} even in the one-dimensional case ($d = 1$), whereas we will generally refer to the speed of environmental change as v (the only non-zero element of \mathbf{v}). We will also refer to σ^{-2} as a measure for the strength of **stabilizing** selection.

For the adaptive-walk approximation, the population is assumed to be monomorphic at all times (i.e., its state is completely characterized by x). Mutations arise at rate $\Theta/2 = N\mu$ (where μ is the *per-capita* mutation rate and $\Theta = 2N\mu$ is a standard population-genetic parameter), and their phenotypic effects α are drawn from a distribution $p(\alpha)$. In this paper $p(\alpha)$ always is a multivariate normal with mean $\mathbf{0}$ and positive definite covariance matrix \mathbf{M} , that is

$$p(\alpha) = \frac{1}{\sqrt{(2\pi)^d \det(\mathbf{M})}} \exp\left(-\frac{1}{2}\alpha' \mathbf{M}^{-1} \alpha\right), \quad (4)$$

even though some of the results by Nassar and Pardoux (2017, 2018) are valid under more general conditions. In the isotropic case, $\mathbf{M} = \omega^2 \mathbf{I}$, where ω^2 is the variance of mutational effects. For general \mathbf{M} , we define

$$\bar{\omega}^2 = \sqrt[d]{\det(\mathbf{M})}, \quad (5)$$

which is the geometric mean of the eigenvalues of \mathbf{M} and can be loosely interpreted as the average variance of mutational effects across phenotypic directions (see Matuszewski et al, 2014).

We neglect the possibility of fixation for deleterious mutations. Yet even beneficial mutations have a significant probability of being lost due to the effects of genetic drift while they are rare. For a mutation with effect α that arises in a population with phenotypic lag \mathbf{x} , **we assume that** its probability of fixation is

$$g(\mathbf{x}, \alpha) = \begin{cases} 1 - \exp(-2s(\mathbf{x}, \alpha)) & \text{if } s(\mathbf{x}, \alpha) > 0, \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where

$$s(\mathbf{x}, \alpha) = \frac{W(\mathbf{x} - \alpha)}{W(\mathbf{x})} - 1 \approx (2\mathbf{x} - \alpha)' \Sigma^{-1} \alpha \quad (7)$$

is the selection coefficient. Equation (6) is a good approximation to the fixation probability derived under a diffusion approximation (Malécot, 1952; Kimura, 1962), as long as the population size N is not too small. Note that Matuszewski et al (2014) used the even simpler approximation $g(\mathbf{x}, \alpha) \approx 2s(\mathbf{x}, \alpha)$ (Haldane, 1927; for more exact approximations of the fixation probability in changing environments, see Uecker and Hermisson, 2011; Peischl and Kirkpatrick, 2012). Once a mutation gets fixed, it is assumed to do so instantaneously, and the phenotypic lag \mathbf{x} of the population is updated accordingly. We call the resulting stochastic process \mathbf{X}_t an “adaptive walk”. Three example realizations are illustrated in Figure 1.

In a rigorous mathematical treatment of the above model, Nassar (2016) and Nassar and Pardoux (2017) have formalized the process \mathbf{X}_t – describing the evolution




Fig. 1 Three example realizations of an adaptive walk in one dimension ($d = 1$), showing the evolution of the lag X_t between the population phenotype \mathbf{z} and a linearly moving optimum $\mathbf{v}t$, for three different speeds of environmental change v . In (A) and (B), the process is recurrent, whereas in (C) it is transient. (A) corresponds to the environmentally-limited regime and (B) to the genetically-limited regime (see main text). Results are from simulations of equation (A1) (“adaptive-walk simulations”) with $\omega^2 = 1$, $\Theta = 1$ and $\sigma = 0.1$.

of the phenotypic lag via the quasi-instantaneous fixation of beneficial mutations – by means of a stochastic differential equation, which we here repeat in Appendix A (eq. A1). For the univariate case ($d = 1$), they show that \mathbf{X}_t is (Harris) recurrent if $v < m$ (see Fig. 1A,B) and transient if $v > m$ (Fig. 1C), where

$$m = \frac{\Theta}{2} \int_{\mathbb{R}_+} \alpha p(\alpha) d\alpha = \Theta \omega / \sqrt{8\pi} \quad (8)$$

is the rate of adaptation that is attained when every beneficial mutation goes to fixation. (The behavior in the limiting case $v = m$ depends on additional technical conditions.) For $d > 1$, we conjecture that the above criterion remains valid in the case of isotropic mutations ($\mathbf{M} = \omega^2 \mathbf{I}$), whereas in the presence of mutational correlations, simulations suggest that the value of v at the boundary to the transient case is less than m .

The aim of this paper is to better understand the behaviour of the process \mathbf{X}_t by using the small-jumps limit developed in Nassar and Pardoux (2018) as an approximation. To assess the performance of this approximation, we will compare its predictions to the results of two types of simulations, using methods developed in Matuszewski et al (2014).

First, “adaptive-walk simulations” (like those shown in Fig. 1) are performed by a straightforward implementation of equation (A1), that is, by repeatedly (i) drawing the waiting time for a new mutation from an exponential distribution with intensity $\Theta/2$; (ii) drawing the size of the mutation from its distribution $p(\alpha)$ (eq. 4); and (iii) accepting the mutation with its fixation probability $g(\mathbf{x}, \alpha)$ (eq. 6); note that, except for the simulation of extinction times, we used the approximate expression for the selection coefficient s given on the right-hand side of eq. (7), which causes fewer numerical problems at high values of v . The calculation of summary statistics for these simulations is described in Appendix B.

Second, individual-based simulations represent a much more realistic model, in which fixations are not instantaneous and multiple mutations may segregate simultaneously. Briefly, we model an initially monomorphic population of N haploid individuals, which are characterized by L genetic loci that additively determine the multivariate phenotype \mathbf{z} and, hence, the lag \mathbf{x} . Generations are discrete and non-overlapping. Each generation comprises (i) viability selection (individuals are removed with probability $1 - W(\mathbf{x})$, eq. 1), (ii) population regulation (random individuals are removed as long as the population size exceeds a carrying capacity K) and (iii) sexual reproduction (surviving individuals are randomly assigned to mating pairs **by sampling without replacement**, and each pair produces B offspring; offspring genotypes are derived from parental genotypes by taking into account recombination at rate r between adjacent loci [**where $r = 0.5$ means free recombination**] and mutation at per-locus rate μ). In all simulations reported here, we used $L = 10$, $K = 1000$, $\omega = 1$, $\mu = 5 \times 10^{-5}$ or $\mu = 5 \times 10^{-4}$ (yielding a population-wide mutation rate $\Theta = 2NL\mu = 1$ or $\Theta = 10$ and mutational variance $V_m = L\mu\omega^2 = 0.0005$ or $V_m = 0.005$), and $B = 2$ or $B = 8$; **in this particular model, the effective population size $N_e \approx 2BN/(2B - 1)$ (Bürger and Lynch, 1995)**; for more details, see Matuszewski et al (2014). Note that individual-based simulations will be applied only to a subset of parameter combinations investigated by means of adaptive-walk simulations.

3 Results

3.1 Evolution of the phenotypic lag

Nassar and Pardoux (2018) have studied a small-jumps limit of the process \mathbf{X}_t , which is obtained by multiplying the jump sizes by ε , dividing the rates by ε^2 , and then letting $\varepsilon \rightarrow 0$ (for details, see Appendix A). In particular, they show that the rescaled process \mathbf{X}_t^ε (eq. A2) converges in probability towards a deterministic solution $\bar{\mathbf{X}}_t$, given by the differential equation

$$\frac{d\bar{\mathbf{X}}_t}{dt} = \mathbf{v} - \sigma^{-2} \Theta \mathbf{M} \bar{\mathbf{X}}_t. \quad (9)$$

For $\mathbf{x}_0 = \mathbf{0}$, its solution is

$$\bar{\mathbf{X}}_t = (1 - \exp(-\sigma^{-2} \Theta \mathbf{M} t)) \frac{\mathbf{M}^{-1} \mathbf{v}}{\Theta \sigma^{-2}}, \quad (10)$$

which converges exponentially to the equilibrium value

$$\bar{\mathbf{X}}_t \xrightarrow{t \rightarrow \infty} \bar{\mathbf{X}}_\infty = \frac{\mathbf{M}^{-1} \mathbf{v}}{\Theta \sigma^{-2}}. \quad (11)$$

Using (2) and (5), equation (11) can be rewritten as

$$\bar{\mathbf{X}}_\infty = \gamma \bar{\omega} \left(\frac{\mathbf{M}}{\bar{\omega}^2} \right)^{-1} e_1, \quad (12)$$

where

$$\gamma = \frac{v/\bar{\omega}}{\Theta(\sigma/\bar{\omega})^{-2}} \quad (13)$$

is the scaled rate of environmental change defined in Matuszewski et al (2014), whose denominator can be interpreted as the ‘‘adaptive potential’’ of the population (see Kopp and Hermisson, 2009b). In the univariate case, $\bar{\omega}^2 = \mathbf{M} = \omega^2$, and the adaptive potential is equal to the constant factor in the second term on the right-hand side of equation (9), which describes phenotypic change in the population due to mutation and selection. Furthermore, in this case, $\bar{\mathbf{X}}_\infty/\omega = \gamma$, that is, the equilibrium mean lag, when measured in units of the typical size of mutations, is simply given by γ . In the multivariate case, the additional term $(\mathbf{M}/\bar{\omega}^2)^{-1} e_1$ corresponds to the first column of the inverse of the scaled mutation matrix. Its entries are related to the partial correlation coefficients between the effects of mutations on trait 1 (whose optimum value is directly affected by v) and each of the other traits (i.e., the partial correlation between trait 1 and trait i is given by $-m_{1i}^{-1}/\sqrt{m_{ii}^{-1} m_{11}^{-1}}$, where the m 's are the elements of the matrix \mathbf{M}). Thus, the equilibrium mean lag depends only on γ and the structure of mutational correlations in the direction of the moving optimum. We note that an alternative way of rewriting equation (11) is

$$\bar{\mathbf{X}}_\infty = \gamma_1 \omega \left(\frac{\mathbf{M}}{\omega^2} \right)^{-1} e_1, \quad (14)$$

where ω^2 the variance of new mutations in the direction of the optimum (i.e., the first entry of the matrix \mathbf{M}) and γ_1 is equal to γ in the one-dimensional case. The difference to equation (12) is that γ_1 does not capture the genetic constraints imposed by mutational correlations (as reflected in the quantity $\bar{\omega}^2$, eq. 5).

To calculate the variance of the lag in the small-jumps limit, Nassar and Pardoux (2018) consider the process

$$\mathbf{U}_t^\varepsilon = \frac{\mathbf{X}_t^\varepsilon - \bar{\mathbf{X}}_t}{\sqrt{\varepsilon}}. \quad (15)$$

They show that, for $\varepsilon \rightarrow 0$, this process converges to an Ornstein-Uhlenbeck process

$$d\mathbf{U}_t = -\sigma^{-2}\Theta\mathbf{M}\mathbf{U}_t dt + \Lambda^{\frac{1}{2}}(\bar{\mathbf{X}}_t)d\mathbf{B}_t. \quad (16)$$

The first term on the right-hand side of equation (16) describes the tendency of the process to return to its mean. It is equal to the mutation-selection term in equation (9) (and, for $d = 1$, also to the adaptive-potential term in the denominator of γ ; eq. 13). The second term describes the tendency of the process to fluctuate around the mean, where \mathbf{B}_t is a d -dimensional standard Brownian motion and the matrix $\Lambda(\bar{\mathbf{X}}_t)$ is the infinitesimal variance of the process (defined in Appendix C, eq. C1). An explicit expression for this term can be obtained only for the special case of isotropic mutations (i.e., $\mathbf{M} = \omega^2\mathbf{I}$) in the limit of large time (when $\bar{\mathbf{X}}_t$ converges to $\bar{\mathbf{X}}_\infty$). For this case, we show in Appendix C that the infinitesimal variance is $4v\omega/\sqrt{2\pi}$ in the direction of the optimum, and half this value in all other directions (covariances are zero). Thus, the infinitesimal variance reflects the factors that drive the process away from the mean, that is, the rate of environmental change and the typical size of mutations (and, hence, jumps). The variance-covariance matrix of the process \mathbf{U}_t is given by its second moment, which for the limit $t \rightarrow \infty$, we denote by $\bar{\mathbf{S}}^2$ (eq. C4). In the isotropic case, it is given by equation (C14), which shows that the **asymptotic** variance around the mean lag is

$$\bar{S}_1^2 = \frac{2v}{\sqrt{2\pi}\Theta\sigma^{-2}\omega} = \frac{2\omega^2\gamma}{\sqrt{2\pi}} \quad (17a)$$

in the direction of the optimum and

$$\bar{S}_{i>1}^2 = \frac{v}{\sqrt{2\pi}\Theta\sigma^{-2}\omega} = \frac{\omega^2\gamma}{\sqrt{2\pi}} \quad (17b)$$

in all other directions, with all covariances equal to zero. Note that the standard deviation of the lag in the direction of the optimum is very close to its mean: $\bar{S}_1 \approx 0.9\bar{X}_1$ (eq. 12, 17a; here and in the following, we abusively denote by \bar{X}_i the i 'th element of the vector $\bar{\mathbf{X}}_\infty$, and by \bar{S}_i^2 the i 'th diagonal element of the matrix $\bar{\mathbf{S}}^2$).

Since \mathbf{U}_t is an Ornstein-Uhlenbeck process, its stationary distribution is Gaussian with mean $\mathbf{0}$ and covariance matrix $\bar{\mathbf{S}}^2$. Similarly, for $\varepsilon \ll 1$ and large t , the process $\mathbf{X}_t^\varepsilon \approx \bar{\mathbf{X}}_t + \sqrt{\varepsilon}\mathbf{U}_t$, and hence $\mathbf{X}_t^\varepsilon \sim \mathcal{N}(\bar{\mathbf{X}}_\infty, \varepsilon\bar{\mathbf{S}}^2)$. We note that in contrast to the original process \mathbf{X}_t , the limiting process \mathbf{X}_t^ε is symmetric and no longer reflects the inherent asymmetry of the original model (where the first element of \mathbf{X}_t increases gradually due to environmental change and decreases suddenly due to fixations). For the original process \mathbf{X}_t with finite jumps, equations (12) and (17) are approximate

predictions for the long-term mean and variance of the phenotypic lag \mathbf{X} if we abusively approximate \mathbf{X}_t by $\bar{\mathbf{X}}_t + \mathbf{U}_t$ (i.e., set $\varepsilon = 1$). It is this approximation that we investigate in the following.

Figure 2 compares the predictions from the small-jumps limit to results from adaptive-walk simulations, for the case of a single evolving trait ($d = 1$). As can be seen, the predictions from the small-jumps limit are fairly accurate if $\bar{\mathbf{X}}_\infty/\omega = \gamma \gtrsim 1$ (and very good for $\gamma \gtrsim 10$), provided v is not too close to m . In other words, the approximation is good if the mean lag is larger than the size of a typical mutation (such that adaptive jumps are small relative to the lag, see Fig. 1B), but the system is not too close to the boundary of the transient case. For $\gamma \lesssim 1$ or $v \rightarrow m$, in contrast, equations (12) and (17) underestimate both the mean size of the lag and its variance, but the reasons are different in the two cases. For $v \rightarrow m$, the approximation does not capture the divergence of the phenotypic lag as the process approaches transience. The reason is that equation (9) assumes weak selection, and in particular, that the fixation probability $g(\mathbf{x}, \alpha) \approx 2s(\mathbf{x}, \alpha)$, whereas the real fixation probability of finitely-sized mutations is lower (see eq. 6) and saturates at 1 as $\mathbf{X}_t \rightarrow \infty$. For $\gamma \lesssim 1$, the small-jumps approximation is invalid because adaptive jumps are large relative to the mean lag (and often overshoot the optimum, leading to $\mathbf{X}_t < 0$, see Fig. 1A). Indeed, simulations show that the mean lag is significantly larger than predicted by equation (12). The reason is that, for $\gamma \ll 1$, adaptive jumps are relatively rare (because only few mutations are beneficial), and the lag will continue to increase until a successful mutation arrives. Figure 2C shows, in addition, that for small γ the variance of the lag converges to the square of the mean, such that the coefficient of variation is close to 1 (Fig. 2D). Finally, the results from adaptive-walk simulations are in very close agreement with those from individual-based simulations (Fig. S1, S2), at least as long as recombination is high and/or the population-wide mutation rate is not too large. In contrast, for $\Theta \gg 1$ and recombination rate $r \rightarrow 0$ (such as in asexual organisms with large population sizes), the lag increases, most likely because co-segregating beneficial mutations compete for fixation.

The above results remain largely valid also in the multivariate case (see Fig. S3 for the case of $d = 4$ traits). In particular, with isotropic mutations ($\mathbf{M} = \omega^2 \mathbf{I}$), the mean lag in the direction of the optimum (\bar{X}_1) is almost identical to the one in the univariate case (except for high values of the scaled selection strength $(\sigma/\omega)^{-2} = \sigma^{-2}\omega^2$, where the lag is slightly increased; Fig. S3A). The lag in all other directions fluctuates around zero (not shown), with a variance that is predicted to be half as large as the one in the direction of the optimum (eq. 17). However, this prediction holds true only for a limited range of parameter values (that is, those for which the mean lag fits the prediction very well; Fig. S3B-D). For $v \rightarrow m$, the variance in the direction of the optimum explodes (Fig. S3B), whereas for small v (i.e., small γ), the ratio of variances $\text{Var}(X_1)/\text{Var}(X_2)$ converges to a value that is close to $\sqrt{2}$ (rather than 2; Fig. S3D) – a finding for which we have no analytical explanation.

When mutational effects on different traits are correlated (parameter ρ in Fig. S4), the key prediction from the adaptive-walk approximation is a bias in those traits whose optimum is not affected by environmental change. For example, in Figure S4B, a positive mutational correlation between traits 1 and 2 creates a positive mean lag in the direction of trait 2 (while having relatively little effect on trait 1, unless ρ is close




Fig. 2 The long-term steady state of the normalized phenotypic lag \bar{X}/ω in the one-dimensional case ($d = 1$). The Figure compares results from adaptive-walk simulations (averaged over 10^6 adaptive steps) to the predictions from the small-jumps limit. (A) The mean phenotypic lag \bar{X}/ω (relative to the mutational standard deviation ω), as a function of the scaled rate of environmental change $v/(\omega\Theta)$ for various values of the scaled strength of selection $\sigma^{-2}\omega^2$. The dotted lines represent the analytical prediction from eq. (12) (from top to bottom for the same values of $\sigma^{-2}\omega^2$ as the simulation data). To reduce the number of parameters, we focus (here and in the following figures) on the variable \bar{X}_t/ω , which depends only on the scaled rate of environmental change $v/(\omega\Theta)$ and the scaled strength of selection $\sigma^{-2}\omega^2$ (this can be shown by measuring the phenotype in units of ω and time in units of Θ^{-1}). (B) The standard deviation of the scaled lag. The dotted lines represent the analytical predictions from eq. (17). (C) The standard deviation as a function of the mean lag. The dotted line marks the analytical prediction according to equations (12) and (17). The solid line marks the main diagonal. (D) The coefficient of variation (standard deviation over mean) of the lag, with the dotted lines again marking the analytical predictions. The vertical line in (A), (B) and (D) marks the boundary between the recurrent and transient cases ($v = m = \Theta\omega/\sqrt{8\pi} \Leftrightarrow v/(\omega\Theta) = 1/\sqrt{8\pi}$).

to 1). This is the “flying-kite effect” described by Jones et al (2004) and Matuszewski et al (2014), **the idea being that moving-optimum selection on trait 1 “pulls” the population and makes it “rise” also in the direction of trait 2.** Adaptive-walk simulations show that it is fully operational only if selection is rather weak and environmental change is fast (such that $\gamma_1 \gtrsim 10$, Fig. S4B, S5B). In these cases, the variance of trait 2 approaches that of trait 1 (Fig. S4C-E, S5C-E), and the lags of the two traits are correlated (Fig. S4F, S5F). Note that both the ratio of variances and the correlation depend only on the strength of mutational correlations (see eq. C9).

3.2 Population survival and extinction

A question of considerable interest for conservation is how much environmental change a population can compensate by adaptive evolution without going extinct (e.g., Bürger and Lynch, 1995; Kopp and Matuszewski, 2014). While our model does not include explicit population dynamics – and, hence, cannot be used to study extinction directly – a simple approach is to assume that extinction risk is strongly elevated once the phenotypic lag exceeds a critical threshold X_{crit} . For example, if individuals that survive selection have, on average, $B > 1$ offspring, population size will start declining once the population (mean) fitness $W(\mathbf{X}_t)$ (eq. 1) drops below $1/B$, that is, once the total size of the lag $\|\mathbf{X}_t\| = \sqrt{\mathbf{X}_t^T \mathbf{X}_t} > X_{\text{crit}}$, where

$$X_{\text{crit}} = \sqrt{\sigma^2 \ln B}. \quad (18)$$

Extinction will usually follow rapidly (Bürger and Lynch, 1995), even though “evolutionary rescue” (Gonzalez et al, 2013) is still possible by the timely arrival and fixation of a beneficial mutation (Orr and Unckless, 2008; Uecker and Hermisson, 2011; Orr and Unckless, 2014). If the process \mathbf{X}_t is transient, the lag will reach the critical size after at most $X_{\text{crit}}/(m - v)$ generations. In contrast, if the process is recurrent, it may spend most of its time below X_{crit} . Note, however, that even in this case, eventual extinction is certain, because the lag has a non-zero probability of reaching any arbitrarily large size. The key question is, therefore: For how long will the size of the lag typically remain below X_{crit} ? Neglecting the possibility of evolutionary rescue, we will call this the “time to extinction” and denote it by T_e .

Obviously, the time to extinction strongly depends on the rate of environmental change (Fig. 3). Typically, extinction risk is negligible if the mean lag in the direction of the optimum is less than one mutational standard deviation ($\gamma < 1$, corresponding to $v/(\omega\Theta) < \sigma^{-2}\omega^2$ in Fig. 3), unless the fitness effect of a single mutation is very high ($\sigma^{-2}\omega^2$ close to 1, Fig. 3E, F). For $\gamma > 1$, we can gain additional insights from the small-jumps limit.

First, extinction risk should certainly be high if the mean size of the lag at equilibrium, $\|\bar{\mathbf{X}}_\infty\|$, is greater than X_{crit} . It follows immediately from equation (12) that this is the case if v exceeds the “critical rate of environmental change” (Bürger and Lynch, 1995)

$$v_{\text{crit}} = \frac{\Theta \sqrt{\sigma^{-2} \ln B}}{\|\mathbf{M}^{-1} \mathbf{e}_1\|}. \quad (19)$$

Note that, in the absence of mutational correlations, this expression simplifies to $v_{\text{crit}} = X_{\text{crit}}\Theta\sigma^{-2}\omega^2$, that is, the product of the critical lag and the ‘‘adaptive potential’’.

If $v > v_{\text{crit}}$, the time to extinction might be estimated by setting the right-hand side of equation (10) equal to $X_{\text{crit}}e_1$ and solving for t (see Bürger and Lynch, 1995). However, the solution diverges as $v \rightarrow v_{\text{crit}}$, and generally overestimates the real time to extinction, because it neglects stochastic fluctuations. Indeed, simulations show that, for $v > v_{\text{crit}}$, the time to extinction is typically of order X_{crit}/v , that is, it is only slightly prolonged by the fixation of beneficial mutations (Fig. 3; see also Gomulkiewicz et al, 2017). If, on the contrary, $v < v_{\text{crit}}$, we can use the fact that the process converges to an Ornstein-Uhlenbeck process around $\bar{\mathbf{X}}_\infty$ (eq. 16). The time to extinction can then be decomposed into the time for \mathbf{X}_t to go from 0 to $\bar{\mathbf{X}}_\infty$ and the additional time from $\bar{\mathbf{X}}_\infty$ to the boundary of the region defined by $\|\mathbf{X}_t\| \leq X_{\text{crit}}$. Unless v is close to v_{crit} , the first part will be much shorter than the second and can be approximated (and slightly underestimated) by $\bar{\mathbf{X}}_\infty/v$. The second part is highly stochastic and can be approximated by the exit time T_f of the process $\bar{\mathbf{X}}_\infty + \mathbf{U}_t$ from the region with $\|\mathbf{X}_t\| \leq X_{\text{crit}}$ when starting at $\bar{\mathbf{X}}_\infty$. Thus, in summary,

$$T_e \approx \begin{cases} \bar{\mathbf{X}}_\infty/v + T_f & \text{if } v \leq v_{\text{crit}}, \\ X_{\text{crit}}/v & \text{if } v > v_{\text{crit}}. \end{cases} \quad (20)$$

To the best of our knowledge, analytical results are available only for $d = 1$, when the exit time corresponds to the first-passage time of the one-dimensional Ornstein-Uhlenbeck process by the point X_{crit} . Following Thomas (1975) and Ricciardi and Sato (1988, see also Finch, 2004), T_f has mean

$$\begin{aligned} \mathbb{E}(T_f) &= \frac{\sqrt{\pi/2}}{\delta} \int_0^{\tilde{X}_{\text{crit}}} \left(1 + \operatorname{erf}\left(\frac{t}{\sqrt{2}}\right)\right) \exp\left(\frac{t^2}{2}\right) dt \\ &= \frac{1}{2\delta} \sum_{k=1}^{\infty} \frac{(\sqrt{2}\tilde{X}_{\text{crit}})^k}{k!} \Gamma\left(\frac{k}{2}\right) \end{aligned} \quad (21)$$

and variance

$$\begin{aligned} \operatorname{Var}(T_f) &= \frac{\sqrt{2\pi}}{\delta^2} \int_0^{\tilde{X}_{\text{crit}}} \int_{-\infty}^t \int_s^{\tilde{X}_{\text{crit}}} \left(1 + \operatorname{erf}\left(\frac{r}{\sqrt{2}}\right)\right) \exp\left(\frac{r^2 + t^2 - s^2}{2}\right) dr ds dt - \mathbb{E}(T_f)^2 \\ &= \mathbb{E}(T_f)^2 - \frac{1}{2\delta^2} \sum_{k=1}^{\infty} \frac{(\sqrt{2}\tilde{X}_{\text{crit}})^k}{k!} \Gamma\left(\frac{k}{2}\right) \left(\phi\left(\frac{k}{2}\right) - \phi(1)\right), \end{aligned} \quad (22)$$

where $\delta = \Theta\sigma^{-2}\omega^2$ is the denominator of γ (‘‘adaptive potential’’), $\tilde{X}_{\text{crit}} = (X_{\text{crit}} - \bar{X}_\infty)/S$ is the normalized mean lag, $\phi(\cdot)$ is the digamma function, and $\operatorname{erf}(x) = 2/\sqrt{\pi} \int_0^x e^{-t^2} dt$ is the Gauss Error Function.

Figure 3 compares these predictions to results from adaptive-walk simulations. As long as $\sigma^{-2}\omega^2 \leq 0.01$, the mean time to extinction is well approximated by (20) and (21), even though it is slightly underestimated in the region where $v \approx v_{\text{crit}}$ (because we neglect adaptive steps before X_t reaches \bar{X}_∞ or X_{crit} , respectively) and for small values of $v/(\omega\Theta)$ (probably as a result of the finite intervals between jumps

in our simulations). Similarly, the variance is well approximated by (20) and (22), as long as $v < v_{\text{crit}}$ (whereas we lack a prediction for the variance in the opposite case $v > v_{\text{crit}}$). Note that, in Figure 3, the predictions were improved (in particular for $\sigma^{-2}\omega^2 = 0.01$) by replacing the mean lag \bar{X}_{∞} according to equation (12) by the value found in simulations (see Fig. 2). Also note that, for small $v/(\omega\Theta)$, the mean and standard deviation of the time to extinction are nearly identical, suggesting that T_e follows an exponential distribution. Finally, for $\sigma^{-2}\omega^2 > 0.01$, the approximation (20) breaks down, because $\gamma < 1$ even for large $v/(\omega\Theta)$ and, hence, the small-jumps approximation does not apply. Extinction is nevertheless fast, because even a small deviation from the optimum (relative to the mutational standard deviation) has dramatic fitness consequences.

The above results are largely confirmed by individual-based simulations (Fig. S6). In particular, for $\sigma^{-2}\omega^2 \leq 0.001$, results from individual-based simulations are very close to those from adaptive-walk simulations. In contrast, for $\sigma^{-2}\omega^2 \geq 0.01$ (and most clearly for $\sigma^{-2}\omega^2 \geq 0.1$), extinction takes longer in individual-based simulations than in adaptive-walk simulations. This might seem surprising, since the assumption of instantaneous fixations in the adaptive-walk simulations should be favorable for adaptation. However, rapid fixation of a single mutation might also prevent the establishment of other mutations that arise slightly later, thus reducing the total number of mutations that contribute to adaptation. Individual-based simulations further show that for $\sigma^{-2}\omega^2 \geq 0.1$, “real” extinction times are substantially larger than the times needed for mean fitness to drop below $1/B$, showing that **(temporary)** evolutionary rescue is common in this regime.

Finally, some simple scaling relations can be obtained for $v < v_{\text{crit}}$ by applying a further approximation to the first-passage time T_f . Indeed, for $\tilde{X}_{\text{crit}} \gtrsim 1.5$, the mean $\mathbb{E}(T_f)$ (eq. 21) is well approximated by

$$\mathbb{E}(T_f) \approx \frac{\sqrt{2\pi}}{\delta\tilde{X}_{\text{crit}}} \exp\left(\frac{\tilde{X}_{\text{crit}}^2}{2}\right), \quad (23)$$

(Ricciardi and Sato, 1988) and the variance $\text{Var}(T_f)$ (eq. 22) by the square of this value (showing again that, for small v , T_e converges to an exponential distribution). For most values of \tilde{X}_{crit} , the **approximation** (23) is dominated by the exponential term

$$\begin{aligned} \exp\left(\frac{\tilde{X}_{\text{crit}}^2}{2}\right) &= \exp\left(\frac{(X_{\text{crit}} - \bar{X}_{\infty})^2}{2S^2}\right) \\ &= \exp\left[\frac{\sqrt{2\pi}\Theta\omega\sigma^{-2}}{4v} \left(\sqrt{\frac{\ln B}{\sigma^{-2}}} - \frac{v}{\Theta\omega^2\sigma^{-2}}\right)^2\right] \\ &= \exp\left[\frac{\sqrt{2\pi}}{4\gamma} \left(\frac{X_{\text{crit}}}{\omega} - \gamma\right)^2\right]. \end{aligned} \quad (24)$$

In particular, as long as $v \ll v_{\text{crit}}$, the difference $X_{\text{crit}} - \bar{X}_{\infty}$ (the squared term in the exponent) depends only weakly on v , Θ and ω . To a first approximation, therefore, the mean time to extinction in this case scales with $\exp(\Theta\omega/v)$ (second line of equation 24). In contrast, the dependence on σ^{-2} is more complex, since σ^{-2} affects both

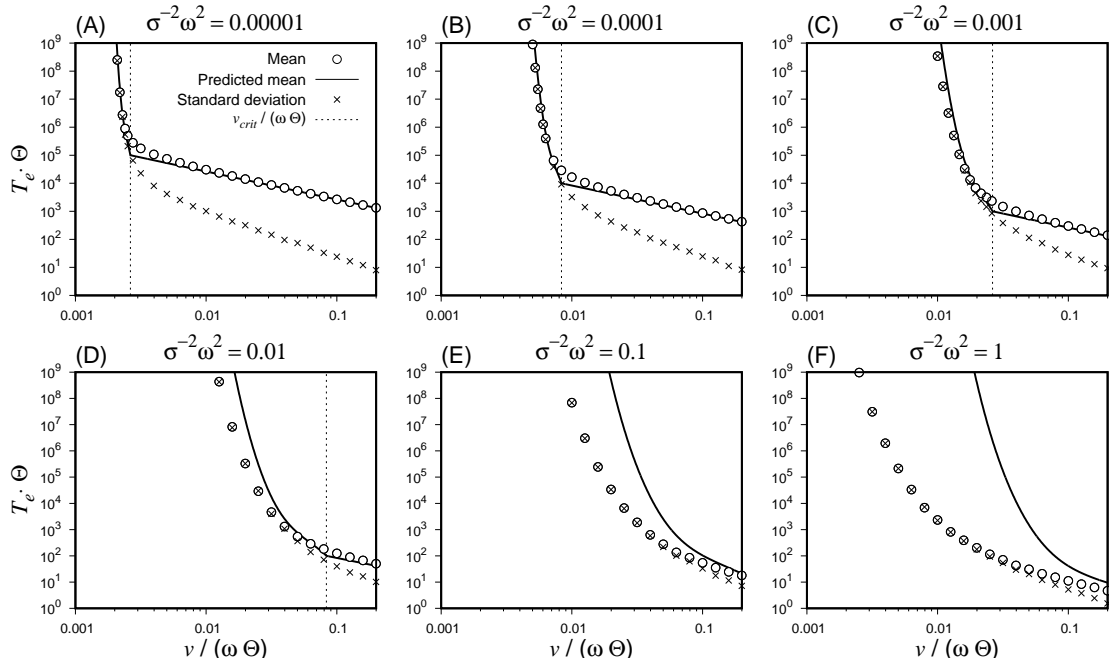


Fig. 3 The time T_e (relative to the mean interval between new mutations, Θ^{-1}) until the population mean fitness $W(\mathbf{X}_t)$ drops below $1/B = 1/2$ for the first time (“time to extinction”), as a function of the scaled rate of environmental change $v/(\omega\Theta)$ for various values of the scaled strength of selection $\sigma^{-2}\omega^2$, in the one-dimensional model ($d = 1$). Circles show the mean and crosses the standard deviation from 1000 replicated adaptive-walk simulations. The solid line shows the analytical prediction for the mean, (eq. 20 and 21, where $\bar{\mathbf{X}}_\infty$ from eq. 12 has been replaced by the mean lag from simulations, see Fig. 2). The prediction for the standard deviation (square root of eq. 22) is not shown, because it is very close to the prediction for the mean, except for $v \rightarrow v_{\text{crit}}$ (eq. 19, vertical dotted line), where it converges to 0 (because our approximation eq. 20 considers stochasticity only after the process has reached its long-term mean $\bar{\mathbf{X}}_\infty$, and at this value, the population is already extinct).

X_{crit} and $\bar{\mathbf{X}}_\infty$, leading to a non-monotonic relation if v is intermediate (Fig. 4). The reason is that, for both low and high σ^{-2} , the mean time to extinction approaches the minimum X_{crit}/v . For small σ^{-2} , selection is so weak that, even though X_{crit} is large, almost no mutations get fixed. For large σ^{-2} , X_{crit} is so small that the population has a high probability of going extinct before the first fixation can occur. In contrast, for intermediate σ^{-2} , X_{crit} is sufficiently large and selection sufficiently efficient to prevent population extinction over long periods due to the fixation of beneficial mutations. Finally, the last line of equation (24) shows that if $X_{\text{crit}} \ll \bar{\mathbf{X}}_\infty$ is treated as a constant and is measured relative to ω then the time to reach this value scales with $\exp(\gamma^{-1})$.

As mentioned above, no analytical results are available for extinction times in the multidimensional model ($d > 1$). Not only do we lack an approximation for exit times in the multidimensional Ornstein-Uhlenbeck process, but the small-jumps limit also fails to capture a crucial effect of multidimensionality, the so-called “cost of complexity” (Orr, 2000). As first noted by Fisher (1930), increasing dimensional-




Fig. 4 The time T_e (relative to the mean interval between new mutations, Θ^{-1}) until the population mean fitness $W(\mathbf{X}_t)$ drops below $1/B = 1/2$ for the first time (“time to extinction”), as a function of the scaled strength of selection $\sigma^{-2}\omega^2$ for an intermediate rate of environmental change $v/(\omega\Theta) = 0.02$. The dashed line marks the minimal time X_{crit}/v . Note that for $\sigma^{-2}\omega^2 < 0.001$, the rate of environmental change $v > v_{\text{crit}}$. For further details, see Fig. 3.

ity causes a decrease in the proportion of beneficial mutations, making adaptation more difficult, but this effect vanishes as mutational effect sizes tend to zero (where there are always 50% beneficial mutations). To study the effect of multidimensionality with finite mutations, we therefore must resort to simulations. Figure S7 shows results for uncorrelated mutations. Simulated extinction times are still very close to the theoretical predictions from the one-dimensional model as long as selection is weak ($\sigma^{-2}\omega^2 \leq 10^{-4}$). In contrast, extinction times are somewhat shorter for moderate selection strengths, $\sigma^{-2}\omega^2 = 10^{-3}$ or 10^{-2} , and they are strongly reduced under strong selection $\sigma^{-2}\omega^2 \geq 10^{-1}$, where they are often close to the theoretical minimum X_{crit}/v . These adaptive-walk results are again confirmed by individual-based simulations (colored symbols in Fig. S7). In summary, the effects of multidimensionality are strongest under strong selection (see Fig. S8).

When mutational effects are correlated (and the dominant eigenvector of the \mathbf{M} -matrix does not point in the direction of optimum), extinction occurs faster than without correlation (Fig. S9), especially when selection is weak ($\sigma^{-2}\omega^2$ small). The reason is that correlations reduce the amount of independent genetic variation that is created in the direction of the optimum, and this limits the response to selection when the lag is large, whereas it is less important when the lag is small (i.e., under strong selection).

4 Discussion

We have studied a stochastic process describing an “adaptive-walk” of an evolving population following a moving phenotypic optimum via the quasi-instantaneous fixation of beneficial mutations. In particular, we used a recently developed small-jumps limit (Nassar and Pardoux, 2018), which allowed us to derive analytical approxima-

tions for the long-term mean and variance of the phenotypic lag in multiple dimensions, and for the mean and variance of the time to extinction in the univariate case. Even though valid only in part of parameter space, these approximations are highly instructive and allow us to place our results in the context of previous studies of the moving optimum model. They also provide new insight into the long-term extinction risk of populations in slowly-changing environments.

Evolution of the phenotypic lag Our analytical results show that the evolution of the phenotypic lag depends strongly on the composite parameter γ (eq. 13; see Kopp and Hermisson, 2009b; Matuszewski et al, 2014), which scales the rate of environmental change relative to the “adaptive potential” of the population (see eq. 9, 16). Indeed, in the small-jumps limit, γ is equal to the long-term mean of the lag, \bar{X}_∞ , when the latter is measured in units of the mutational standard deviation ω (and mutational correlations are absent). Comparison to simulations (Fig. 2, S3) shows that the predictions from the small-jumps limit perform reasonably well as long as $\gamma \gtrsim 1$, that is, as long as the mean lag exceeds the effect size of a typical mutation (and, in addition, environmental change is not too close to the boundary of the transient case, in which continued adaptation is impossible). This observation conforms nicely to the classification introduced in Kopp and Hermisson (2009b), who stated that for $\gamma \ll 1$, the adaptive process is “environmentally-limited”, whereas for $\gamma \gg 1$, it is “genetically-limited”. The idea is that for $\gamma \ll 1$, the mean lag is so small ($\bar{X}_\infty \ll \omega$) that large mutations are usually selected against (as they would overshoot the optimum by too much), and hence, which mutations are fixed depends primarily on the rate of environmental change. In contrast, for $\gamma \gg 1$, the mean lag is large ($\bar{X}_\infty \gg \omega$), so most mutations that have positive effects in the direction of the optimum ($\alpha_1 > 0$) are positively selected and their rate of fixation depends primarily on genetic factors (i.e., their rate of appearance; for a discussion of the boundary between these two regimes, see Supporting Information 3 in Matuszewski et al, 2014). It is, thus, in part of the genetically-limited regime (i.e., the part with intermediate v) that the small-jumps approximation is most accurate. More precisely, the small-jumps approximation requires that $\sigma^{-2}\bar{\omega}^2$ is small and $v/(\bar{\omega}\Theta)$ has intermediate values. This means either weak selection (small σ^{-2}) or small mutations (small $\bar{\omega}^2$) in combination with slow environmental change or a high mutation rate. We note that realistic values of $\sigma^{-2}\bar{\omega}^2$ are probably in the middle of the range we investigated. For example, Bürger and Lynch (1995), based on published empirical estimates, assumed $\bar{\omega}^2 = 0.05$ and considered values of σ^{-2} between 0.005 and 0.5, implying $\sigma^{-2}\bar{\omega}^2$ between 0.00025 and 0.025 (or 0.0125 if selection is “diluted” by non-genetic phenotypic variation).

In the environmentally-limited regime ($\gamma \ll 1$), the small-jumps approximation fails, because most mutations are large relative to the phenotypic lag. Indeed, Kopp and Hermisson (2009b) and Matuszewski et al (2014) proposed a different approximation for this case: Since most large mutations are selected against, the successful mutations resulting in adaptive jumps come from the center of the distribution of new mutations, which can be approximated by a uniform distribution with appropriate density. Unfortunately, this approximation did not allow us to obtain results about the long-term behavior of the lag. However, some scaling relations can be obtained by focusing on the *first* jump of the adaptive walk, using results from Matuszewski et al

	Genetically-limited regime (unless close to transient case)	Environmentally-limited regime
Parameter range	$\gamma \gtrsim 1$	$\gamma \lesssim 1$
With isotropic mutation		
Mean of X_1/ω	equal to γ	proportional to $\gamma^{1/(d+3)}$
Variance of X_1/ω	proportional to γ	proportional to $\gamma^{2/(d+3)}$
CV of X_1/ω	proportional to $1/\sqrt{\gamma}$	close to 1
Ratio of variances X_1 to X_2	equal to 2	close to $\sqrt{2}$
With correlated mutations		
Flying kite effect	present	weak or absent
Correlation within lag	present	weak or absent

Table 1 The long-term behavior of the phenotypic lag in the genetically- and environmentally-limited regime, respectively. X_i stands for the lag in the direction of trait i . CV is the coefficient of variation (standard deviation over mean).

(2014). In particular, immediately before and after the first jump, the mean lag in the direction of the optimum is proportional to $\gamma^{1/(d+3)}$ and its variance to $\gamma^{2/(d+3)}$, which explains the initial slope of the curves for $\sigma^{-2}\omega^2 = 1$ in Figures 2A,B and S3A,B (for further details, see Appendix D). A comparison between the environmentally- and genetically-limited regimes is given in Table 1.

Time to population extinction We then applied our results to investigate the time until the population reaches a dangerously high level of maladaptation, entailing a significant risk of extinction. Following Lynch and Lande (1993) and Bürger and Lynch (1995), we calculated a “critical rate of environmental change” v_{crit} , beyond which the equilibrium phenotypic lag becomes too large for the population to tolerate. The result is very simple: In the absence of mutational correlations, the critical rate of environmental change equals the critical phenotypic lag times the adaptive potential of the population.

However, even below this critical rate, the population will ultimately go extinct due to stochastic fluctuations. So far, the time until this event had been studied only by simulations (e.g., Bürger and Lynch, 1995). Here, we used the fact that, in the small-jumps limit, the adaptive walk converges to an Ornstein-Uhlenbeck process around the expected mean lag. We then used known results on the first-passage time of this process to derive analytical predictions for the time to extinction, at least in the univariate case. A simple approximation yields that the time to extinction is roughly proportional to $\exp(\Theta\omega/v)$, that is, it is exponential in the mutation rate, mutational standard deviation and the inverse of the speed of environmental change (e.g., Fig. 3). In contrast, the dependence on the strength of stabilizing selection is more complex and non-monotonic, since this parameter influences not only the adaptive potential but also the critical phenotypic lag (Fig. 4). Finally, the distribution of the time to extinction is approximately exponential if v is sufficiently below v_{crit} . The best fit is reached for slowly changing environments, where extinction times are on a palaeon-

tological timescale. Indeed, there is empirical support for an exponential distribution of species life times (Stenseth and Smith, 1984; Pigolotti et al, 2005). The traditional explanation goes back to van Valen (Van Valen, 1973), stating that Red Queen dynamics prevent any one species from attaining a long-term fitness advantage, such that extinction is mostly based on demographic stochasticity. Our model provides an alternative explanation: species go extinct due to the eventual failure (even for large populations) of catching up with an ever-changing (biotic or abiotic) environment. This should remain true even if environmental change is non-linear (e.g., sinusoidal on a large timescale).

Effect of multidimensionality An obvious draw-back of our small-jumps limit is its inability to capture the “cost of complexity” (Orr, 2000), since the latter vanishes as the phenotypic effect size of mutations tends to zero: independently of the number of traits under selection, there will always be 50% beneficial mutations. As a consequence, in the isotropic model, the predicted mean lag in the direction of the optimum (as well as its variance) are independent of the number of traits d . In contrast, adaptive-walk simulations with finite mutations show that both the mean lag and its variance increase with d (compare Fig. 2 and S3), and this effect is strongest if $\sigma^{-2}\omega^2$ is large, that is, mutational effect sizes are large or selection is strong (meaning that the fitness landscape has a strong curvature). The same result holds for the time to extinction, which is strongly decreased in complex organisms (large d) if selection is strong (or mutations are large) but not if selection is weak (or mutations are small; Fig. S8).

Environmental change affecting a single trait also induces variation in other traits that are pleiotropically affected by the same set of genes. If mutation is isotropic (no correlation between the effects on different traits), the small-jumps approximation predicts that these latter traits have a lag with mean zero and a variance that is exactly half the variance of the lag of the first trait. If mutations are correlated, the lag in the first trait induces a correlated lag in the other traits (Fig. S4, S5), a phenomenon that has been called the “flying-kite effect” (Jones et al, 2004; Matuszewski et al, 2014). Strong mutational correlations also increase the lag of the first trait (Fig. S4A) and decrease extinction times (Fig. S9), showing that they function as a genetic constraint to adaptation, analogous to the one caused by genetic correlations in the standing genetic variation (Walsh and Blows, 2009; Chevin, 2012).

Comparison to previous models Both the adaptive-walk approximation and its small-jumps limit are closely related to approaches used in the theory of “adaptive dynamics” (a theory of mutation-limited evolution mostly used to study eco-evolutionary dynamics and frequency-dependent selection, e.g., Geritz et al, 1998). Indeed, our deterministic equation (9) for the mean lag can be seen as a version of the “canonical equation” (e.g., Champagnat et al, 2002; Champagnat, 2006) of adaptive dynamics (even though it lacks a factor 1/2 that is usually present in the latter). Similarly, our equation (17a) for the variance in the direction of the optimum follows from an instance of the “fluctuation equation” developed by Boettiger et al (2010). Our analysis and simulations thus provide some guidance for when adaptive-dynamics type approximations yield quantitatively accurate predictions for trait dynamics away from

equilibria or “singular strategies” (Geritz et al, 1998) – namely, when the optimal phenotype cannot be reached by a single mutation, but the fixation probability is still approximately linear in the selection coefficient.

More generally, and independently of the small-jumps limit, the good fit between adaptive-walk and individual-based simulations (Fig. S1, S2, S6, S7) shows that – at least in the simple scenario studied here – long-term evolution of a polygenic trait can be accurately predicted by just focusing on the selection gradient and the “incoming” genetic variance $(\Theta/2)\omega^2$ (which equals the product of the mutational variance $\mu\omega^2$ and the [effective] population size N), while neglecting other genetic details. This reflects results from Kopp and Hermisson (2009b), who had already shown that, in the moving-optimum model, the adaptive-walk approximation produces accurate predictions for the distribution of adaptive substitutions (i.e., jump sizes) even in situations with moderately high mutation rate. An exception occurs, however, when the population-wide mutation rate is high (such that several beneficial mutations enter the population each generation) and recombination is low or absent (e.g., in asexuals; Fig. S1, S2). It is well-known that, in this case, beneficial mutations that arise on different genetic backgrounds compete for fixation (a phenomenon called clonal interference or Hill-Robertson effect; Hill and Robertson, 1966; Gerrish and Lenski, 1998), which reduces the overall rate of adaptation and may strongly increase extinction risk in temporarily variable environments (not investigated here, but already shown by Bürger, 1999).

Short-term adaptation, in contrast, occurs mainly from standing genetic variation (Hermisson and Pennings, 2005; Barrett and Schluter, 2008; Matuszewski et al, 2015) and is most often modelled using approaches from quantitative genetics (Lande, 1976). This includes most previous applications of the moving-optimum model to questions about the phenotypic lag and population extinction risk (Lynch and Lande, 1993; Bürger and Lynch, 1995; Gomulkiewicz and Holt, 1995; Chevin et al, 2010). These models have the same gradient-structure as equation (9), and it is therefore not surprising that their results for the mean lag \bar{X}_∞ and the critical rate of environmental change v_{crit} are analogous to ours. In particular, in the model by Bürger and Lynch (1995) the mean lag in the one-dimensional case is given by $\bar{X}_\infty = v(\sigma_g^2 + V_s)/\sigma_g^2$, where σ_g^2 is the additive (standing) genetic variance, and V_s equals $\sigma^2/2$ in our notation. Replacing σ_g^2 by the “incoming” variance $(\Theta/2)\omega^2$ yields $\bar{X}_\infty/\omega = v/\omega + \gamma$. The difference to our equation (12) (i.e., the term v/ω) stems from the fact that standing genetic variation reduces the effective strength of selection (and, hence, increases the lag). Similarly, the approximate expression for the critical rate of environmental change in the Bürger and Lynch (1995) model, $v_{\text{crit}} = \sigma_g^2 \sqrt{2 \ln B / V_s}$ (see eq. A6 in Kopp and Matuszewski, 2014), is identical to equation (19) when again setting $\sigma_g^2 = (\Theta/2)\omega^2$.

It is tempting to use these similarities for a comparison of, for example, critical rates of environmental change when adaptation is based on either new mutations or standing genetic variation. However, such a comparison is problematic because the amount of standing genetic variation is difficult to predict (which makes the results by Bürger and Lynch, 1995 “deceptively simple” in the authors’ own words). If one assumes that the environment was stable before $t = 0$ (i.e., the optimum was stable for

a sufficiently long period before the onset of change), the amount of genetic variation at mutation-selection-drift balance is bounded above by the variance of a neutral trait, $\Theta\omega^2 = 2V_mN$ (Lynch and Hill, 1986), which is exactly twice the variance coming in from new mutations ($\Theta\omega^2/2$). This seems to suggest that, under weak selection, the presence of standing genetic variation might increase the critical rate of environmental change by at most a factor of 2. However, standing variance might be higher due to some sort of balancing selection, and we did not assess the effect of its immediate availability on the chances of evolutionary rescue under strong selection. Overall, the role of standing variation versus new mutations in preventing population extinctions is a topic that requires further study.

Unlike for the mean value of the lag, the predictions for its long-term variance differ markedly between quantitative genetic models and ours (or those from adaptive dynamics). In particular, a slightly simplified version of equation (8b) in Bürger and Lynch (1995) states that $\lim_{t \rightarrow \infty} \text{Var}(\bar{\mathbf{X}}) = V_s/(2N) + \sigma_g^2/(2V_s)$, with σ_g^2 and V_s as defined above. This expression is independent of the speed of environmental change, whereas our equation (17a) is linear in v . This difference likely reflects the fact that evolution in quantitative genetic models (from standing genetic variation) is “smooth”, whereas in our model, it happens in jumps, such that faster environmental change leads to larger jumps (Kopp and Hermisson, 2009b). Interestingly, this effect remains valid even in the limit of infinitesimally small jumps. Furthermore, our individual-based simulations (Fig. S2) show that the variance of the mean lag does indeed increase with v (and that eq. 17a is very accurate). The only exception occurs for weak selection and slow change, where the variance of the mean lag becomes independent of v , and indeed, approaches the value predicted for a constant optimum in quantitative genetics models ($V_s/(2N_e)$, Bürger and Lande, 1994; grey dashed lines in Fig. S2).

A key limitation of our model is its restriction to a deterministically and linearly moving optimum. We suspect that mild non-linearities will not change our qualitative conclusions, as long as the movement of the optimum remains monotonic. The effect of stochastic or deterministic fluctuations around a linear trend will depend upon their timescale. On the one hand, fluctuations that are faster than the time between adaptive steps essentially dilute the effect of selection (Bürger and Lynch, 1995), while leaving the rest of the model unchanged. For example, if the dynamics of a one-dimensional optimum are given by $v_t + \varepsilon_t$, where the ε_t are i.i.d. random variables drawn each generation from a Gaussian distribution with mean 0 and variance σ_v^2 , then this essentially increases the selection parameter from σ^2 to $\sigma^2 + \sigma_v^2$ (Bürger and Lynch, 1995; note, however, that large fluctuations will introduce additional demographic stochasticity, which is likely to increase extinction risk). On the other hand, fluctuations that are much slower than the time between adaptive steps will not alter the predictions over intermediate timescales. In contrast, fluctuations at a timescale similar to that of adaptive steps, or fluctuations without a trend (e.g. Bürger, 1999), are likely to have different effects, whose exploration is beyond the scope of this paper.

Conclusions The small-jumps limit to the adaptive-walk approximation provides easily interpretable analytical results about long-term adaptation in gradually changing environments. In particular, evolution in such environments resembles an Ornstein-

Uhlenbeck process around an average phenotypic lag, which allows to make predictions about the scaling of extinction risk on long (potentially) palaeontological timescales. The predictions are most accurate in parts of parameter space where selection is weak and the speed of environmental change is intermediate. For very slow environmental change, some additional insights can be gained from another approximation that assumes that all beneficial mutations are equally likely. Beyond the small-jumps limit, the adaptive-walk approximation itself, which neglects fixation time and co-segregation of alleles and is very easy to simulate, provides accurate predictions over a very wide range of parameters, provided that standing genetic variation is absent or has been depleted.

Acknowledgements We thank two anonymous reviewers for helpful comments on the manuscript.

References

- Barrett RDH, Schluter D (2008) Adaptation from standing genetic variation. *Trends Ecol Evol* 23:38–44
- Boettiger C, Dushoff J, Weitz JS (2010) Fluctuation domains in adaptive evolution. *Theoretical population biology* 77(1):6–13
- Bürger R (1999) Evolution of genetic variability and the advantage of sex and recombination in a changing environment. *Genetics* 153:1055–1069
- Bürger R, Lande R (1994) On the distribution of the mean and variance of a quantitative trait under mutation-selection-drift balance. *Genetics* 138:901–912
- Bürger R, Lynch M (1995) Evolution and extinction in a changing environment: a quantitative genetic analysis. *Evolution* 49:151–163
- Champagnat N (2006) A microscopic interpretation for adaptive dynamics trait substitution sequence models. *Stochastic Processes and their Applications* 116:1127–1160
- Champagnat N, Ferrière R, Ben Arous G (2002) The canonical equation of adaptive dynamics: a mathematical view. *Selection* 2(1):73–83
- Chevin LM (2012) Genetic constraints on adaptation to a changing environment. *Evolution* pp no–no, DOI 10.1111/j.1558-5646.2012.01809.x
- Chevin LM, Lande R, Mace GM (2010) Adaptation, plasticity, and extinction in a changing environment: towards a predictive theory. *PLoS Biology* 8(4):e1000357
- Collins S, Rost B, Rynearson TA (2014) Evolutionary potential of marine phytoplankton under ocean acidification. *Evolutionary applications* 7(1):140–155
- Finch S (2004) Ornstein-Uhlenbeck process. <http://www.people.fas.harvard.edu/~sfinch/csolve/ou.pdf>
- Fisher R (1930) *The genetical theory of natural selection*. Clarendon Press, Oxford
- Geritz S, E K, Meszéna G, Metz J (1998) Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary ecology* 12(1):35–57
- Gerrish PJ, Lenski RE (1998) The fate of competing beneficial mutations in an asexual population. *Genetica* 102/103:127–144

- Gillespie JH (1983a) A simple stochastic gene substitution model. *Theor Pop Biol* 23:202–215
- Gillespie JH (1983b) Some properties of finite populations experiencing strong selection and weak mutation. *The American Naturalist* 121(5):691–708
- Gomulkiewicz R, Holt RD (1995) When does Evolution by Natural Selection Prevent Extinction? *Evolution* 49(1):201–207
- Gomulkiewicz R, Houle D (2009) Demographic and genetic constraints on evolution. *Am Nat* 174(6):E218–229
- Gomulkiewicz R, Krone SM, Remien CH (2017) Evolution and the duration of a doomed population. *Evolutionary Applications* 10:471–484
- Gonzalez A, Ronce O, Ferrière R, Hochberg M (2013) Evolutionary rescue: an emerging focus at the intersection of ecology and evolution. *Philos Tr R Soc Lon B* 368:20120404
- Hairston NG, Ellner SP, Geber MA, Yoshida T, Fox JA (2005) Rapid evolution and the convergence of ecological and evolutionary time. *Ecology Letters* 8:1114–1127
- Haldane JBS (1927) A mathematical theory of natural and artificial selection. Part V: selection and mutation. *Proc Camb Philos Soc* 23:838–844
- Hermisson J, Pennings PS (2005) Soft sweeps: molecular population genetics of adaptation from standing genetic variation. *Genetics* 169:2335–2352
- Hill WG, Robertson A (1966) The effect of linkage on limits to artificial selection. *Genet Res* 8:269–294
- Jones AG, Arnold SJ, Bürger R (2004) Evolution and stability of the G-matrix on a landscape with a moving optimum. *Evolution* 58:1636–1654
- Kimura M (1962) On the probability of fixation of mutant genes in a population. *Genetics* 47(6):713
- Kopp M, Hermisson J (2007) Adaptation of a quantitative trait to a moving optimum. *Genetics* 176:715–718
- Kopp M, Hermisson J (2009a) The genetic basis of phenotypic adaptation I: Fixation of beneficial mutations in the moving optimum model. *Genetics* 182:233–249
- Kopp M, Hermisson J (2009b) The genetic basis of phenotypic adaptation II: The distribution of adaptive substitutions in the moving optimum model. *Genetics* 183:1453–1476
- Kopp M, Matuszewski S (2014) Rapid evolution of quantitative traits: theoretical perspectives. *Evol Appl* 7:169–191
- Lande R (1976) Natural selection and random genetic drift in phenotypic evolution. *Evolution* 30(2):314–334
- Lynch M, Hill WG (1986) Phenotypic evolution by neutral mutation. *Evolution* 40(5):915–935
- Lynch M, Lande R (1993) Evolution and extinction in response to environmental change. In: Kareiva P, Kingsolver JG, Juey RB (eds) *Biotic interactions and global change*, Sinauer, Sunderland, pp 43–53
- Malécot G (1952) Les processus stochastiques et la méthode des fonctions génératrices ou caractéristiques. *Publications de l’institut de Statistiques de l’Université de Paris* 1:1–16
- Martin G, Lenormand T (2006) A general multivariate extension of Fisher’s geometrical model and the distribution of mutation fitness effects across species. *Evolution*

- 60:893–907
- Matuszewski S, Hermisson J, Kopp M (2014) Fisher’s geometric model with a moving-optimum. *Evolution* 68:2571–2588
- Matuszewski S, Hermisson J, Kopp M (2015) Catch me if you can: Adaptation from standing genetic variation to a moving phenotypic optimum. *Genetics* 200(4):1255–1274
- Nassar E (2016) Modèles probabilistes de l’évolution d’une population dans un environnement variable. PhD thesis, Aix-Marseille University
- Nassar E, Pardoux E (2017) On the long time behaviour of the solution of an SDE driven by a Poisson Point Process. *Advances in Applied Probability* 49:344–367
- Nassar E, Pardoux E (2018) Small jumps asymptotic of the moving optimum Poissonian SDE, in revision
- Orr HA (1998) The population genetics of adaptation: the distribution of factors fixed during adaptive evolution. *Evolution* 52:935–949
- Orr HA (2000) Adaptation and the cost of complexity. *Evolution* 54:13–20
- Orr HA (2005) The genetic theory of adaptation: a brief history. *Nat Rev Gen* 6:119–127
- Orr HA, Unckless RL (2008) Population extinction and the genetics of adaptation. *The American Naturalist* 172(2):160–169
- Orr HA, Unckless RL (2014) The population genetics of evolutionary rescue. *PLoS Genetics* 10:e1004551, DOI 10.1371/journal.pgen.1004551
- Peischl S, Kirkpatrick M (2012) Establishment of new mutations in changing environments. *Genetics*
- Pigolotti S, Flammini A, Marsili M, Maritan A (2005) Species lifetime distribution for simple models of ecologies. *Proc Natl Acad Sci USA* 102:15747–15751
- Ricciardi LM, Sato S (1988) First-passage-time density and moments of the Ornstein-Uhlenbeck process. *J Appl Prob* 25:43–57
- Stenseth NC, Smith JM (1984) Coevolution in ecosystems: Red queen evolution or stasis? *Evolution* 38(4):870–880
- Tenaillon O (2014) The utility of fisher’s geometric model in evolutionary genetics. *Annual Review of Ecology, Evolution, and Systematics* 45:179–201
- Thomas MU (1975) Mean first-passage time approximations for the Ornstein-Uhlenbeck process. *J Appl Prob* 12:600–604
- Uecker H, Hermisson J (2011) On the fixation process of a beneficial mutation in a variable environment. *Genetics* 188(4):915–930
- Van Valen L (1973) A new evolutionary law. *Evol Theory* 1:1–30
- Walsh B, Blows MW (2009) Abundant genetic variation + strong selection = multivariate genetic constraints: A geometric view of adaptation. *Annu Rev Ecol Evol Syst* 40(1):41–59, DOI 10.1146/annurev.ecolsys.110308.120232

5 Appendix A: Stochastic differential equations for the phenotypic lag

According to Nassar (2016) and Nassar and Pardoux (2017), the evolution of the phenotypic lag \mathbf{X}_t of the population can be described by the stochastic differential

equation

$$\mathbf{X}_t = \mathbf{x}_0 + \mathbf{v}t - \int_{[0,t] \times \mathbb{R}^d \times [0,1]} \alpha \Gamma(\mathbf{X}_{s^-}, \alpha, \xi) \mathcal{N}(ds, d\alpha, d\xi). \quad (\text{A1})$$

Here, \mathcal{N} is a Poisson point process over $\mathbb{R}_+ \times \mathbb{R}^d \times [0, 1]$ with intensity $ds v(d\alpha) d\xi$ where

$$v(d\alpha) = \frac{\Theta}{2} p(\alpha) d\alpha$$

and

$$\Gamma(\mathbf{x}, \alpha, \xi) = \mathbf{1}_{\{\xi \leq g(\mathbf{x}, \alpha)\}}.$$

Recall that $g(\mathbf{x}, \alpha)$ is the fixation probability of a mutation of size α that hits the population when the lag is \mathbf{x} , as defined by (6) and (7). The points of the Poisson point process $(T_i, \mathbf{A}_i, \Xi_i)$ are such that the (T_i, \mathbf{A}_i) form a Poisson point process over $\mathbb{R}_+ \times \mathbb{R}^d$ of the mutations that hit the population with intensity $ds v(d\alpha)$, and the Ξ_i are i.i.d. $\mathcal{U}[0, 1]$, globally independent of the Poisson point process of the (T_i, \mathbf{A}_i) . T_i 's are the times when mutations are proposed and \mathbf{A}_i 's are the effect sizes of those mutations. The Ξ_i are auxiliary variables determining fixation: a mutation gets instantaneously fixed if $\Xi_i \leq g(\mathbf{X}_{T_i}, \mathbf{A}_i)$, and is lost otherwise.

The rescaled process \mathbf{X}_t^ε (Nassar and Pardoux, 2018) is given by

$$\mathbf{X}_t^\varepsilon = \mathbf{x}_0^\varepsilon + \mathbf{v}t - \int_{[0,t] \times \mathbb{R}^d \times [0,1]} \varepsilon \alpha \Gamma(\mathbf{X}_{s^-}^\varepsilon, \varepsilon \alpha, \xi) \mathcal{N}_\varepsilon(ds, d\alpha, d\xi), \quad (\text{A2})$$

where the intensity measure of the Poisson point process N_ε is $\varepsilon^{-2} ds \times v(d\alpha) \times d\xi$, and the small-jumps limit is obtained by letting $\varepsilon \rightarrow 0$. An intuitive explanation for why the intensity measure (and, hence, the mutation rate Θ) needs to be rescaled by ε^{-2} as mutational effects (and, hence, jump sizes) α are rescaled by ε is that smaller mutations have a double effect on the rate of adaptation: Not only do they induce smaller phenotypic changes, but they also have lower fixation probability (approximately linear in ε).

6 Appendix B: Summary statistics for adaptive-walk simulations

Summary statistics for adaptive-walk simulations depend on the evolution of the phenotypic lag between adaptive steps. Let t_k be the time between steps $k-1$ and k , ξ_k the size of the lag immediately after step $k-1$ and $\zeta_k = \xi_k + \mathbf{v}t_k$ the size of the lag just before step k . The mean lag over n steps is given by

$$\bar{X} = \frac{\sum_{k=1}^n t_k (\xi_k + \zeta_k) / 2}{\sum_{k=1}^n t_k},$$

and the variance is

$$\text{Var}(X_1) = \frac{\sum_k |\xi_{k,1}^3 - \zeta_{k,1}^3|}{3v \sum_k t_k} - \bar{X}_1^2$$

in the direction of the optimum (i.e., for trait 1), and

$$\text{Var}(X_{i>1}) = \frac{\sum_k t_k \xi_{k,i}^2}{\sum_k t_k} - \bar{X}_i^2$$

in all other directions. Covariances involving trait 1 are

$$\text{Cov}(X_1, X_{i>1}) = \frac{\sum_k t_k \xi_{k,i} (\xi_{k,1} + \zeta_{k,1}) / 2}{\sum_k t_k} - \bar{X}_1 \bar{X}_i$$

and those not involving trait 1 are

$$\text{Cov}(X_{i>1}, X_{j>1}) = \frac{\sum_k t_k \xi_{k,i} \xi_{k,j}}{\sum_k t_k} - \bar{X}_i \bar{X}_j.$$

7 Appendix C: The large-time variance in the small-jumps limit

As shown in Nassar and Pardoux (2018), for $\varepsilon \rightarrow 0$, the process U_t^ε (15) converges to the Ornstein-Uhlenbeck process (16), where the infinitesimal variance is given by

$$\Lambda(\mathbf{x}) = 2\Theta\sigma^{-2} \int_{(\mathbf{x}|\alpha) \leq 0} |(\mathbf{x}|\alpha)| \alpha \otimes \alpha p(\alpha) d\alpha \quad (\text{C1})$$

(here, $(\cdot|\cdot)$ denotes the inner product of two vectors and \otimes the outer product). Thus,

$$\mathbf{U}_t = \int_0^t e^{-\Theta\sigma^{-2}(t-s)\mathbf{M}} \Lambda^{\frac{1}{2}}(\bar{\mathbf{X}}_s) d\mathbf{B}_s, \quad (\text{C2})$$

and its second moment is given by

$$\mathbb{E}(\mathbf{U}_t \otimes \mathbf{U}_t) = \int_0^t e^{-\Theta\sigma^{-2}(t-s)\mathbf{M}} \Lambda(\bar{\mathbf{X}}_s) e^{-\Theta\sigma^{-2}(t-s)\mathbf{M}} ds, \quad (\text{C3})$$

or, by a change of variables and writing $\mathbf{A} = \Theta\sigma^{-2}\mathbf{M}$,

$$\mathbb{E}(\mathbf{U}_t \otimes \mathbf{U}_t) = \int_0^t e^{-s\mathbf{A}} \Lambda(\bar{\mathbf{X}}_{t-s}) e^{-s\mathbf{A}} ds. \quad (\text{C4})$$

We call $\bar{\mathbf{S}}^2$ the limit of this second moment for $t \rightarrow \infty$ (when $\bar{\mathbf{X}}_t \rightarrow \bar{\mathbf{X}}_\infty$):

$$\bar{\mathbf{S}}^2 = \lim_{t \rightarrow \infty} \mathbb{E}(\mathbf{U}_t \otimes \mathbf{U}_t) = \int_0^\infty e^{-s\mathbf{A}} \Lambda(\bar{\mathbf{X}}_\infty) e^{-s\mathbf{A}} ds. \quad (\text{C5})$$

An alternative way of characterizing $\bar{\mathbf{S}}^2$ is as the solution to the equation

$$\mathbf{A}\bar{\mathbf{S}}^2 + \bar{\mathbf{S}}^2\mathbf{A} = \Lambda(\bar{\mathbf{X}}_\infty) \quad (\text{C6})$$

(this follows from equation (C5) by noting that $d/dt(e^{-t\mathbf{A}}\Lambda e^{-t\mathbf{A}}) = -\mathbf{A}e^{-t\mathbf{A}}\Lambda e^{-t\mathbf{A}} - e^{-t\mathbf{A}}\Lambda e^{-t\mathbf{A}}\mathbf{A}$ and was used for the numerical calculation of predictions in Fig. S4 and S5).

We now write $\mathbf{M} = \bar{\omega}^2 \mathbf{M}_0$ (see eq. 5). $\Lambda(\bar{\mathbf{X}}_\infty)$ can then be written as

$$\Lambda(\bar{\mathbf{X}}_\infty) = \frac{2\nu}{\bar{\omega}^2} \int_{(\mathbf{M}_0^{-1} e_1 | \alpha) \leq 0} |(\mathbf{M}_0^{-1} e_1 | \alpha)| \alpha \otimes \alpha p(\alpha) d\alpha, \quad (\text{C7})$$

and hence

$$\mathbf{M}_0 \bar{\mathbf{S}}^2 + \bar{\mathbf{S}}^2 \mathbf{M}_0 = \frac{2\sigma^2 \nu}{\Theta \bar{\omega}^4} \int_{(\mathbf{M}_0^{-1} e_1 | \alpha) \leq 0} |(\mathbf{M}_0^{-1} e_1 | \alpha)| \alpha \otimes \alpha p(\alpha) d\alpha. \quad (\text{C8})$$

Therefore,

$$\bar{\mathbf{S}}^2 = \frac{\sigma^2 \nu}{\Theta \bar{\omega}^4} \Phi(\mathbf{M}_0), \quad (\text{C9})$$

where $\Phi(\mathbf{M}_0)$ is a matrix that depends only on \mathbf{M}_0 , that is, only on the structure of the mutation matrix. This shows, in particular, that in Fig. S4E,F where $\mathbf{M} = \omega^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$, the ratio of variances of the lag in different dimensions and the correlations between these lags depend only on ρ .

Finally, an explicit expression for $\bar{\mathbf{S}}^2$ is possible in the particular case where the mutation matrix is isotropic ($\mathbf{M} = \omega^2 \mathbf{I}$). In this case,

$$\bar{\mathbf{S}}^2 = \frac{1}{2\Theta \sigma^{-2}} \mathbf{M}^{-1} \Lambda(\bar{\mathbf{X}}_\infty), \quad (\text{C10})$$

and the matrix $\Lambda(\bar{\mathbf{X}}_\infty)$ evaluates to

$$\begin{aligned} \Lambda(\bar{\mathbf{X}}_\infty) &= \frac{2\nu}{\bar{\omega}^2} \int_{\mathbb{R}^{d-1}} \int_0^\infty \alpha_1 \alpha \otimes \alpha p(\alpha) d\alpha \\ &= \frac{2\nu}{(2\pi)^{\frac{d}{2}} \bar{\omega}^{d+2}} \int_{\mathbb{R}^{d-1}} \int_0^\infty \alpha_1 \alpha \otimes \alpha e^{-\frac{1}{2\bar{\omega}^2} \sum_{i=1}^d \alpha_i^2} d\alpha. \end{aligned} \quad (\text{C11})$$

The first element of this matrix is

$$\Lambda_{1,1}(\bar{\mathbf{X}}_\infty) = \frac{2\nu}{\sqrt{2\pi} \bar{\omega}^3} \int_0^\infty \alpha_1^3 e^{-\frac{\alpha_1^2}{2\bar{\omega}^2}} d\alpha_1 = \frac{4\nu \bar{\omega}}{\sqrt{2\pi}}. \quad (\text{C12})$$

For $i = 2, \dots, d$,

$$\Lambda_{i,i}(\bar{\mathbf{X}}_\infty) = \frac{2\nu}{2\pi \bar{\omega}^4} \int_{\mathbb{R}} \alpha_i^2 e^{-\frac{\alpha_i^2}{2\bar{\omega}^2}} d\alpha_i \int_0^\infty \alpha_1 e^{-\frac{\alpha_1^2}{2\bar{\omega}^2}} d\alpha_1 = \frac{2\nu \bar{\omega}}{\sqrt{2\pi}}.$$

For $2 \leq i < j \leq d$,

$$\Lambda_{i,j}(\bar{\mathbf{X}}_\infty) = \Lambda_{j,i}(\bar{\mathbf{X}}_\infty) = \frac{2\nu}{(2\pi)^{\frac{3}{2}} \bar{\omega}^5} \int_{\mathbb{R}} \alpha_i e^{-\frac{\alpha_i^2}{2\bar{\omega}^2}} d\alpha_i \int_{\mathbb{R}} \alpha_j e^{-\frac{\alpha_j^2}{2\bar{\omega}^2}} d\alpha_j \int_0^\infty \alpha_1 e^{-\frac{\alpha_1^2}{2\bar{\omega}^2}} d\alpha_1 = 0.$$

Similarly, for all $i \neq 1$,

$$\Lambda_{1,i}(\bar{\mathbf{X}}_\infty) = \Lambda_{i,1}(\bar{\mathbf{X}}_\infty) = 0.$$

Hence,

$$\Lambda(\bar{\mathbf{X}}_\infty) = \frac{2v\omega}{\sqrt{2\pi}} \begin{pmatrix} 2 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 \end{pmatrix}. \quad (\text{C13})$$

It follows that

$$\bar{\mathbf{S}}^2 = \frac{v}{\sqrt{2\pi}\Theta\sigma^{-2}\bar{\omega}} \begin{pmatrix} 2 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 \end{pmatrix} = \frac{\omega^2\gamma}{\sqrt{2\pi}} \begin{pmatrix} 2 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 \end{pmatrix}, \quad (\text{C14})$$

because, for $\mathbf{M} = \omega^2\mathbf{I}$, $\bar{\omega} = \omega$.

8 Appendix D: The environmentally-limited regime

As argued in the main text, $\gamma \ll 1$ corresponds to the environmentally-limited regime, in which the lag is small relative to the size of new mutations and the small-jumps approximation fails. Matuszewski et al (2014) (see also Kopp and Hermisson, 2009b) showed that in this regime the distribution of new mutations can be approximated by a uniform distribution with density equal to $p(\mathbf{0})$, the value of the density at $\alpha = \mathbf{0}$. Unfortunately, this approximation does not allow to calculate the long-term moments of the process. We can, however, gain some insights from focusing on the lag before and after the first adaptive substitution. In particular, it follows from equation (S18) and (S28) in Matuszewski et al (2014) that, just before the first substitution, the lag in the direction of the optimum has the cumulative distribution function

$$\mathcal{P}(X_1^{1-} \leq x) = 1 - \exp\left(-\frac{\eta(d)p(\mathbf{0})}{\gamma}x^{d+3}\right) \quad (\text{D1})$$

with mean

$$\mathbb{E}(X_1^{1-}) = \left(\frac{\gamma}{\eta(d)p(\mathbf{0})}\right)^{\frac{1}{d+3}} \Gamma\left(\frac{d+4}{d+3}\right) \quad (\text{D2})$$

and variance

$$\text{Var}(X_1^{1-}) = \left(\frac{\gamma}{\eta(d)p(\mathbf{0})}\right)^{\frac{2}{d+3}} \left[\Gamma\left(\frac{d+5}{d+3}\right) - \Gamma\left(\frac{d+4}{d+3}\right)^2 \right], \quad (\text{D3})$$

where

$$\eta(d) = \frac{\pi^{\frac{d}{2}}}{(d+3)\Gamma(2+\frac{d}{2})} \quad (\text{D4})$$

and

$$p(\mathbf{0}) = \left(\frac{1}{\sqrt{2\pi\bar{\omega}^2}}\right)^d. \quad (\text{D5})$$

Obviously, the lag in all other directions is zero before the first jump. Immediately after the first jump, the lag in any direction has mean 0 and variance

$$\text{Var}(X_i^{1+}) = \frac{1}{d+4} \left(\frac{\gamma}{\eta(d)p(\mathbf{0})} \right)^{\frac{2}{d+3}} \Gamma\left(\frac{d+5}{d+4}\right) \quad (\text{D6})$$

(equation S33 of Matuszewski et al, 2014). These results explain several patterns seen in the environmentally-limited regime.

First, the mean lag in the direction of the optimum is proportional to $\gamma^{1/(d+3)}$ (eq. D2). For the log-log plots in Figures 2A and S3A, this explains the initial slope of the curves for large $\sigma^{-2}\omega^2$. Indeed, the slope between the first two values of $v/(\omega\Theta)$ for $\sigma^{-2}\omega^2 = 1$ (red curve) is 1/3.95 for $d = 1$ (Fig. 2A) and 1/6.32 for $d = 4$ (Fig. S3A). In contrast, in the genetically-limited regime, the slope equals 1, and close to the boundary of the transient case, it tends towards infinity.

Similarly, the variance of the lag in the direction of the optimum both before and after the first step is proportional to $\gamma^{2/(d+3)}$ (eq. D3 and D6), and in consequence, the standard deviation is proportional to $\gamma^{1/(d+3)}$ (same as the mean). This explains the initial slope of the $\sigma^{-2}\omega^2 = 1$ curves in Figures 2B and S3B, which is 1/3.93 for $d = 1$ (Fig. 2B) and 1/6.67 for $d = 4$ (Fig. S3B). It also explains why, in the environmentally-limited regime, the coefficient of variation (i.e., the standard deviation divided by the mean) tends to be independent of γ , and when using equation (D3), why it is close to one (Fig. 2C,D and Fig. S3C; note that the gamma functions in eq. D2 and D3 are close to unity and can be neglected).

In contrast, in the multivariate case of the environmentally-limited regime, the variance of the lag in the direction of the optimum seems to exceed the variance in the other directions by about a factor of $\sqrt{2}$ if there are no mutational correlations (Fig. S3D), but we do not have an analytical explanation for this observation. Note that, due to symmetry, the variance of the lag is identical in all directions immediately after the first jump (eq. D6), but in between jumps, X_1 increases while the other components stay constant.

Finally, the above results can explain the dependence of extinction time on the speed of environmental change in the case of very strong selection (the only case where a non-negligible extinction risk exists in the environmentally-limited regime). For very high $\sigma^{-2}\omega^2$ (Fig. 3E,F), extinction times from simulations appear to be approximately exponential in $1/v$. This is in agreement with equation (D1), according to which the probability that the process reaches X_{crit} before the first jump (when starting at 0) is given by $\exp(-X_{\text{crit}}^4/(3\sqrt{2\pi}\omega^2\gamma))$, and hence, the time until this event occurs for the first time should arguably scale with $\exp(1/v)$ (but not $\exp(1/\gamma)$, since X_{crit} depends on σ^2).

9 Supplementary Figures




Fig. S1 Comparison of the normalized mean phenotypic lag \bar{X}/ω predicted by equation (12) (dashed line), adaptive-walk simulations (black symbols; same data as in Fig. 2A) and individual-based simulations with and without recombination (red symbols: $\Theta = 1, r = 0.5$; blue symbols: $\Theta = 10, r = 0$; averaged over 10^6 generations) with $B = 8$ offspring per individual, in the one-dimensional model ($d = 1$). Missing values at low values of $v/(\omega\Theta)$ are those where the mean lag in the individual-based simulations was (slightly) negative (i.e., the mean phenotype was overshooting the optimum). Missing values at high values of $v/(\omega\Theta)$ are those where the individual-based population went extinct. Error bars represent the standard deviation of the mean also shown in Figure S2. For further details, see Figure 2.

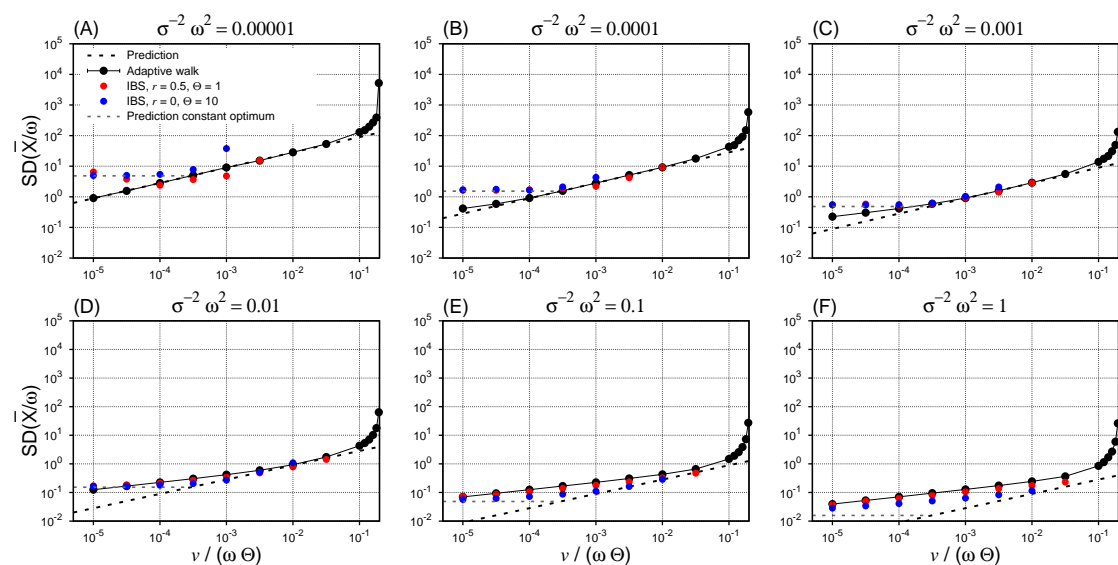


Fig. S2 Comparison of the standard deviation of the normalized phenotypic lag X/ω predicted by equation (17) (dashed line), adaptive-walk simulations (black symbols; same data as in Fig. 2B) and individual-based simulations (red and blue symbols) with $B = 8$ offspring per individual, in the one-dimensional model ($d = 1$). The grey dashed line is the prediction for a constant optimum in quantitative genetic models ($\sqrt{V_s}/(2N_e)$, see Discussion). For further details, see Figure S1.




Fig. S3 The long-term steady state of the normalized phenotypic lag \mathbf{X}/ω in the model with $d = 4$ traits and isotropic mutations ($\mathbf{M} = \omega^2 \mathbf{I}$). (A) to (C) are as in Fig. 2 for the lag in the direction of the optimum (that is, for trait 1, whose optimum is directly affected by environmental change, see eq. 2). (D) shows the ratio of the variances of the lag with respect to traits 1 and 2 (the latter being orthogonal to the direction of the optimum). The horizontal dashed line shows the prediction from eq. (17). The solid vertical lines mark the boundary of the transient case for $d = 1$, which should be unchanged in the multidimensional case as long as mutations are isotropic, even though this conjecture has not been proved. For further details, see Fig. 2.




Fig. S4 The long-term steady state of the normalized phenotypic lag \mathbf{X}/ω in the model with $d = 2$ traits, for fast environmental change, $v/(\omega\Theta) = 0.01$, and various values of the scaled selection strength $\sigma^{-2}\omega^2$, as a function of the strength of mutational correlations, ρ , when the mutation matrix is $\mathbf{M} = \omega^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$. Dashed lines are predictions from the small-jumps limit (eq. 12 and C5), whereas points show results from adaptive-walk simulations (averaged over 1000000 adaptive steps). (A) The mean lag in the direction of the optimum (i.e., for trait 1). (B) The ratio of mean lags for traits 2 (orthogonal to the direction of the optimum) and 1, showing the “flying-kite effect” (with our choice of M , the predicted ratio is simply equal to ρ). (C) and (D) The standard deviations of the lag for traits 1 and 2 respectively. (E) The ratio of the variances of the lag for traits 1 and 2. (F) The correlation of the lags for traits 1 and 2.




Fig. S5 The long-term steady state of the normalized phenotypic lag \mathbf{X}/ω in the model with $d = 2$ traits and moderate mutational correlations, that is, a mutation matrix $\mathbf{M} = \omega^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$ with $\rho = 0.5$. Trait 1 is in the direction of the optimum (eq. 2), whereas trait 2 is orthogonal to this direction. (A) shows the lag for trait 1 with the predictions (dashed lines) from eq. (12), whereas (B) shows the ratio of the lags for traits 2 and 1 (“flying-kite effect”). (C) and (D) show the standard deviations of the lag for trait 1 and 2, respectively, with the predictions obtained by numerical evaluation of eq. (C5). (E) shows the ratio of the variances of the lag for traits 1 and 2. Note that the predicted value depends only on ρ . (F) shows the correlation of the lag in directions 1 and 2, with the prediction again depending only on ρ . High values of $v/(\omega\Theta)$ are close to the transient case (see Fig. 2), but the exact boundary is unknown. Adaptive-walk simulations were run for between 10^5 and 10^7 adaptive steps (for low and high $v/(\omega\Theta)$, respectively), with a burn-in period of one tenth of this duration. For further details, see Fig. 2.




Fig. S6 Comparison of the “time to extinction”, T_e (relative to the mean interval between new mutations, Θ^{-1}), in adaptive-walk and individual-based simulations for a single evolving trait ($d = 1$). Results shown in black are the same as in Fig. 3 (solid line: analytical predictions from eq. 20 and 21; circles and crosses: means and standard deviations of the time to reach fitness $1/B = 1/2$ in adaptive-walk simulations). Colored symbols show results from individual-based simulations (averages over 20 replicates). Red circles and crosses are the means and standard deviations of the first time the population mean fitness dropped to $1/B$. Blue circles and crosses are for the times the population actually went extinct. No individual-based results are shown for parameter values where the extinction time in some replicates exceeded 2×10^7 generations.




Fig. S7 Comparison of the “time to extinction”, T_e (relative to the mean interval between new mutations, Θ^{-1}), in adaptive-walk and individual-based simulations for a model with three pleiotropic traits ($d = 3$). The black line is the analytical prediction from the one-dimensional model ($d = 1$, eq. 20 and 21, same as in Fig. 3). Black circles and crosses represent the means and standard deviations of the time to reach fitness $1/B = 1/2$ in adaptive-walk simulations. Red symbols show the same results for individual-based simulations (averages over 20 replicates). Blue circles and crosses are for the time the individual-based population actually went extinct. For further details see Fig. 3 and S6.

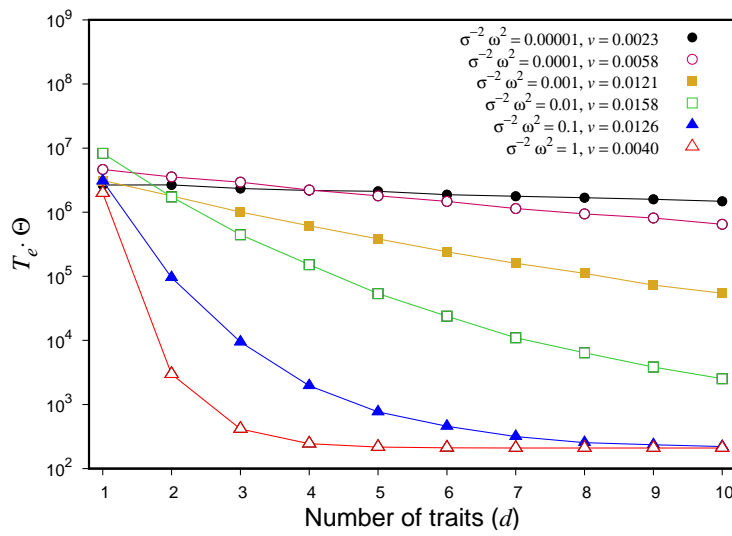


Fig. S8 The time T_e (relative to the mean interval between new mutations, Θ^{-1}) until the population mean fitness $W(\mathbf{X}_t)$ drops below $1/B = 1/2$ for the first time (“time to extinction”) as a function of the number of traits, d , for different values of the scaled selection strength $\sigma^{-2}\omega^2$. The speed of environmental change, ν was chosen among those used in Fig. 3 such that $T_e\Theta$ for $d = 1$ is between 10^6 and 10^7 . Results are means from 1000 adaptive-walk simulations.




Fig. S9 The time T_e (relative to the mean interval between new mutations, Θ^{-1}) until the population mean fitness $W(X_t)$ drops below $1/B = 1/2$ for the first time (“time to extinction”) in the model with $d = 2$ traits as a function of mutational correlation ρ , when the mutation matrix $\mathbf{M} = \omega^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$. The dotted line shows the analytical prediction for the one-dimensional case ($d = 1$, eq. 20), and the dashed line the simulation results for this case. Results are means from 1000 adaptive-walk simulations.