*A copy-edited version of this paper will appear in the journal Artificial Life, published by MIT Press.*

# Simulating the Effect of Environmental Change on Evolving Populations

John A. Bullinaria

School of Computer Science
University of Birmingham
Birmingham, B15 2TT, UK

*j.a.bullinaria@cs.bham.ac.uk*

**Abstract:** This study uses evolutionary simulations to explore the strategies that emerge to enable populations to cope with random environmental changes in situations where lifetime learning approaches are not available to accommodate them. In particular, it investigates how the average magnitude of change per unit time and the persistence of the changes (and hence the resulting autocorrelation of the environmental time series) affect the change tolerances, population diversities and extinction timescales that emerge. Although it is the change persistence (often discussed in terms of environmental noise colour) that has received most attention in the recent literature, other factors, particularly the average change magnitude, interact with this and can be more important drivers of the adaptive strategies that emerge. Moreover, when running simulations, the choice of change representation and normalization can also affect the outcomes. Detailed simulations are presented that are designed to explore all these issues. They also reveal significant dependences on the associated mutation rates and the extent to which they can evolve, and clarify how evolution often leads populations into strategies with higher risks of extinction. Overall, this study shows how modeling the effect of environmental change requires more care than may have previously been realized.

**Keywords:** Agent-based models, Evolution, Variable environments, Change persistence, Mutation rates, Extinction timescales.

## 1. Introduction

The important issue of understanding the effect of environmental change on evolving populations is a subject of long-standing interest that has already received considerable attention in the literature (e.g., Bürger & Lynch, 1995; Gilljam et al., 2019; Haaland, Wright & Ratikainen, 2020; Halley, 1996; Halley & Inchausti, 2004; Inchausti & Halley, 2003; Morales, 1999; Petchey, Gonzalez & Wilson, 1997; Potts, 2013; Ruokolainen et al., 2009; Steele, 1985; Vasseur & Yodzis, 2004). This has involved some debate about which aspects of environmental change are most crucial to understanding the outcomes, with various suggestions of how best to model the key factors, and somewhat inconsistent conclusions. In the hope of providing further clarification, this has recently led to the exploration of such matters using a different approach, namely agent-based artificial life simulations (Bullinaria, 2018, 2022; Grove et al., 2020, 2021), which has proved successful in the past for modeling evolving ecosystems (e.g., Wilke et al., 2001; Yedid, Ofria & Lenski, 2009).

When lifetime learning is available for individuals in the population, suitable learning strategies may evolve to cope with the environmental changes, and the simulations of Bullinaria (2022) have shown that in this case it is the average total magnitude of change per unit time that has the main effect on the evolutionary outcomes, and there is very little dependence on the actual distribution of individual change magnitudes. However, when effective lifetime learning is not possible, survival of the species in question will depend on evolution resulting in other ways to deal with the environmental changes. It is already known that the life history traits emerging from evolution for different species depend on the nature of their environmental conditions (Roff, 2002; Stearns, 1992). One obvious approach that might emerge here would be to maintain a population that has sufficient diversity that there are always some individuals able to survive whatever environmental change may occur, so that the population as a whole can track the changes. This could, for example, occur by evolving to produce large numbers of offspring which have much variation and require little parental care. Another, potentially more efficient, approach might be to maintain less diversity but have all individuals develop a generalist tolerance to the changing environmental conditions, though generalism tends to have its own costs (Haaland et al., 2020). It is fairly obvious that different strategies will emerge depending on the magnitude and persistence of the changes, but it is not clear what tolerance levels will be best in different circumstances, nor how they will interact with other aspects of the species' evolution such as the selection pressures and mutation rates that control the population diversity.

The agent-based simulations of Grove et al. (2020, 2021) explored the evolution of such a tolerance in a minimalist model involving the key factors, and found a clear dependence on the persistence of the changes as measured by the autocorrelation or colour of the environmental time series. However, an issue with those simulations was the way they involved normalizing the magnitude of the whole time series, whereas the simulations of Bullinaria (2022) indicated that to achieve reliable results it is important to explore separately the effects of the change magnitude and distribution, along with the relative timescales of the change events and processes by which the changes are accommodated by the modeled species. The aim of this paper is to look again at the models of Grove et al. (2020, 2021) and separate the effects of the average change magnitude per unit time and the change persistence, and to explore how other factors, such as the change representation and mutation rates, affect the outcomes. It also investigates how the various model details affect the crucial extinction timescales that emerge in the simulations.

The remainder of this paper is organized as follows: The next section describes and justifies the simplified evolutionary model on which the necessary simulations are based. Sections 3 and 4 then present the key simulation results, first with mutation rates optimized by evolution, and then with fixed sub-optimal mutation rates. The effect of normalizing the change magnitude in those simulations is studied in Section 5. Then the modeling of extinction events in the various simulated cases is explored in Section 6. The paper ends with some conclusions and discussion.

## 2. The Simplified Evolutionary Model

As usual with artificial life simulations, the simplest possible model of the scenario of interest should be employed – one that allows exploration of all the relevant factors while abstracting out any unnecessary details that may confuse the outcomes. A standard agent-based evolutionary modeling approach will therefore be utilized that maintains a population of abstract simulated individuals, each of which is specified by a minimal set of evolvable parameters representing their crucial innate properties (Eiben & Smith, 2016). The key factors underlying the model required here are a changing environment and a measure of the performance of individuals responding to that environment. Following Grove et al. (2020, 2021), the environment will be represented by a single parameter $E$ that varies randomly over time, and the innate properties of each individual are an evolved parameter $R$ which represents their response to the prevailing environmental conditions, and an evolved tolerance parameter $T$ that represents their tolerance to having their response deviate from what is optimal for the current environment. This minimalist framework was originally formulated and justified by Grove (2014) with the tolerance representing the trade-off between specialism and generalism/versatility for differing levels of environmental variability. Whether $E$ summarizes a whole complex environment, or just a single physical aspect of it such as temperature or rainfall, or any number of factors related to interactions with other species, needs not be specified. The performance $P$ of each individual at each point in time can then, as argued by Lynch & Gabriel (1987), be quantified as a standard normalized Gaussian function

$$P = \frac{1}{T\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{E-R}{T}\right)^2} \tag{1}$$

that is maximum when $R$ exactly matches $E$ and falls off more slowly as the distribution "width" $T$ increases, but penalizes having $T$ larger than necessary as generalism will generally involve some kind of cost. The scale of $E$, $R$ and $T$ is essentially arbitrary, as is the baseline or initial value for $E$ and $R$ which can be taken to be zero so that $E$ represents the cumulative environmental change at each stage.

## 2.1 Representing the changing environment

If the environment $E$ were fixed, individuals would clearly evolve to maximize their performance by specializing with $R$ as close to $E$, and $T$ as close to zero, as the evolutionary crossover and mutation processes allow. The more realistic case of interest here is when the environment $E$ varies randomly over time and the individual responses $R$ have to evolve to keep up with the changes, and their tolerances $T$ have to evolve so they become more generalist, performing as well as possible as the environment changes. The simplest representation of a randomly changing environment $E$ uses a standard first-order autoregressive model leading to an autocorrelated time series varying with simple additive noise

$$E_{t+1} = \alpha E_t + m\varepsilon_{t+1} \tag{2}$$

The autocorrelation parameter $\alpha$ specifies the extent to which the random changes persist over time, with $\alpha = 0$ corresponding to no persistence and a totally new random environmental deviation at each time step, and $\alpha = 1$ corresponding to complete persistence with the random changes building up over time as a random walk. In effect, individual random changes $m\varepsilon_t$ are introduced at each time step $t$ and decay exponentially over time with rate $1-\alpha$. Cleary it is important for artificial life models to explore the full range of possibilities for $\alpha$ if they are to provide ecologically realistic models, though negative values of $\alpha$ do not appear to correspond to real environments (Halley, 1996; Steel, 1985) so they will not be considered here. The noise term consists of a series of random variates $\varepsilon_t$, with a scale factor $m$ controlling the magnitude of change at each time-step. One obvious distribution to use for the noise would be an inverse power law as that is known to represent the individual changes in many kinds of natural event (Newman, 2005), but because many independent factors will generally be contributing to the environment and its changes, central limit theorem arguments indicate that it would be better here to use a

Gaussian distribution with zero mean and unit variance, as in many previous environment models (e.g., Heino, Ripa & Kaitala, 2000; Ripa & Lundberg, 1996; Wichmann et al., 2005).

In much of the existing literature, the issue of environmental change has been discussed within the framework of noise colour in inverse-power-law models rather than the autoregressive model of Equation 2 (e.g., Gilljam et al., 2019; Grove et al. 2020, 2021; Halley, 1996; Halley & Inchausti, 2004; Inchausti & Halley, 2003; Morales, 1999; Ruokolainen et al., 2009; Steele, 1985). One reason for this may be that environmental variability is most easily measured in practice as a randomly changing noise signal laid upon a stable (or predictable) underlying state, and that noise time series can be conveniently analyzed using spectral density decomposition which reveals that the power spectra can be approximated by inverse power laws (Vasseur & Yodzis, 2004). The exponent $\beta$ for each particular inverse power law is referred to as the noise colour, with the extreme persisting-change random-walk case corresponding to red/Brown noise with mostly low frequencies (exponent 2), and the opposite extreme case of rapidly-decaying-change corresponding to white noise with all frequencies contributing equally (exponent 0). So the exponent $\beta$ can be seen as an alternative measure of environmental change persistence.

One can carry out the same spectral density decomposition of the autoregressive time series produced by Equation 2 and find that they are also approximated by inverse power laws. The average relationship between $\alpha$ and $\beta$ is non-linear, but there is a monotonic increase from $\beta = 0$ for $\alpha = 0$ to $\beta \sim 1.8$ for $\alpha = 1$ (Grove et al., 2021). Nevertheless, the $\beta$-based inverse power law (IPL) time series and the $\alpha$-based autoregressive time series are known to have fundamentally different characteristics, and it is likely that neither fully reflect real environmental changes (Cuddington & Yodzis, 1999). Moreover, they have led to different model outcomes in the past (Cuddington & Yodzis, 1999), so it is not clear whether they are actually good enough approximations of each other that either would be suitable for use in simple models of the type studied here. Grove et al. (2020, 2021) chose to utilize $\beta$-based inverse power law time series to model their environment $E$, but the simpler autoregressive time series of Equation 2 has also been used in the past (e.g., Heino, Ripa & Kaitala, 2000; Ripa & Lundberg, 1996; Wichmann et al., 2005). To obtain reliable results in what follows, it is clear that both will need to be simulated and tested.

## 2.2 Implementing the environmental time series

Implementing the autoregressive time series simply requires generating random Gaussian numbers and adding them into the environment $E$ using Equation 2. The average root-mean-squared (rms) change then builds up over time at a rate depending on the persistence $\alpha$, as shown in the left graph of Figure 1 for representative values of $\alpha$ and fixed imposed change level $m = 1$ averaged over 1000 runs. Any $\alpha < 1$

corresponds to a decay of past changes that leads to an eventual leveling off of the overall change. As $\alpha$ increases above zero, the settling period is longer and the final stable cumulative change is higher. The extreme $\alpha = 1$ case is a true random walk that has the change diverging in proportion to the square root of the number of generations.

Implementing the inverse power law (IPL) time series is not so straightforward. The Fast Fourier Transform (FFT) of any real valued vector of length $N = 2^X$ representing a time series gives its frequency space representation consisting of a vector of complex numbers of the same length, the first half of which are the amplitudes and phases of its $N/2$ frequency components, and the second half is the complex conjugate of the first half in reverse order. So an IPL time series can be generated by creating its frequency representation and performing an Inverse FFT on that. That means the FFT amplitude for each component $k$ should vary according to an IPL with exponent $\beta/2$, namely $1/f_k^{\beta/2}$ where $f_k = k/N$, and then the randomness of the time series comes from introducing uniformly random phases $\theta_k \in [0, 2\pi]$ for each component. Thus the frequency space vector components need to be

$$V_k = \frac{e^{i\theta_k}}{f_k^{\beta/2}} \quad , \quad V_{N-k} = V_k{}^* \quad , \quad f_k = \frac{k}{N} \qquad \text{for} \quad k = 1, \dots, \frac{N}{2} - 1 \tag{3}$$

with $V_0 = V_{N/2} = 0$. Then the Inverse FFT of this gives the aimed-for IPL time series $E$ that is real valued and exactly follows the required IPL power spectrum with exponent $\beta$. If required, it is easy to introduce additional Gaussian randomness into the frequency space amplitudes, as has been done in some previous studies (e.g., Cuddington & Yodzis, 1999; Morales, 1999), but, in the interests of avoiding potential confounding factors, such deviations from the true IPL will not be adopted here. The normalization of the time series will be discussed in detail in Section 5, but for now it is simply scaled so that the rms change per time step is set equal to a parameter $m$ to match the autoregressive time series. The time series length $N$ needs to be a power of 2, and long enough for all potential simulations, so $N = 2^{22}$ was used.

One difficulty when comparing the effects of using IPL and autoregressive time series for $E$ is that the IPL time series specified by Equation 3 will not generally start from zero, and setting suitable initial values can be difficult (Halley, 2003). In principle, that is not a problem for the models here, because it is straightforward to start the individual responses $R$ at the same point as their environment $E$. For comparison purposes though, it is better to generate a much longer than needed time series and select the first subset of it that starts at a crossing of zero. While particular subsets of the full time series will no longer exactly follow the imposed IPL spectrum, on average they will. The average rms change $E$ then builds up over time in a similar manner to the autoregressive time series at a rate depending on the exponent $\beta$, as shown in the right graph of Figure 1 for representative values of $\beta$ and fixed imposed rms change $m = 1$ averaged over 1000 runs. The growth here only reaches a stable limit for $\beta < 1$, so long-term averages diverge for $\beta \geq 1$, which can be problematic for normalization purposes.

### 2.3 Extreme cases not requiring simulation

For either type of environmental time series $E$, if the individual responses $R$ do not change, higher cumulative changes $E$ will result in lower associated performance levels $P$ that depend on the change tolerance $T$, as shown for the IPL time series in the left graph of Figure 2. If the tolerances $T$ can evolve, then natural selection will easily find the peak in the curve for each persistence level $\alpha$ or $\beta$, leading to performances that fall with $\alpha$ or $\beta$, and tolerances that increase, as shown for the IPL time series in the right graph of Figure 2. If, as is often done, the time series are normalized to give equal average cumulative change (rather than equal average change per generation) across all $\alpha$ or $\beta$, the tolerance will simply match the chosen cumulative change and have no variation with $\alpha$ or $\beta$.

The other extreme scenario is when the responses $R$ perfectly track the changing environment $E$. In this case, the individuals only need a tolerance $T$ to deal with the average change during each generation. Then if the time series are normalized so that average change is equal across all $\alpha$ or $\beta$, any evolved tolerance will simply match the chosen change and have no variation with $\alpha$ or $\beta$. If the time series are instead normalized to have equal cumulative change across all $\alpha$ or $\beta$, the average change needed each generation to achieve that will fall with increasing $\alpha$ or $\beta$, and the tolerances will fall to match those, as shown for the IPL time series in the right graph of Figure 2.

These are the simplest non-trivial cases that represent the extremes of what can happen, and there is no need for any evolutionary simulations to understand them. What happens in the intermediate situations, in which individuals partially track their changing environment, will depend on many interacting factors, such as the potential rates of evolutionary change and the emerging population diversities, and full evolutionary simulations will be required to explore the outcomes for those cases.

### 2.4 Simulating evolving populations

The outcomes are difficult to predict when the individual responses $R$ to the environment $E$ co-evolve with the tolerances $T$. This is where explicit evolutionary simulations are required, particularly because it is difficult to know how the natural random variability between individuals and across simulation runs affect what evolves. The simulations of Grove et al. (2020, 2021) allowed $R$ and $T$ to co-evolve, but for $E$ they only used the IPL time series normalized to give equal average magnitude for all persistence levels $\beta$. This is a sensible starting point, but, because that normalization combines the effects of the individual change magnitudes and their persistence, and past studies have found differences depending on whether an IPL or autoregressive times series is used, it could potentially lead to misleading results. It also renders it impossible to explore the relative importance of the change magnitude and persistence. Consequently, a more exhaustive approach is adopted here that allows an exploration of the separate effects of the two environmental change parameters for both types of time series. To do this, a range of change magnitude values $m$ are tested for each of a range of autocorrelations $\alpha$ and IPL exponents $\beta$. To be specific, $m$ is varied across values from the logarithmically distributed set {0.01, 0.032, 0.1, 0.32, 1.0, 3.2, 10.0} and either $\alpha$ is varied across the linearly distributed set {0, 0.25. 0.5, 0.75, 1.0} or $\beta$ is varied across the linearly distributed set {0, 0.5. 1.0, 1.5, 2.0}.

As usual, simulating the evolutionary process involves maintaining a population of individuals each represented by their key innate parameters, which in this case are simply their response $R$ to the current environment $E$, and their tolerance to change $T$. Because this study only concerns cases in which there is no lifetime learning, a simple generational approach (Eiben & Smith, 2016) will suffice, which is much simpler and more computationally efficient than the steady state approach of Bullinaria (2018, 2022) that is necessary when learning can improve performance during an individual's lifetime. This is conveniently implemented by the fairly standard process of replacing the worst performing half of the population at each generation with children of random pairs of the best half. That does not lead to the strongest possible selection pressure, but it is simple and sufficient for good performance to evolve quickly. (A variation of this will be tested in Section 6 to see how this choice affects the outcomes.) The crossover and mutation processes for creating the children follow Bullinaria (2018, 2022) with each child's innate parameter values chosen randomly from the ranges spanned by their two parents, and mutations adding random constants so the inherited values can fall outside the parental ranges. It is generally found in models as simple as those here that having the children inherit their innate parameters from only one parent with random mutations leads to similar qualitative steady-state outcomes, though in more complex biologically realistic models there can be significantly different outcomes for asexual rather than sexual reproduction that also depend on how the innate parameters and crossovers are implemented.

### 2.5 Setting or evolving the mutation rates

Because mutations are needed to enable $R$ to track the changing environment $E$, and the rate at which that occurs will affect the optimal evolved tolerances $T$, and the mutation rates also affect the diversity of the evolving population, achieving the best evolved performance requires the sizes of the mutations to be set carefully and appropriately for the given pattern of environmental change. This is further complicated by the fact that the best mutation rates also depend on the evolutionary selection pressures, which are fixed in the simplified simulations here, but will co-evolve themselves in real populations. The simplest way to handle all this reliably is to allow the two mutation rates $M_R$ and $M_T$ themselves to co-evolve with $R$ and $T$. That still leaves suitable mutation rates to be set for the evolving $M_R$ and $M_T$, but they are less crucial as they are largely independent of the environmental change details, and sub-optimal values can easily be compensated for by simply allowing the simulations to run for plenty enough generations that they always settle into a stable state. For the current study, it is not necessary to specify nor model the details of how the mutation rate evolution might occur in practice, so it is sufficient to simply treat the mutation rates as further parameters that are free to be optimized by the natural selection. Consequently, the first set of simulations (presented in Section 3) will involve each individual in the population having four evolvable innate parameters $\{R, T, M_R, M_T\}$.

The evolution of mutation rates and evolvability is actually another important topic in the field of environmental change that has received considerable interest recently (e.g., André & Godelle, 2006; Good & Desai, 2016; Johnson, 1999; Lynch, 2011; Metzgar, 2007; Ogbunugafor, 2023; Romero-Mujalli & Tiedemann, 2021; Sniegowski et al., 2000; Wei et al., 2022), and the simulations of this study have been designed to cast light on this matter too. There is no doubt that real species do evolve over time in response to changing environmental pressures, and that the evolutionary changes ultimately derive from mutations (Metzgar, 2007). However, there is an associated trade-off that is known to be important for both real and artificial populations (Wilke et al., 2001). Large mutation rates will facilitate the tracking of environmental change, but having them too large will lead to too many individuals with poor performance. The evolutionary simulations in Section 3 will show how the selection pressures related to those two factors can be balanced to result in appropriate compromise mutation rates emerging for each particular environment.

An important issue here, though, is that in biological populations there will generally be more complex limitations on how quickly the species can mutate, and it is the interaction of this rate of evolutionary change with the rate of environmental change that is a key factor affecting the performance of real species in real environments. Biological mutations primarily come from genetic factors, such as DNA replication errors, and are predominately deleterious, with few of them resulting in useful variation in particular high level traits such as tolerance to environmental variation (Eyre-Walker & Keightley, 2007; Metzgar, 2007; Sniegowski et al., 2000). This implies a rebalancing of the selection pressures associated with high mutation rates, casts doubt on the reliability of the results obtained by simulating the mutation rate evolution as described above, and suggests the simulations may over-estimate the evolved mutation rates.

On the other hand, real species have evolved various mechanisms to deal effectively with random mutations (Good & Desai, 2016), and do cope with considerable diversity in their high level abilities, possibly as a consequence of the complex indirect relation between the genome and numerous inter-related high level abilities. Moreover, it is possible that real mutations may not be totally random, but occur with increased evolved rates targeting only the portions of the genome where they are most needed (Martincorena & Luscombe, 2013). There will also inevitably be physiological costs involved in reducing the mutation rates to very low levels (Sniegowski et al., 2000), and random genetic drift may ultimately be responsible for the effective lower bounds (Lynch, 2011). These issues suggest that the evolutionary simulations could sometimes under-estimate the mutation rates. It is also known that factors enabling increased mutation rates can hitchhike on the effect of beneficial mutations, but at different rates depending on whether the reproduction is sexual or asexual (André & Godelle, 2006; Johnson, 1999), and it is not clear how realistically such mechanisms might be captured in the simulations.

So, to explore the potential effects of all these complications in evolving the key mutation rates in the simulations, a second set of simulations (presented in Section 4) will be run with each individual having only two evolvable innate parameters $\{R, T\}$ to understand what happens when the mutation rates $M_R$ and $M_T$, and hence the evolution rates, are fixed at a range of values that are larger or smaller than the baseline "optimal" values found by simulating their evolution in Section 3. This still leaves the problem of identifying appropriate mutation rates for each real species being modeled, but simulating a wide range of fixed values will show the consequences of all the options.

### 2.6 Running the simulations

Running all the evolutionary simulations, with either two or four evolving parameters, follows a standard and straightforward process – it is simply a matter of updating an array of two or four numbers for each individual for each new generation according to the specification above, and re-computing the individual performance levels after each environmental change. In each case, the simulated environment is updated once per generation, with each simulation run using a fixed population size of 1000 individuals for 100,000 generations, which proved to be many more than needed for all the runs to stabilize before their half way point for even the most extreme parameter value variations tested and severely sub-optimal starting points for the evolving parameters. All the evolutionary outcomes are recorded once every 500 generations which provides data points with a good degree of independence, and averaged over the stable final half of each run which are sufficiently long to provide accurate means despite the random nature of the simulated environmental changes. Then to achieve reliable statistics, each simulation is run 50 times using different random number seeds, with means and standard deviations computed and plotted for all the outcomes. This results in error bars on the graphs corresponding to approximately seven standard errors, so any non-overlapping error bars indicate differences of high statistical significance. The only exception is for the study of extinction timescales in Section 6 and the associated graphs in Figures 9 and 10. Here the variance across runs is large due to their direct dependence on rare random environmental changes, so 1000 runs are carried out for each scenario, which enables suitably small standard errors to be plotted, again indicating high statistical significance.

## 3. Simulations with Evolving Mutation Rates

The first set of simulations allowed the four key innate parameters$\{R, T, M_R, M_T\}$ to evolve for a range of values of the imposed change magnitude $m$ and persistence $\alpha$ or $\beta$ as described above. It soon became apparent that presenting the evolutionary outcomes was complicated by the fact that the effective magnitude of change that drives the need for a non-zero change tolerance $T$ is not the imposed change magnitude $m$, but how far the current environment $E$ has deviated from the current response $R$, and that will clearly depend on the persistence $\alpha$ or $\beta$ of the changes as well as their magnitude $m$. The simplest way to quantify the effective change magnitude $m_{eff}$ for each simulation is as the mean absolute difference between the responses $R$ and the associated environment $E$ during the stable final half of each run. Plotting the evolutionary outcomes against this $m_{eff}$ has the additional advantage of allowing the dependence on both $m$ and $\alpha$ or $\beta$ to be observed on the same graphs, so this was adopted here. (Using rms difference rather than absolute difference was found to lead to qualitatively similar results.)

The key simulation outcomes as a function of change magnitude $m_{eff}$ are presented in Figure 3, which show that the differences between the IPL and autoregressive time series results are actually rather small in this case. The evolved tolerance $T$ (top-left) increases approximately linearly (on the log-log plot) with $m_{eff}$, so that $T$ is consistently between about 0.6 and 0.8 of $m_{eff}$. This is not surprising given that larger tolerances are necessary to cope with larger discrepancies between the individuals responses and their current environment, and the evolutionary natural selection is able to optimize that. As the tolerances increase linearly, the corresponding performance levels $P$ (top-right) fall linearly in line with their definition in Equation 1, which is what provides the selection pressure to prevent the tolerances rising larger than is cost effective. As the change magnitude is decreased towards zero, both the tolerance $T$ and response mutation rate $M_R$ are seen to fall linearly to zero, and the corresponding performance $P$ rises towards infinity. The limiting case at zero change is achieved because the population has fully evolved to perfectly match the fixed environment, with no need for any tolerance or further mutations.

Each cluster of points in Figure 3 corresponds to a particular value of imposed change magnitude $m$, and follows the same pattern in each case, as one would expect given that all the relevant parameters are free to optimize themselves on the same arbitrary scale. The consequences of breaking that scale invariance will be apparent in the fixed-mutation-rate simulations presented in the next section. Within each cluster there is a clear dependence on the change persistence $\alpha$ or $\beta$. This is most pronounced in the plots of the evolved mutation rates as a function of $m_{eff}$. For each persistence $\alpha$ or $\beta$, the evolved response mutation rate $M_R$ (bottom-left) increases approximately linearly with the change magnitude (both $m$ and $m_{eff}$). This is again to be expected given the scale invariance in the representation of the environments and responses, and the response mutation rate can compensate for any changes in the environmental scale. The tolerance mutation rates $M_T$ (bottom-right) vary much less and more noisily with the environmental change parameters. Unlike the response $R$ which needs to adjust continuously throughout each run in line with the environment changes, the tolerance $T$ should eventually settle down and not need to change as long as the average magnitude of change is constant, so it is not surprising that the associated mutation rate should end up very low and largely independent of the change magnitude. The effective lower bound in this evolution towards zero for $M_T$ (and $M_R$ at $\alpha$ or $\beta = 0$, or $m = 0$) is the one aspect where details of the associated mutation rates and crossover mechanism do affect these models.

The main finding here is that the environmental change magnitude has a much greater effect on the outcomes than the change persistence. However, it is still important to understand how varying the persistence $\alpha$ or $\beta$ affects the outcomes for a given imposed magnitude of change $m$. This aspect of the simulation results is shown in more detail for the representative case of $m = 1$ in Figure 4. The most striking feature is how flat the plots of $T$, $m_{eff}$ and $P$ are against $\alpha$ or $\beta$. This is very different to the large monotonic decrease in tolerance and large monotonic increase in fitness (i.e. performance) with increasing change exponent $\beta$ that was observed in the earlier simulations of Grove et al. (2020, 2021). There are two potential reasons for this difference. One (which will be explored in Section 4) is that the earlier study did not allow the mutation rates to co-evolve with the other parameters to facilitate the evolutionary response to the changes. The other (which will be explored in Section 5) concerns the way in which the environmental time series were normalized in the earlier study, because that effectively combined the effects of persistence and magnitude so that higher persistence led to lower imposed change magnitude which has been seen here to be the dominant effect lowering the tolerance.

The differences between the IPL and autoregressive cases observed in Figure 4 are not as significant as they appear. First, the relationship between $\alpha$ and the associated IPL exponents $\beta$ in the autoregressive case is non-linear, with $\alpha = 1$ actually corresponding to $\beta \sim 1.8$. This means the $\alpha$ scales on the graphs should really be shifted non-linearly by around 10% to the left to accurately match the $\beta$ scales. Also, the imposed change $m$ in the autoregressive case only matches the change per generation $m$ in the IPL case for $\alpha = 1$. For lower values of $\alpha$ there will also be a contributing change per generation arising from the change decay term $(1-\alpha)E$, meaning that the change $m_{eff}$, tolerance $T$ and performance $P$ values for the autoregressive case would all need to fall by corresponding amounts for the scales to match. Taking these factors into account, the IPL and autoregressive results here are remarkably similar.

A natural expectation is that the optimal response $R$ mutation rates $M_R$ will increase with $\alpha$ or $\beta$ because increasing the environmental change persistence inevitably increases the cumulative change that the mutations are required to compensate for (as seen in Figure 1). Figure 4 (bottom-right) confirms that the simulations do show a clear monotonic increase in the average evolved mutation rate $M_R$ with $\alpha$ and $\beta$, and very little variation across runs, in line with the greater build up of environmental changes that needs to be tracked. It is the resulting successful tracking of the environment that allows the variation of the average effective change magnitudes $m_{eff}$ (bottom-left) and the average individual performance levels $P$ (top-right) to be kept fairly flat against $\alpha$ and $\beta$, though there is noticeable variation across runs arising from the random nature of the environmental changes. The associated evolved tolerance $T$ (top-left) is also fairly flat against $\alpha$ and $\beta$, but there is an interesting and significant decrease as the persistence increases from 0 to about 0.1, before settling into a slow monotonic increase after that. This reflects the trade-off between the effect of the two appearances of $T$ in the specification of the performance measure $P$ (in Equation 1) that drives the evolutionary selection pressures. Interestingly, when the persistence $\alpha$ or $\beta$ is higher, more longer term variations in the environment will arise and it will often prove worthwhile for individuals in those cases to adjust their tolerances more to track them. This is what underlies the variation in tolerance mutation rates $M_T$ seen for each value of $m$ in Figure 3 (bottom-right).

In summary, the simulations here have shown that the main driver of the evolved environmental change tolerances and associated performance levels is the mean effective rate of environmental change, with approximately linear relationships on the log-log plots as long as the mutation rates are allowed to co-evolve. What happens when the mutation rates cannot, for some practical biological reason, take on the optimal values achieved in these simulations remains to be explored in the next section.

## 4. Simulations with Fixed Mutation Rates

The simulations so far have involved mutation rates that are free to be evolutionarily optimized so that the populations are best able to evolve to cope with their changing environments. In practice, the selection pressures and reproductive cross-over mechanisms will also affect the rate at which a population can evolve to track its changing environment, but simply varying the mutation rates as a proxy for all such factors is a sufficient approximation for this simplified study. This mutation rate evolution provides an "optimal" baseline performance, and it is certainly true that rates of mutation have evolved to vary across real species (Bromham, 2009), but it leaves the question of what happens when the rate of evolution (determined here by the mutation rates) is, for some practical reason or another, unable to adapt to deal optimally with the imposed pattern of environmental changes.

Fortunately, the simulations of the previous section can easily be repeated with the mutation rates $M_R$ and $M_T$ fixed at any chosen values. In this way, cases can be modeled which have the mutation rates fixed at values which vary from being too large to too small as the applied environmental change magnitude and persistence levels are varied across the same ranges as before. Specifically, $M_R$ was set at 0.1 and $M_T$ was set at 0.01, which are round numbers near the centers of the mutation rate graphs of Figure 3. The key evolutionary outcomes in this case are plotted in Figure 5 using the same scales for comparison with the evolving mutation rate case of Figure 3. Similar approximately linear relations exist here for the average evolved tolerances $T$ (left) and associated performance levels $P$ (right) with respect to the effective change magnitudes $m_{eff}$. However, the relation is somewhat noisier here than in Figure 3, the effective change $m_{eff}$ corresponding to each imposed magnitude $m$ is larger, and the resulting performance levels are lower (with some data points now actually falling off the graphs).

The effect of having sub-optimal mutation rates can be seen more clearly in the plots of the outcomes against persistence $\alpha$ or $\beta$ presented in Figure 6. For the case of imposed change $m = 0.01$ (left), for which the mutation rate $M_R$ is larger than optimal for all $\alpha$ or $\beta$, tracking the environmental changes is never a problem, and the outcomes for $T$, $m_{eff}$ and $P$ are fairly flat against $\alpha$ or $\beta$. This is similar to the pattern of outcomes for the evolved mutation rate case presented in Figure 4, but having the mutation rate $M_R$ too large means the response $R$ cannot be fine-tuned as well as it was there. The larger than necessary mutations introduce more variation into the populations and that has a negative effect on the evolved performance averages. For the larger change magnitude case of $m = 1$ (right), the fixed mutation rate $M_R$ becomes increasingly smaller than optimal as $\alpha$ or $\beta/2$ rises above ~0.12, and tracking the environmental changes becomes a problem when the persistence is large enough to allow the changes to build up too quickly over time. As $\alpha$ tends to 1 in this case, the effective change magnitude $m_{eff}$ becomes increasingly large and the performance levels quickly drop towards zero. There are similar, though less sharp, trends as $\beta/2$ tends to 1. So this is one important situation in which the persistence $\alpha$ or $\beta$ has a big effect on what emerges, and there are also significant differences between the IPL and autoregressive simulation outcomes, though increasing both forms of persistence here still affects the outcomes $T$ and $P$ in the opposite direction to that found in the earlier simulations of Grove et al. (2020, 2021).

It is clear that the key indicator for the outcomes of each case simulated in this study is the effective change magnitude $m_{eff}$ as that has a direct approximately-liner effect on the tolerances and performance levels that emerge. Consequently, knowing how $m_{eff}$ varies with the imposed change magnitude $m$ in each case is crucial for understanding everything else that emerges from the simulations. This relationship is plotted for four representative IPL cases on the left graph of Figure 7. When the mutation rates are free to evolve, there is a scale invariance and a simple linear relationship emerges between $m$ and $m_{eff}$, with only the small dependence on persistence $\beta$ previously seen in Figure 4. Fixing the mutation rates at intermediate values breaks the scale invariance and leads to more complex outcomes. For low $m$, the fixed mutation rate of $M_R = 0.1$ is higher than optimal for all $\beta$, so $m_{eff}$ is larger than in the corresponding evolved mutation rate cases, which in turn leads to larger tolerances $T$ and lower performance levels $P$ emerging. For larger values of $m$, the persistence $\beta$ has a bigger effect on the outcomes. For $\beta = 0$, the fixed mutation rate is close enough to optimal for the higher values of $m$ that the emergent values of $m_{eff}$ are very close to those of the corresponding evolved mutation rate case. However, as $\beta$ increases, the fixed mutation rate $M_R$ becomes increasing unable to keep up with the buildup of changes, leading to significantly larger values of $m_{eff}$ with the associated rise of $T$ and fall of $P$ seen in Figure 6.

The effect of the response mutation rate $M_R$ on the evolutionary outcomes can be clarified further by investigating how the performances $P$ themselves vary for a particular value of persistence $\beta$ with more extreme fixed values of $M_R$. The right graph of Figure 7 shows how the outcomes for $\beta = 2$ vary across fixed $M_R \in \{0.001, 0.1, 10.0\}$ compared with the evolved $M_R$ values of Figure 4. For each value of the imposed change magnitude $m$, the evolved mutation rates naturally result in best performance. The fixed mutation rate $M_R = 0.001$ is far too low and leads to much worse performance for all $m$ shown. The mutation rate $M_R = 0.1$ is near optimal for $m \simeq 0.06$, but leads to worse performance for lower $m$ because it is too large, and worse performance for higher $m$ because it is too small. Finally, the mutation rate $M_R = 10$ comes close to optimal for $m \simeq 6$, and again leads to worse performance for lower $m$ because it is too large, and for higher $m$ because it is too small.

## 5. Normalizing the Change Magnitude

Perhaps the most notable aspect of the simulation results in this study so far is that the evolved tolerances $T$ and associated performance levels $P$ vary with the environmental change persistence $\alpha$ or $\beta$ in the opposite direction to that found in the simulations of the original Grove et al. (2020, 2021) study. However, those earlier simulations normalized their environmental noise time series $E$ by adjusting the imposed change magnitude to give an equal average variance across all values of $\beta$ in a manner justified by reference to the study of Wichmann et al. (2005) that investigated how best to apply environmental variance scaling. The simulations presented above have deliberately treated the imposed magnitude $m$ and persistence $\alpha$ or $\beta$ as independent aspects of the environmental noise and not carried out any normalization or scaling as that is known to confuse those two aspects (e.g., Heino et al., 2000). Nevertheless, it is worth checking the effect of persistence for some kind of fixed cumulative magnitude, rather than fixed imposed magnitude, and informative to attempt a replication of the reversed pattern of results found by Grove et al. (2020, 2021), so the consequences of change normalization will now be explored here.

The main practical difficulty with carrying out the normalization is that there are numerous reasonable ways to define the associated change magnitude that is fixed, with many depending on the simulation timescales involved, and the choice can make significant qualitative differences to the modeling outcomes (Heino et al., 2000; Wichmann et al.,2005). Perhaps the most natural normalization approach for the simulations of this study would be to scale to the expected asymptotic variance, as has been done many times in the past (e.g., Cuddington & Yodzis, 1999; Petchey et al., 1997; Ripa & Lundberg, 1996). This involves adjusting the imposed magnitude $m$ for each $\alpha$ or $\beta$ so that in graphs of the average buildup over time of the rms environmental change (like in Figure 1) the cumulative change endpoints coincide for all values of $\alpha$ or $\beta$. The problem with this is that it relies on all the simulations running long enough to reach a stable endpoint, and for $\alpha = 1$ or $\beta \geq 1$ that can never happen. The obvious alternative would be to instead adjust $m$ to normalize the average rms environmental change over the chosen simulation length, but doing that is also problematic in that it leads to outcomes that depend on what are essentially arbitrary simulation lengths. All these normalization difficulties can be avoided by not normalizing at all, as has been done so far in this study. Nevertheless, it is worth presenting one representative set of results to illustrate how the above simulations would be affected by normalization.

It is evident from Figure 1 that scaling to the expected asymptotic variance can give reliable run-length-independent outcomes for the autoregressive simulations with $\alpha < 0.999$, so those are the cases that will be studied here. Interpolating the outcome plots presented above, or re-running the simulations using the required normalizing imposed change magnitude $m$ for each $\alpha$, gives the outcomes for normalized change magnitudes shown in Figure 8. It is not difficult to understand what the normalization has changed. Figure 1 (left) shows that the reductions in $m$ necessary to achieve normalization are modest for low $\alpha$ (< 0.5) but increase rapidly as $\alpha$ tends to 1, so the autoregressive graphs of Figures 4 and 6 are little changed for low $\alpha$, but there are massive decreases in tolerance $T$ and increases in performance $P$ for high $\alpha$. Populations with evolving mutation rates naturally perform better than the fixed mutation rate populations, and the detailed outcomes for the fixed mutation rate cases depend on the mutation rates used (as they did in Figure 6), but the general trend of $T$ falling and $P$ rising with persistence $\alpha$ is now in line with the earlier simulations of Grove et al. (2020, 2021). This reversed trend is not surprising – for higher $\alpha$, the magnitude $m$ of each imposed change needs to be smaller to give the same cumulative change, which means the changes are easier to track, resulting in the need for smaller tolerances which in turn lead to better performances. Unfortunately, this simple explanation totally obscures the complexity of the outcomes revealed by the un-normalized simulations. Using different normalization criteria will lead to different qualitative outcomes, and larger mutation rates for $T$ would cause the falls in average $T$ and associated rises in average $P$ to level off sooner, but the general trends here for $T$ and $P$, and their causes, remain the same.

The conclusion here is that it is important to be clear about what the simulations are aiming to demonstrate, because carrying out environmental change normalization in different ways, rather than not at all, can lead to completely different dependences of the evolutionary outcomes $T$ and $P$ on the imposed change persistence $\alpha$ or $\beta$, and obscure the true complexity of the outcomes. However, as long as the change magnitude $m$ and persistence $\alpha$ or $\beta$ are varied independently, as has been done throughout this study, it will always be straightforward to read off any required model outcomes by interpolating from the simulation results produced.

## 6. Modeling Extinction Events

Of course, the aim of artificial life models is not just to produce graphs of evolutionary outcomes, but to aid our understanding of what has happened, or what might happen in the future, for real populations. Usually, the key worry about environmental change is population extinction or extirpation, and evolutionary models of this are not new (e.g., Bürger & Lynch, 1995; Yedid, Ofria & Lenski, 2009). Intuitively, higher environmental change persistence suggests that a "bad" change will last more generations and hence lead to a higher extinction risk, but what happens in practice is actually more complex than that. The relationship between extinction and change persistence has already attracted considerable interest, but earlier studies have led to conflicting conclusions (e.g., Cuddington & Yodzis, 1999; Halley & Kunin, 1999; Heino et al., 2000; Inchausti & Halley, 2003; Johst & Wissel. 1997; Kamenev, Meerson & Shklovskii, 2008; Morales, 1999; Petchey, Gonzalez & Wilson, 1997; Ripa & Lundberg, 1996; Wichmann et al., 2005), and the hope is that models of the type studied in this paper can clarify matters.

When there is only one species being modeled, the evolutionary simulations can continue to run even if the performance of all individuals in the population drop to very low levels. However, if the simulated species really needs to compete with other species that can cope more effectively with the environmental changes, they are likely to be driven by them to extinction. In fact, even if there are no competing species, there will still generally be minimum performance levels required to survive in a given environment, related to matters such as food gathering or the ability to keep warm or cool enough. So, in agent-based simulations like those of this study, if all the individual performances $P$ drop below a certain level, the population should really be regarded as unable to survive and an extinction event recorded. In previous simulations of populations that are able to use lifetime learning to deal with changing environments, interesting relationships were found between the evolved learning strategies and the average length of time the populations survive before suffering an extinction event (Bullinaria, 2018, 2022). The aim now is to explore the factors that affect the times to extinction in the models studied above.

Modeling survival times for evolving populations is not as straightforward as it sounds. One cannot, for example, simply run the simulations as described above and check at each time-step whether the individuals exceed a certain minimum performance level. The problem is that real populations emerge from earlier viable populations, whereas the simulations start from random initial states that are far from optimal, and the extinction events tend to occur well before the evolved populations of interest emerge. The extinction timescales can therefore easily end up depending more on how close the random initial population is to the near-optimal evolved population than on the factors of interest such the environmental change persistence, and this is one issue likely to be behind the conflicting results of earlier studies.

The kind of extinction events normally of interest are when the population has already evolved and is surviving well in its current environment, but a particular environmental change, or sequence of changes, cause the performance levels to plummet to below that required to survive. This would, for example, happen when an unusually large change, or combination of changes, takes the environment far outside the population's evolved tolerance range. With Gaussian distributions for the random variations, this kind of extinction event will be inevitable at some point. The question is how frequently does this happen for different levels of change magnitude and persistence, and for different mutation rates and minimum performance levels. So, rather than running the extinction simulations from random initial states, it is appropriate to start them from the evolutionary endpoints found above. The Gaussian distribution of changes will still lead to large variances in the survival times, but with enough simulation runs it will be possible to determine reliable mean survival times.

Given the arbitrariness of the magnitude scale in the simulations, the survival times are most easily investigated by keeping the imposed change magnitude or normalized cumulative change magnitude fixed and varying the mutation rates. Then it is just a matter of running the simulations starting with the evolved tolerances and mutation rates at the ends of the earlier simulations, and recording the number of generations that pass before the performance level drops below a particular minimum value $P_{min}$. For reasons noted above, the variances in the times to extinction are large, often of a similar magnitude to the means, so means and standard errors were computed from 1000 independent runs.

Simulations were first carried out for the autoregressive environmental time series to give the results presented in Figure 9. It is immediately noticeable that the extinction timescales are usually considerably less that the 100,000 generations used to generate the main simulation results, even though they start with near-optimal evolved parameters. Imposed magnitude $m = 1$ cases are shown in the top two graphs, with extinction taken to have occurred when the performance of all individuals in the population falls below $P_{min} = 0.001$ (left) and $P_{min} = 0.00001$ (right). The number of generations to extinction obviously rises as the chosen minimum required performance level $P_{min}$ is reduced, but the general pattern of outcomes is found to take the same form for any reasonable choice of $P_{min}$. The "optimal" evolved case exhibits extinction times increasing with change persistence $\alpha$ in line with the increasing diversity resulting from the increasing mutation rates $M_R$ seen in Figure 4, and there is also a dip for low values of $\alpha$ in line with the dip in tolerances $T$ seen in Figure 4. So this pattern of results is understandable.

The effect of having different fixed mutation rates also follows understandable trends. Generally, when the mutation rates are less than the near-optimal evolved values, the diversity will be lower and the times to extinction will be shorter, and when they are higher, the diversity will be higher and the times to extinction will be longer. So given the rise in optimal mutation rates with persistence $\alpha$ seen in Figure 4, a fall in times to extinction with $\alpha$ is expected when the mutation rates are reasonably high. However, when the mutation rates are so low that the effective change magnitudes $m_{eff}$ and tolerances $T$ diverge, as seen on the right graphs of Figure 6, the large tolerances then render extinction rare. Whether such very high evolved tolerances $T$ and associated low performances $P$ really do correspond to low extinction rates, rather than a totally uninhabitable environment, is something that will depend on the species and environment being modeled.

It is noticeable that the fully evolved populations are often more prone to extinction than the corresponding fixed-mutation-rate populations, even though natural selection has given them higher performance levels. A similar situation was found in earlier models which looked at the evolution of lifetime learning strategies in changing environments (Bullinaria, 2018, 2022) – evolution led to strategies that had the best performance levels most of the time but led to much shorter times to extinction. Clearly, evolution cannot optimize a species to avoid extinction when there are no individuals left to breed after the first extinction event, but there may be other selection pressures that can help avoid extinctions. One effective way to explore the evolutionary pressures involved is to vary the selection process from that used in the main simulations above. Instead of taking the best performing half of the population to survive and produce children to replace the others, one can require individuals to have a certain minimum performance level for them to survive and procreate, while making sure that at least the best two individuals survive to create each new generation. Repeating the above simulations doing this leads to similar patterns of results, but with larger evolved mutation rates $M_R$ (particularly for low change persistence $\alpha$) so there is more diversity in the population to cope with larger changes, and the evolved tolerances $T$ are slightly larger too, which also aids coping with larger changes. This in turn leads to the larger times to extinction seen in all the graphs of Figure 9 for the "Evolved (min $P$)" lines for all $\alpha$, which shows that getting the selection pressure right is another important factor here. In fact, in addition to selection pressure, there are various other life history traits (such as reproduction timescales, offspring numbers and lifespans) that will generally co-evolve with the tolerances and mutation rates to optimize the outcomes in changing environments of different types and timescales (Roff, 2002; Stearns, 1992), and these factors will need to be accommodated in more realistic species-specific simulations.

Extinction outcomes for the normalized cumulative change magnitude approach studied in Section 5 are shown for $P_{min} = 0.001$ in the bottom-left graph of Figure 9 and are also easily understood. For the evolved mutation rate cases, natural selection keeps the mutation rates and tolerances near optimal independent of the normalization factor, so the outcomes here are very little changed from the un-normalized case above. For the fixed mutation rate cases, the sharp rise in the average performance for high change persistence $\alpha$ seen in Figure 8 now occurs for all mutation rates and results in the longer times to extinction seen for higher $\alpha$ here. However, there are known difficulties with using normalized time series to model extinction risk (Heino, et al., 2000), not least because the effects of change magnitude and persistence become confused. As argued above, it is probably best to not normalize at all, but to model the magnitude and persistence as independent parameters.

The final graph of Figure 9 (bottom-right) shows how the average extinction times typically vary with the imposed change magnitude $m$, specifically for fixed intermediate persistence $\alpha = 0.5$ and $P_{min} = 0.001$. The general decrease seen in the times to extinction as the environmental variability increases is consistent with intuition, empirical real-world data, and many previous models (Inchausti & Halley, 2003). For the evolved mutation rate cases, the relatively slow fall of times to extinction with increasing magnitude are in line with the expectation that bigger changes are more likely to cause problems for the population and the average performance fall with imposed magnitude (as seen in Figure 3) brings the population closer to the performance cutoff point $P_{min}$. Similarly, the increased population diversity and associated ability to track larger magnitude changes for larger fixed mutation rates explains the relation between the results across the fixed mutation rate cases. It is also evident that the evolutionary drive for better individual performance leads to evolved mutation rates for low change magnitudes that limit the population diversity so much that it leads to increased risk of extinction over all the fixed mutation rate cases, though evolved mutation rates do lead to longer times to extinction for large change magnitudes.

Finally, all that remains is to check how having an inverse power law, rather than autoregressive, environmental time series affects the outcomes. Figure 10 shows the two key graphs for comparison with Figure 9. Contrary to earlier suggestions (e.g., Cuddington & Yodzis, 1999), the type of time series has surprisingly little effect on the outcomes. A precise match between the extinction times for the two cases is not to be expected, because of the previously-noted non-linear relations between the change persistence parameters $\alpha$ and $\beta$ and between the change magnitude definitions, but the general pattern of the outcomes is clearly the same in both cases. The only practical difference is in how the times to extinction increase for high persistence when the mutation rates are low, which can be traced back to the differences in evolved tolerances seen in the top-right graph of Figure 6.

## 7. Conclusions and Discussion

The study described in this paper has carried out an extensive series of agent-based simulations using variations of a simplified model originally developed by Grove et al. (2020, 2021) that represents randomly changing environments by a single parameter, and evolves populations of individuals with responses and tolerances to that environment also represented by single parameters. It focused on the effect of change persistence, as that has attracted considerable recent interest, but also investigated other factors that may potentially affect the outcomes. Throughout, intuitive explanations of the evolutionary outcomes have been identified that are not limited by the particular choices made for the various parameter values or other simulation details, and can therefore be applied to future scenarios without recourse to further simulations. The main conclusion is that, despite the simplicity of the model, the outcomes are complex and depend crucially on apparently minor details such as the chosen mutation rates and normalization approaches. Other factors, such as the choice between autoregressive and inverse power law environmental time series, have been found to have much less effect on the outcomes than previously suggested (e.g., Cuddington & Yodzis, 1999). Generally, it is found that the average magnitude of the environmental changes has the biggest effect on the evolutionary outcomes, and that swamps the effects of change persistence however that is modeled. However, this does not mean that the environmental change persistence or noise colour is not important, because it clearly does have a big effect on the average effective change magnitude.

In the process of developing the artificial life simulations here to better understand the processes going on in real populations, it was found that particular care was needed when taking the environmental change persistence into account. This is because considerable divergences in outcomes emerge depending on how exactly it is done, as seen in the differences between the un-normalized simulation results of Sections 3 and 4 of this paper and the normalized simulations of Section 5 of this paper and Grove et al. (2020, 2021). Indeed, it appears that the importance of change persistence may have been over-estimated in previous studies because they have attempted to normalize the environmental time series in ways which confuse the effects of change magnitude and persistence. This has led to the recommendation here that it is best to explore the separate effects of magnitude and persistence without carrying out any normalization at all. Independently varying the magnitude and persistence certainly makes it easier to interpolate other conclusions from the same results without the need to run further simulations. In fact, that approach can be used to explore the effect of different forms of magnitude normalization as discussed in Section 5. This is just one aspect of a much more general issue of how best to model the factors affecting performance and extinction rates in real changing environments using a single environmental time series. For example, the optimal simulation specification will clearly depend on the details of the empirical real-world data being modeled, and using measured environmental time series may be more appropriate than using either the pure autoregressive or pure inverse power law time series that are usually studied.

The simulations presented in this paper have also clarified the relation between the environmental change magnitude $m$ imposed at each time-step and the effective change magnitude $m_{eff}$ defined as the mean absolute difference between the current individual responses $R$ and environment $E$. It is the effective change magnitude $m_{eff}$ that needs to be tracked by the evolving response $R$, or accommodated by an evolved tolerance $T$, so it has been argued that $m_{eff}$ is a more useful quantity for understanding how populations cope in changing environments than the imposed change magnitude $m$ or persistence $\alpha$ or $\beta$. For example, an individual's response will be optimal for a particular environment, and it is how far that environment deviates from the current environment that is more important than how far the current environment has deviated from the environment at some particular point in the evolutionary past. The clear linear relations between both $T$ and $P$ and $m_{eff}$ in Figures 3 and 5 support this idea.

A related issue explored in more detail in this paper than earlier studies is the importance of setting appropriate mutation rates in the simulations. Optimal values for those mutation rates can be established by treating them as additional evolvable parameters, as in the models of Section 3. However, as discussed in Section 2.5, there are numerous complications concerning mutation rate evolution in real biological populations which suggest the evolved rates found in those simulations could easily under- or over-estimate the values appropriate for use in reliable models of real species. For example, the hitchhiking of increased mutation rates on beneficial mutations is only really effective in asexual species, so mutation rates in sexual species will only evolve to exceed the lower bounds set by physiological costs and genetic drift in certain circumstances, and it is not clear when that might occur in practice (Romero-Mujalli & Tiedemann, 2021). Consequently, Section 4 of this paper explored in some detail how the evolutionary outcomes in changing environments are affected by having fixed mutation rates that are larger or smaller than the simulation-optimal values, and showed quite clearly how setting different mutation rates in the models has a significant effect on the simulation outcomes. In future models it will be crucial to set the mutation rates carefully, or at least explore the outcomes over the full range of plausible values.

A key issue with changing environments is that associated incidents of worsened performance can lead to higher death rates which in extreme cases may result in population extirpation or extinction. Section 6 of this study ran simulations to clarify the complex dependencies underlying extinction risk in changing environments, with particular reference to the much studied effect of change persistence. It is clear from the simulated extinction outcomes presented in Figures 9 and 10 that there are considerable differences depending on the details of the simulations which need to be tied down by empirical date from real populations. By separating the underlying factors, and simulating them over large ranges of values, it was found that in some situations the extinction risk increases with the change persistence, while in others it decreases. Similar differences have been observed across earlier studies with some (e.g., Johst & Wissel, 1997; Petchey et al.,1997) supporting an increased risk, some (e.g., Ripa & Lundberg, 1996) supporting a decreased risk, and others reporting mixed results (e.g., Cuddington & Yodzis, 1999). Such qualitative differences have also been observed before in the context of exploring different normalization approaches (Heino et al., 2000). By separating and independently varying all the crucial details (including the change magnitude, change persistence, mutation rates, and normalization approach), the simulations presented in this paper have clarified the sources of the conflicting results reported elsewhere.

Ultimately, artificial life models need to be validated by empirical data from real populations. Looking at extinction timescales is one way of doing that, but it is clear from the simulations of Section 6 that Gaussian distributions of changes lead to large variance in the survival times, so one has to be careful about interpreting empirical data from biological populations unless there are large numbers of them, which unfortunately is rarely the case. Comparing the generalist-specialist spectrum associated with environmental tolerance in real and modeled populations is another way. A complication here is that one normally separates environmental noise from signal on the grounds of predictability, but this may not always be appropriate for models of the evolution of environmental tolerance. In some cases, coping in the real environment will requires a tolerance of seasonal variation that may be larger than the required tolerance of environmental noise, so the predictable seasonal variation needs to be included in the environment signal $E$ even though it would normally be excluded. Further complications arise when the average magnitude of the environmental variation increases or decreases over time. To simplify the analysis, each simulation in the study here has been run with constant average magnitude, but it is clear from their outcomes what would happen if they were not. If the simulated species can evolve to track the change, for example by having large enough mutation rates, they will be fine. If they cannot, performance will fall and extinction may eventually be the consequence.

It is possible in situations where purely evolutionary mechanisms for coping with environmental variation are insufficient (for example, due to constraints on the evolution of mutation rates as discussed in Section 2.5) that non-evolutionary mechanisms such as lifetime learning could be adopted instead or in parallel. In some ways, this could be regarded as the evolution of a special form of tolerance, and many species are known to use learning to cope with change, but there are many associated costs, and evolving such abilities has its own difficulties. The evolution of various lifetime learning strategies in changing environments has already been simulated and discussed in some detail elsewhere (Bullinaria, 2018, 2022), and the dependence on the environmental change distribution was found to be small relative to that of the average change magnitude (Bullinaria, 2022). The crucial evolutionary factors discussed above, such as mutation rates, naturally have much less effect in these cases, and the learning can easily adjust to accommodate differing degrees of change persistence, but, rather than making the simulations simpler, these modeling complexities are replaced by others, such as the distinction between social and individual learning, and the fidelity of information transfer between individuals.

As noted earlier, the models studied in this paper were deliberately kept as simple a possible, and that leads to the obvious questions of whether they may have been simplified too far, and whether further variations in the details may affect the outcomes. In the process of refining and testing the models, many variations have already been explored, such as using different selection processes, crossover details, population sizes, update timescales, run lengths, and so on. Such changes do indeed generally lead to slightly different outcomes, as with the selection and minimum performance level variations discussed in Section 6, but none of them led to the broad pattern of outcomes being different from those presented above. Models containing or comparing multiple species that have different generation times (or reproduction rates) will clearly need refining to take better account of the timescales involved, as conveniently specifying everything in relation to an average magnitude of environmental change per generation will no longer work. Future models will need to tie all such details to empirical data for the particular species and environments of interest with view to obtaining more than broad patterns of results. In this regard, as discussed above, this study has highlighted the particular importance of setting or evolving appropriate mutation rates in the simulations.

Finally, it is important to note that there are limitations of the general simulation framework employed here and by Grove et al. (2020, 2021) that are less easy to explore. Representing the whole environment and the individual responses to it by single numbers on the same scale is an enormous simplification, as is the idea that mutations of the given species and their tolerance to environmental changes can also be represented by further single numbers on the same scale. There is also the issue of how the distribution and average magnitude of the environmental changes might vary over time with different timescales, and how the distinction between sexual and asexual reproduction might need to be modeled more realistically. But simulations of this kind need to start somewhere. The principal factors have been explored and there are no remaining unnecessary details left to confuse matters in the models presented above. The key issues appear to have been captured despite the simplifications, but it will certainly be interesting to see if future simulations that represent more of the details and complications of real populations and their environments lead to deviations from the pattern of outcomes observed in this study.

## References

André, J. B. & Godelle, B. (2006). The evolution of mutation rate in finite asexual populations. *Genetics*, *172*, 611-626.

Bromham, L. (2009). Why do species vary in their rate of molecular evolution? *Biology Letters, 5*, 401-404.

Bullinaria, J. A. (2018). Evolution of learning strategies in changing environments. *Cognitive Systems Research*, *52*, 429-449.

Bullinaria, J. A. (2022). Effect of environmental change distribution on artificial life simulations. *Artificial Life*, *28*, 134-153.

Bürger, R. & Lynch, M. (1995). Evolution and extinction in a changing environment: A quantitative-genetic analysis. *Evolution, 49*, 151-163.

Cuddington, K. M. & Yodzis, P. (1999). Black noise and population persistence. *Proceedings of the Royal Society of London B*, *266*, 969-973.

Eiben, A. E. & Smith, J. E. (2016). *Introduction to Evolutionary Computing*. Berlin, Germany: Springer.

Eyre-Walker, A. & Keightley, P. D. (2007). The distribution of fitness effects of new mutations. *Nature Reviews Genetics*, *8*, 610-618.

Gilljam, D., Knape, J., Lindén, A., Mugabo, M., Sait, S. M. & Fowler, M. S. (2019). The colour of environmental fluctuations associated with terrestrial animal population dynamics. *Global Ecology and Biogeography*, *28*, 118-130.

Good, B. H. & Desai, M. M. (2016). Evolution of mutation rates in rapidly adapting asexual populations. *Genetics*, *204*, 1249-1266.

Grove, M. (2014). Evolution and dispersal under climatic instability: a simple evolutionary algorithm. *Adaptive Behavior*, *22*, 235-254.

Grove, M., Borg, J. M. & Polack, F. (2020). Coloured noise time series as appropriate models for environmental variation in artificial evolutionary systems. In *Proceedings of the 2020 Conference on Artificial Life,* (pp. 292-299). MIT Press.

Grove, M., Timbrell, L., Jolley, B., Polack, F. & Borg, J. M. (2021). The importance of noise colour in simulations of evolutionary systems. *Artificial Life, 27*, 164-182.

Haaland, T. R., Wright, J. & Ratikainen, I. I. (2020). Generalists versus specialists in fluctuating environments: A bet-hedging perspective. *Oikos*, *129*, 879-890.

Halley, J. M. (1996). Ecology, evolution and 1/f noise. *Trends in Ecology and Evolution*, *11*, 33-37.

Halley, J. M. (2003). The initial-value spectrum of 1/f noise: Notes on its importance and its estimation. In *Proceedings of the 17th International Conference on Noise and Fluctuations*, Josef Sikula (Ed.), (pp. 497-500). Brno, Czech Republic: CNRL.

Halley, J. M. & Inchausti, P. (2004). The increasing importance of 1/f-noises as models of ecological variability. *Fluctuation and Noise Letters*, *4*, R1-R26.

Halley, J. M. & Kunin, W. E. (1999). Extinction risk and the 1/f family of noise models. *Theoretical Population Biology*, *56*, 215-230.

Heino, M., Ripa, J. & Kaitala, V. (2000). Extinction risk under coloured environmental noise. *Ecography*, *23*, 177-184.

Inchausti, P. & Halley, J. (2003). On the relation between temporal variability and persistence time in animal populations. *Journal of Animal Ecology*, *72*, 899-908.

Johnson, T. (1999). Beneficial mutations, hitchhiking and the evolution of mutation rates in sexual populations. *Genetics*, *151*, 1621-1631.

Johst. K. & Wissel. C. (1997). Extinction risk in a temporally correlated fluctuating environment. *Theoretical Population Biology*, *52*, 91-100.

Kamenev, A., Meerson, B. & Shklovskii, B. (2008). How colored environmental noise affects population extinction. *Physical Review Letters*, *101*, 268103.

Lynch, M. (2011). The lower bound to the evolution of mutation rates. *Genome Biology and Evolution*, *3*, 1107-1118.

Lynch, M. & Gabriel, W. (1987). Environmental tolerance. *The American Naturalist*, *129*, 283-303.

Martincorena, I. & Luscombe, N. M. (2013). Non-random mutation: The evolution of targeted hypermutation and hypomutation. *BioEssays*, *35*, 123-130.

Metzgar, D. (2007). Mutation Rates: Evolution. *Encyclopedia of Life Sciences*. Wiley.

Morales, J. M. (1999). Viability in a pink environment: Why "white noise" models can be dangerous. *Ecology Letters*, *2*, 228-232.

Newman, M. E. J. (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, *46*, 323-351.

Ogbunugafor, C. B. (2023). Mutations that enhance evolvability may open doors to faster adaptation. *Nature Communications*, *14*, 6310.

Petchey, O. L., Gonzalez, A. & Wilson, H. B. (1997). Effects on population persistence: The interaction between environmental noise colour, intraspecific competition and space. *Proceedings of the Royal Society of London B*, *264*, 1841-1847.

Potts, R. (2013). Hominin evolution in settings of strong environmental variability. *Quaternary Science Reviews*, *73*, 1-13.

Ripa, J. & Lundberg, P. (1996). Noise colour and the risk of population extinctions. *Proceedings of the Royal Society of London B*, *263*, 1751-1753.

Roff, D. A. (2002). *Life History Evolution*. Sunderland, MA: Sinauer Associates.

Romero-Mujalli, D. & Tiedemann, R. (2021). Evolution of the mutation rate in sexual species. *eLS*, *2*, 140-146.

Ruokolainen, L., Linden, A., Kaitala, V. & Fowler, M. S. (2009). Ecological and evolutionary dynamics under coloured environmental variation. *Trends in Ecology and Evolution*, *24*, 555-563.

Sniegowski, P. D., Gerrish, P. J., Johnson, T. & Shaver, A. (2000). The evolution of mutation rates: separating causes from consequences. *BioEssays*, *22*, 1057-1066.

Stearns, S. C. (1992). *The Evolution of Life Histories*. Oxford, UK: OUP.

Steele, J. H. (1985). A comparison of terrestrial and marine ecological systems. *Nature*, *313*, 355-358.

Vasseur, D. A. & Yodzis, P. (2004). The color of environmental noise. *Ecology*, *85*, 1146-1152.

Wei, W., Ho, W. C., Behringer, M. G., Miller, S. F., Bcharah, G. & Lynch, M. (2022). Rapid evolution of mutation rate and spectrum in response to environmental and population-genetic challenges. *Nature Communications*, *13*, 1-10.

Wichmann, M. C., Johst, K., Schwager, M., Blasius, B. & Jeltsch, F. (2005). Extinction risk, coloured noise and the scaling of variance. *Theoretical Population Biology*, *68*, 29-40.

Wilke, C. O., Wang, J. L., Ofria, C., Lenski, R. E. & Adami, C. (2001). Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature*, *412*, 331-333.

Yedid, G., Ofria, C. A. & Lenski, R. E. (2009). Selective press extinctions, but not random pulse extinctions, cause delayed ecological recovery in communities of digital organisms. *The American Naturalist*, *173*, E139-E154.

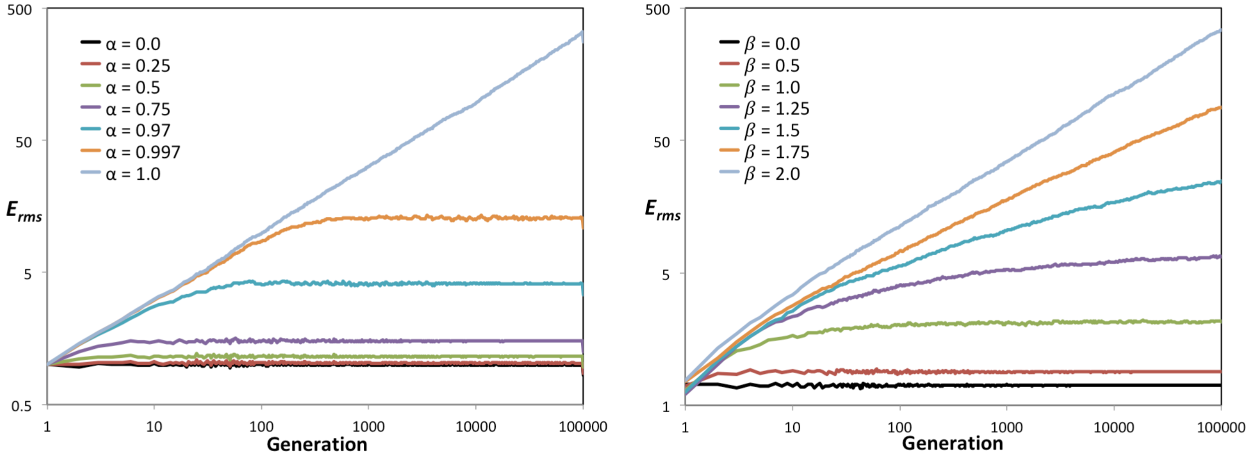

Figure 1. The average buildup of root-mean-squared (rms) environmental change $E_{rms}$ over generations for autoregressive environmental time series with representative values of change persistence $\alpha$ and fixed imposed change level $m = 1$ (left), and inverse power law time series with representative values of change exponent $\beta$ and fixed rms change per generation $m = 1$ (right).

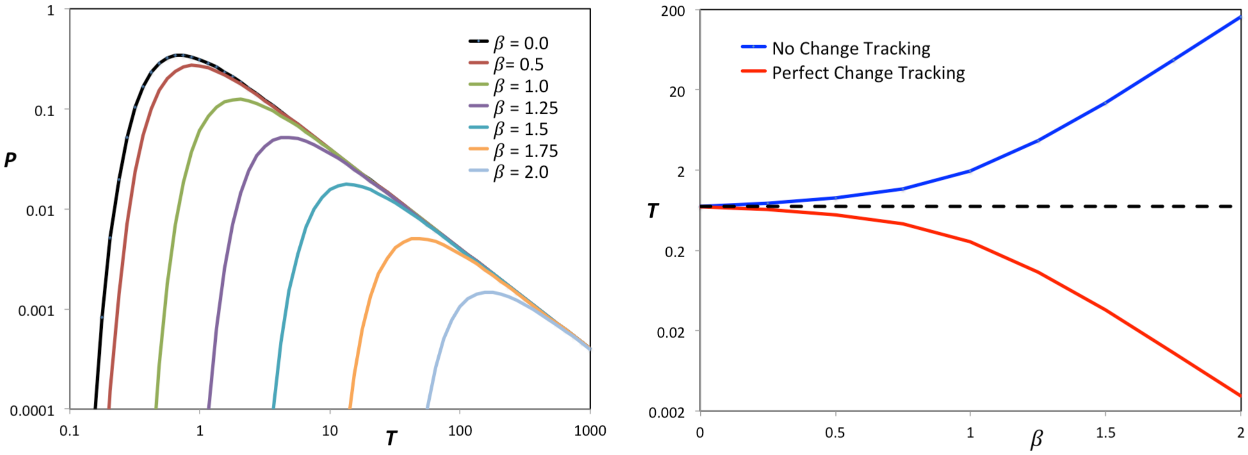

Figure 2. The performance $P$ as a function of tolerance $T$ when individual responses do not track the changes at all, for stabilized rms environmental change levels of the IPL time series normalized to have fixed change per generation (left). This leads to the increasing evolved $T$ with $\beta$ shown (right), along with the corresponding decrease of evolved $T$ with $\beta$ when the population responses track the changes perfectly and the environmental time series is normalized to have fixed rms cumulative change levels. The dashed line represents the lack of variation resulting from switching the normalization type in each case.

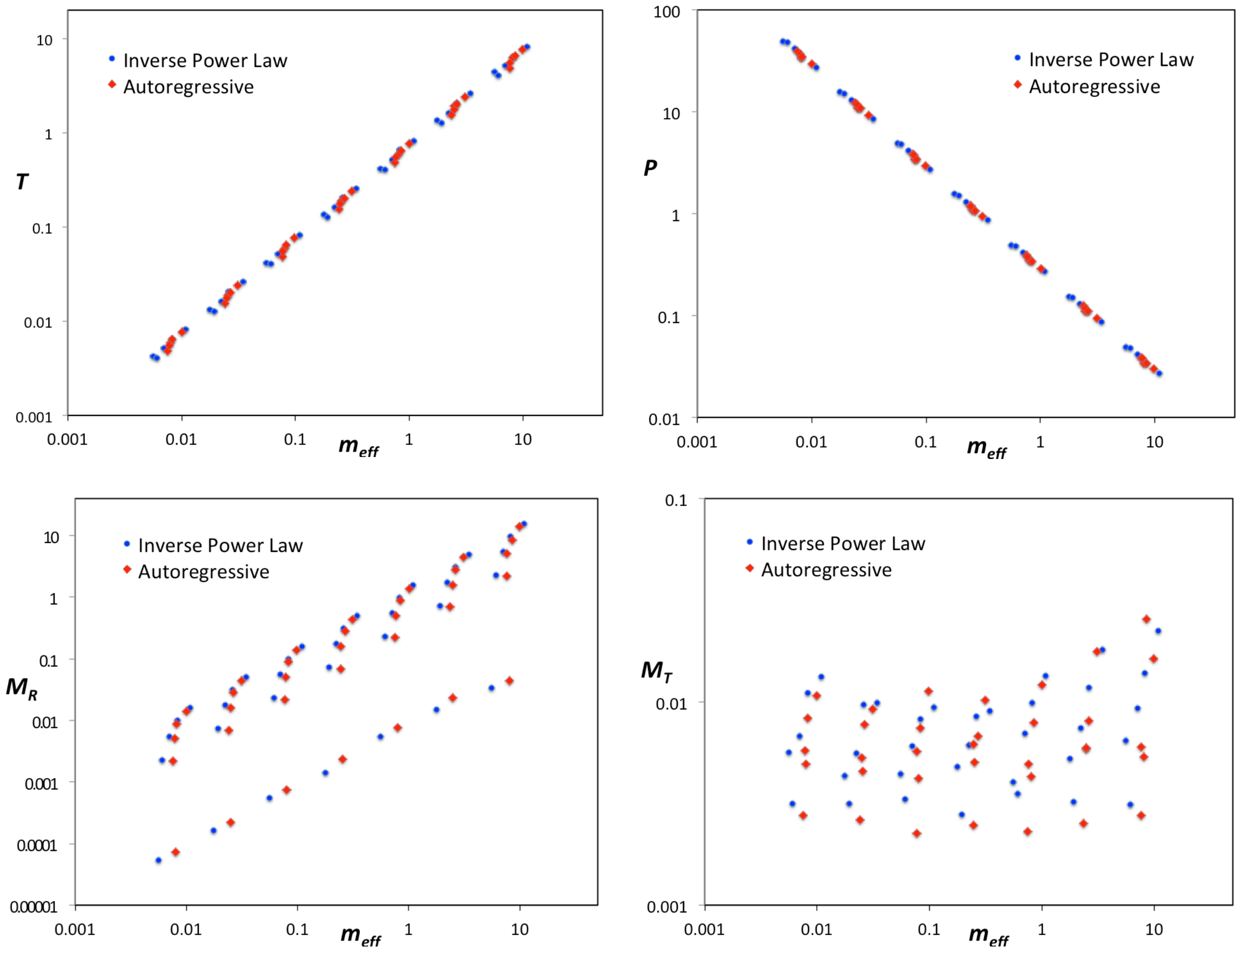

Figure 3. The average evolved environmental change tolerance $T$ (top-left), associated performance level $P$ (top-right), response mutation rate $M_R$ (bottom-left) and tolerance mutation rate $M_T$ (bottom-right), as a function of the average effective change magnitude $m_{eff}$ when the mutation rates are allowed to co-evolve. Each cluster corresponds to a particular imposed change magnitude $m \in \{0.01, 0.032, 0.1, 0.32, 1.0, 3.2, 10.0\}$, and the variation within each cluster reflects the effect of different autoregressive change persistence $\alpha \in \{0, 0.25, 0.5, 0.75, 1\}$ or inverse power law exponent $\beta \in \{0, 0.5, 1.0, 1.5, 2\}$ as shown more clearly in Figure 4.

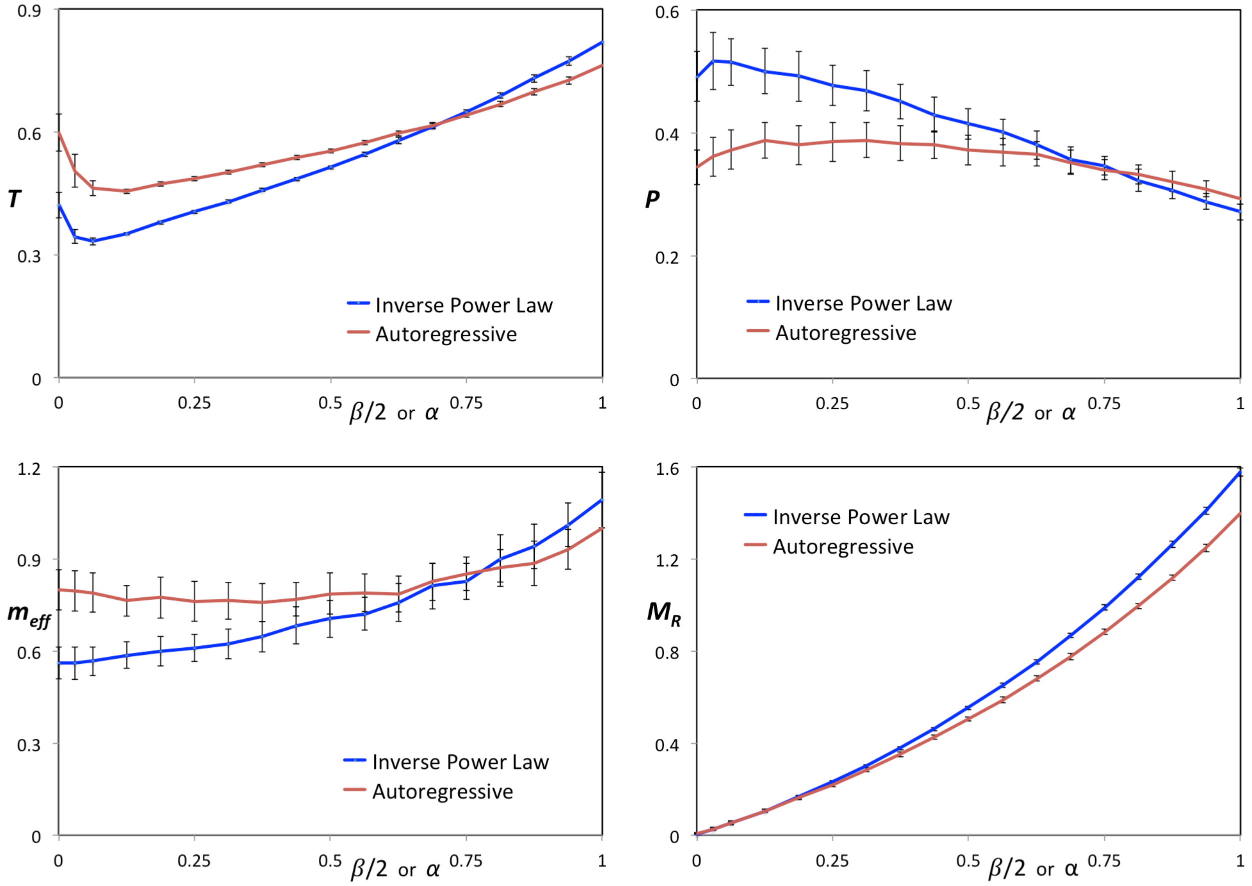

Figure 4. The variation with change persistence $\alpha$ or $\beta$ of the average evolved tolerance $T$ (top-left), the average individual performance level $P$ (top-right), the average effective change magnitude $m_{eff}$ (bottom-left), and the average evolved mutation rate $M_R$ (bottom-right), for the case of evolutionary optimized mutation rates $M_R$ and $M_T$ for a representative imposed change magnitude ($m = 1$).

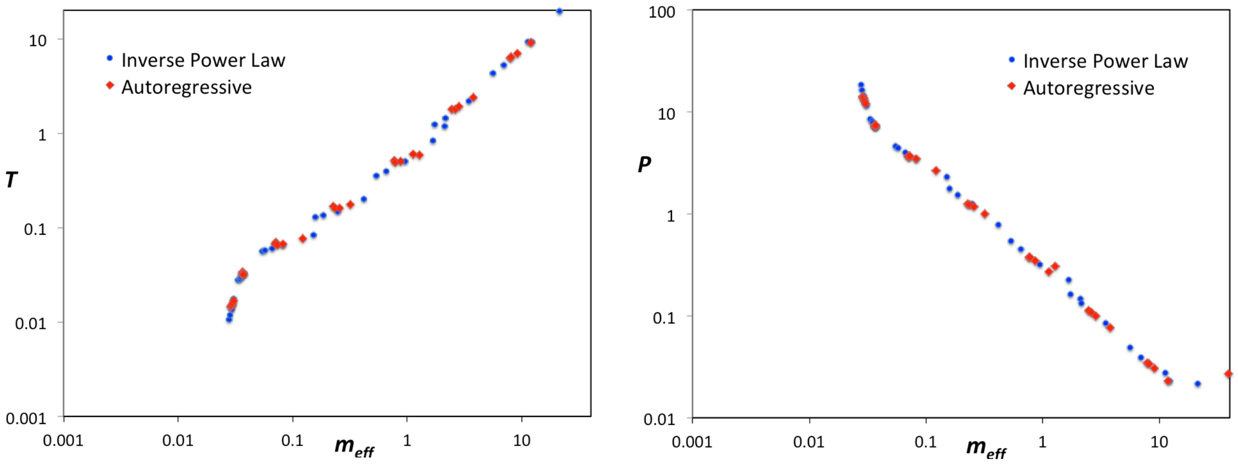

Figure 5. The average evolved environmental change tolerance $T$ (left) and associated average performance level $P$ (right), as a function of the average effective change magnitude $m_{eff}$ when the mutation rates are fixed at $M_R = 0.1$ and $M_T = 0.01$. Each cluster (now less clear than in Figure 3) corresponds to a particular imposed change magnitude $m \in \{0.01, 0.032, 0.1, 0.32, 1.0, 3.2, 10.0\}$, and the variation within each cluster reflects the effect of different autoregressive change persistence $\alpha \in \{0, 0.25, 0.5, 0.75, 1\}$ or inverse power law exponent $\beta \in \{0, 0.5, 1.0, 1.5, 2\}$ as shown more clearly in Figure 6.

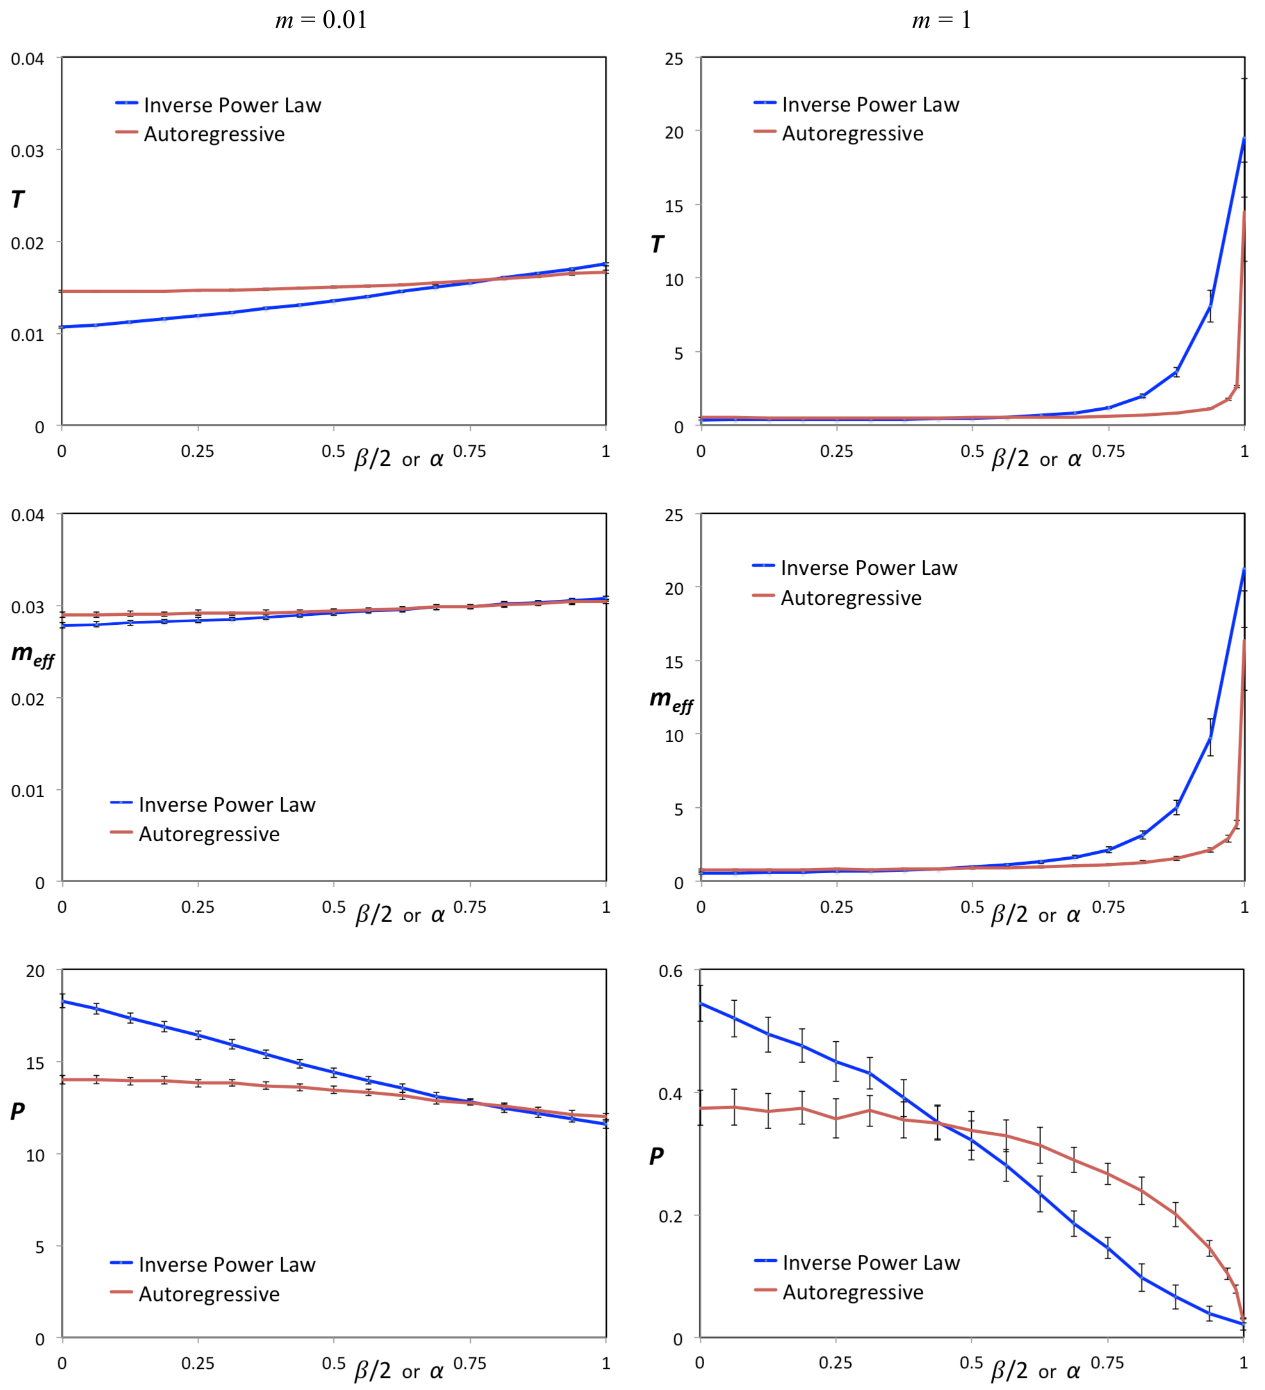

Figure 6. The variation with change persistence $\alpha$ or $\beta$ of the average evolved tolerance $T$ (top), the average effective change magnitude $m_{eff}$ (middle) and the average individual performance level $P$ (bottom), with fixed mutation rates $M_R$ and $M_T$ for two representative fixed imposed change magnitudes $m$. For $m = 0.01$ (left) the mutations are larger than optimal for all $\alpha$ or $\beta$, for $m = 1$ (right) they are less than optimal for all $\alpha$ and $\beta$ /2 above about 0.12.

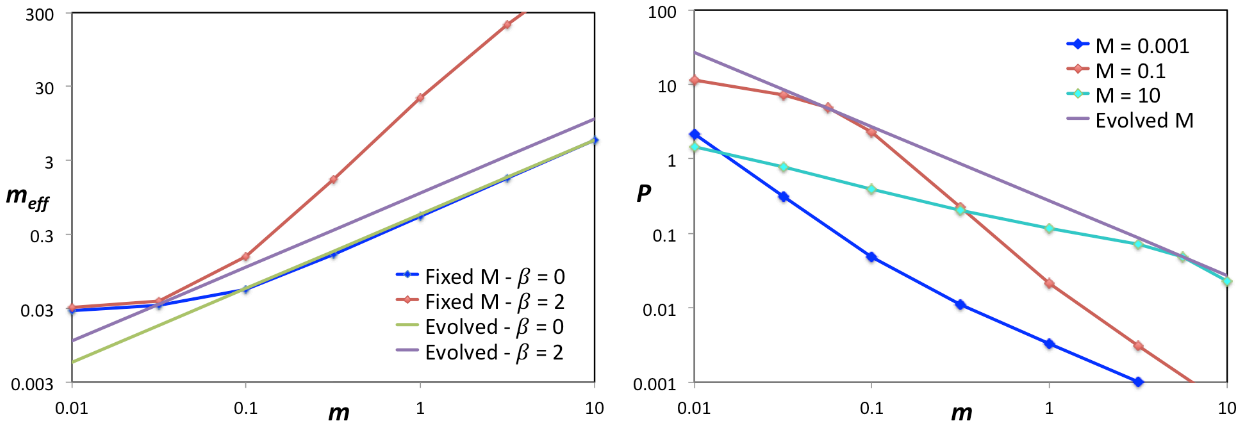

Figure 7. The average effective environmental change magnitude $m_{eff}$ when the mutation rates $M_R$ are either evolved or fixed at 0.1 with fixed change IPL persistence $\beta = 0$ or 2 (left), and the average performance $P$ achieved when the mutation rates $M_R$ are either evolved or fixed at 0.001, 0.1 or 10 with fixed IPL persistence $\beta = 2$ (right). Both sets of results are plotted against the imposed environmental change magnitude $m$.

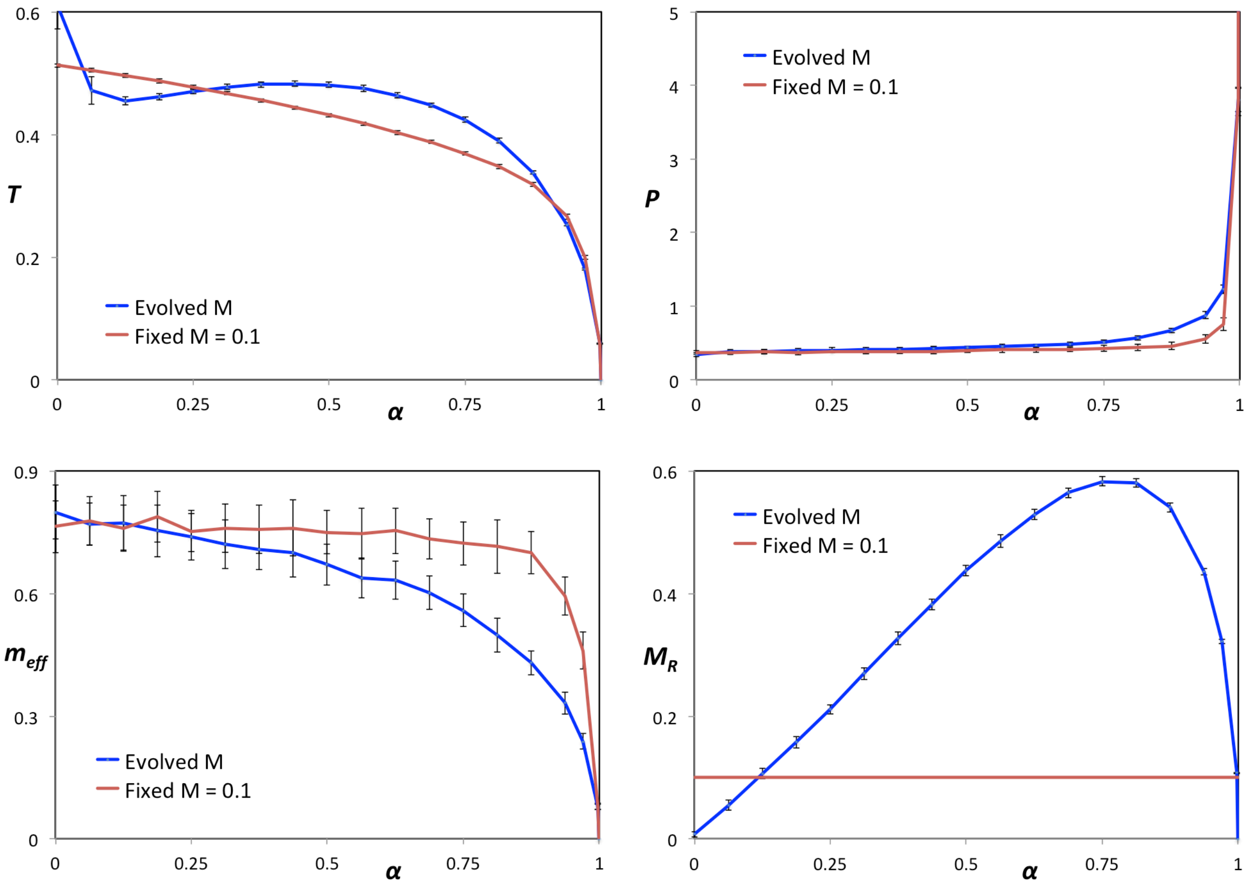

Figure 8. The variation with autoregressive change persistence $\alpha$ of the average tolerance $T$ (top-left), the average individual performance level $P$ (top-right), the average effective change magnitude $m_{eff}$ (bottom-left), and the average mutation rate $M_R$ (bottom-right), for the case of normalized cumulative change magnitudes with either evolved or fixed mutation rates.

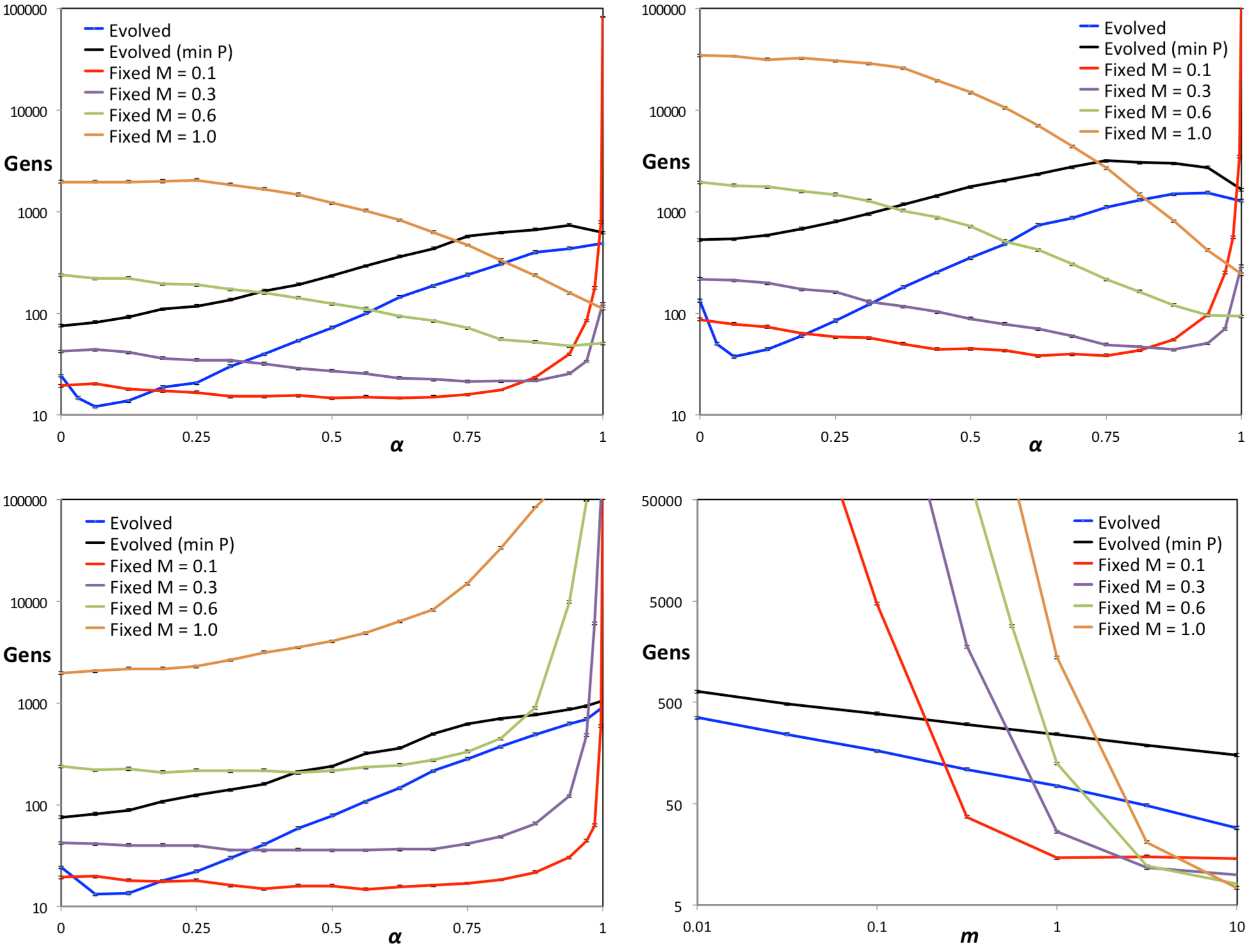

Figure 9. The number of generations before extinction as a function of autoregressive change persistence $\alpha$, for fixed imposed change magnitude $m$ = 1 with $P_{min}$ = 0.001 (top-left) and $P_{min}$ = 0.00001 (top-right), and normalized cumulative magnitude with $P_{min}$ = 0.001 (bottom-left). Also shown are the generations before extinction as a function of imposed change magnitude $m$ for intermediate fixed $\alpha$ = 0.5 with $P_{min}$ = 0.001 (bottom-right). In each case, the results are shown for the standard and more selective (min P) evolved mutation rates, and fixed mutation rates with $M_R \in \{0.1, 0.3, 0.6, 1.0\}$. The variances are large, so means and standard errors from 1000 runs are shown.

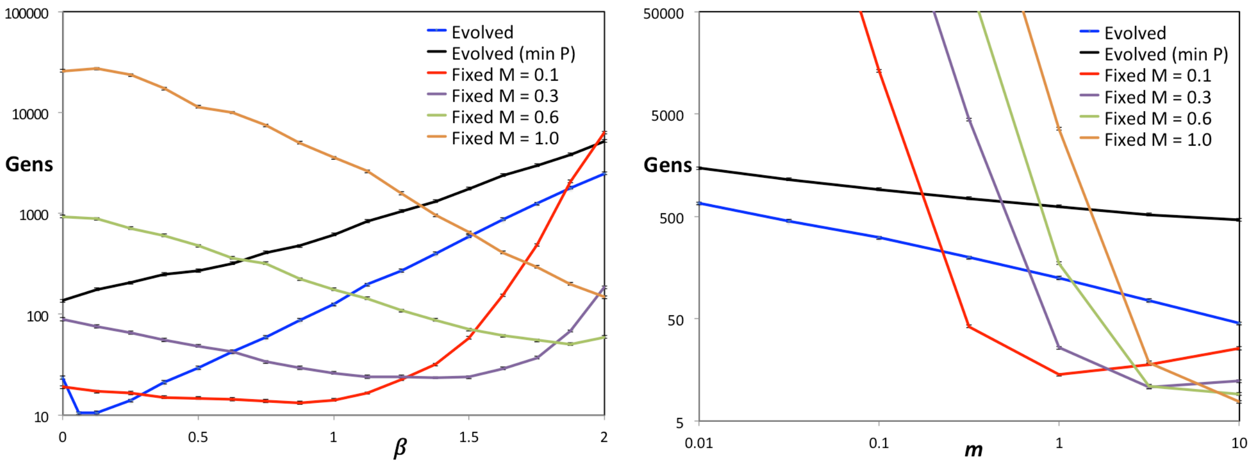

Figure 10. The number of generations before extinction as a function of inverse power law change persistence $\beta$ for fixed imposed change magnitude $m = 1$ with $P_{min} = 0.001$ (left), and the generations before extinction as a function of imposed change magnitude $m$ for intermediate fixed $\beta = 1.0$ with $P_{min} = 0.001$ (right). In each case, the results are shown for the standard and more selective (min P) evolved mutation rates, and fixed mutation rates with $M_R \in \{0.1, 0.3, 0.6, 1.0\}$. The variances are large, so means and standard errors from 1000 runs are shown.