Evolution of evolvability

Master thesis of Jana M. Riederer

s3507130

February 2019 – December 2019

(50 ECTS)

Supervised by prof. dr. F.J. Weissing & T.J.B van Eldijk

MSc Ecology & Evolution – Track Evolutionary Biology

Theoretical Research in Evolutionary Life Sciences

Rijksuniversiteit Groningen

rijksuniversiteit groningen

MARM GROUP

Table of Contents

Evolution of evolvability

Preface

When I started this MSc project, I set out to investigate the impact of population turn-over on evolvability, by modelling the evolution of death rates. Now, at the end thereof, I find myself with three distinct projects, all approaching the topic of evolvability from different angles: the evolution of mutation rates, the evolution of death rates, and the evolution of a plastic trait under episodic selection. I would therefore like to take this opportunity to explain how these various projects came about, and how they fit together.

I had no prior experience with computer models, hence as a stepping stone I started with a model examining evolvability in its simplest form: the spread of strains with different mutation rates. This was intended as a mere exercise which would allow me to explore the evolution of evolvability in fluctuating environments and would eventually inform me on how to build my model on death rate evolution. However, this simplified mutation rate model surprisingly yielded interesting results, prompting me to take mutation rate evolution as a second route (in addition to death rate evolution) –exploring the evolution of evolvability from a different angle. Both of these approaches highlight the role that population dynamics play in the evolution of evolvability. Any research into the evolution of evolvability through mutation rate is written up in Part I of the project.

Subsequently I modelled my original idea concerning the impact of population turn-over on evolvability, exploring whether selection for evolvability can elicit an evolutionary response in the individual mortality through kin selection. Specifically, I tested my hypothesis that selection for evolvability increases death rates. Research into evolution of evolvability through population turn over is written up in Part II of the project.

A few months into my MSc project, I was offered the exciting opportunity to work on a project on Dollo's law, in collaboration with the research group of Dr. Bertanne Visser (Université catholique de Louvain, Belgium). As this fit well with my general overarching topic of evolvability, I developed a theoretical model together with Stephano Tiso and Timo van Eldijk (University of Groningen) to complement the empirical research carried out by Dr. Visser. This was a wonderful learning opportunity which allowed me to greatly improve my modelling skills as well as expand my theoretical knowledge. The paper resulting from this collaboration has since been submitted for publication to PNAS. Generally, this project explored how the evolvability of traits under episodic selection is mediated by plasticity, and how this explains apparent violations of Dollo's law. The resulting manuscript is here submitted as Part III of the project.

Overall, all three parts of my MSc project consider the question of how evolvability itself may evolve: evolvability achieved through mutation rate (part I), evolvability achieved through population turnover (part II), and the impact of episodic selection on evolvability (part III). In these projects I investigated how the evolution of evolvability is influenced by population dynamics, plasticity, and the assumption of the underlying genetic system, thus uniting the three different approaches in one overarching theme.

Evolution of evolvability

Part I: Evolution of the mutation rate

Master thesis of Jana M. Riederer

s3507130

February 2019 – December 2019

(50 ECTS)

Supervised by prof. dr. F.J. Weissing & T.J.B van Eldijk

MSc Ecology & Evolution – Track Evolutionary Biology

Theoretical Research in Evolutionary Life Sciences

Rijksuniversiteit Groningen

rijksuniversiteit groningen

MARM GROUP

Evolution of evolvability I: Evolution of the mutation rate

Introduction

Evolvability has recently come to the forefront of evolutionary biology. This is partly due to the growing popularity of the "Extended Evolutionary Synthesis (EES)" initiative, which has brought some previously neglected topics, such as evolvability, into the spotlight (Laland et al. 2015, Nuño de la Rosa 2017). Further to this, evolvability is also starting to play a role in the conservation of threatened ecosystems (Wright et al. 2019) and climate change research (Stojanova et al. 2019).

The concept of evolvability itself, however, is by no means new. Various papers from the 20th century explored this concept (see for example Hanson 1966, Schwabauer 1976, Swartout & Neches 1986, Kauffman 1990, Houle 1992 and Rowe et al. 1994), yet the use of the term "evolvability" differed widely in those early papers. Wagner and Altenberg (1996) subsequently clarified the terminology, defining evolvability as "a genome's ability to produce adaptive variants". This interpretation thus includes not only variation but crucially also variability, and can be rephrased as a system's ability to change in response to selection pressures (Pigliucci 2008). Evolvability can be influenced by various traits and processes, including, but not limited to modularity (Parsons et al 2018, Verd et al. 2019), the type of genotype-phenotype mapping (Wagner & Zhang 2011), or plasticity (Van Gestel & Weissing 2016). But perhaps the most intuitive change in evolvability is a change in mutation rate: a higher mutation rate increases variation and variability of traits, thus providing the raw material for evolution in greater amounts. In the first part of the report, I therefore want to focus on evolvability through mutation rate.

One of the most important questions in relation to evolvability is whether evolvability itself can evolve – that is, change not only as a byproduct of selection on other traits, but in response to selection pressure for increased or decreased evolvability. A striking example of mutation rate evolution comes from experimental evolution: After 12 strains of an ancestral population of *Escherichia coli* evolved independently for 50,000 generations, 6 of the strains evolved hyper-mutability, with mutation rates orders of magnitude larger than those observed in the ancestral phenotype (Tenaillon 2016). This leads to the question of what selective forces shape mutation rates – are high mutation rates mal-adaptive side effects? Or are these high mutation rates adaptive, perhaps even an example of evolution of evolvability?

One of the most prominent explanations for the maintenance of high but mal-adaptive mutation rates is based on the concept of "cost of fidelity". This hypothesis postulates that there is a trade-off between replication speed and replication fidelity: Proof-reading mechanisms can be time consuming and therefore slow down cell division, imposing the cost of slower reproduction – a cost that can be critical especially in prokaryotes (Furió et al. 2005, Baer et al. 2007, Belshaw et al. 2008). In other words, selection for faster replication speed leads to the unfortunate but unavoidable side-effect of higher mutation rates. However, whilst this is a common and widely accepted explanation for high mutation rates in viruses, it is doubtful that it can satisfactorily explain high mutation rates in eukaryotic, and especially in multicellular organisms, in which speed of reproduction is dependent on many factors other than speed of cell division (Regoes et al. 2013). In addition to "cost of fidelity", various other hypotheses have been put forward aiming to demonstrate how high mutation rates can arise without being adaptive. These include diverse mechanisms, including but not limited to variation in exposure to mutagens (Baer et al. 2007), specific aspects of DNA structure such as repetitive elements (Jeffreys et al. 1988), or relaxed selection on mutation rates due to buffers such as heat shock

proteins compensating for deleterious mutations (Rutherford 2003, Baer et al. 2007). Overall, all of these mechanisms, including the "cost of fidelity" hypothesis explain high mutation rates as a maladaptive by-product of other processes.

However, there is also evidence that high mutation rates can be adaptive. Strains or populations with higher mutation rates outperform those with lower mutation rates in a variety of contexts, as demonstrated in both experimental evolution of mutator strains engineered in the lab, as well as in experiments using natural variation in mutation rates (see for example Desai et al. 2007, Maharjan et al. 2013, or Sprouffske et al. 2018). Furthermore, environmental stress can prompt a process termed stress-induced mutagenesis, a transient up-regulation of the mutation rate brought about by environmental stressors (Matic 2013). It is important to emphasize that these increased mutation rates are not due to a reduction of function caused by stress, but rather through evolved mechanisms linking stress and increased mutation rate, such as activation of error-prone DNA polymerases and down-regulation of error-correcting enzymes (Foster 2007, MacLean et al. 2013). This mechanism is part of the organism's defense mechanism, and facilitates faster adaptation to a changed environment. Overall, these lines of evidence indicate that higher mutation rates can indeed be adaptive, and might evolve in response to a selection pressure for evolvability.

Yet this adaptiveness of high mutation rates can be of a transient nature: once a solution to a novel environmental challenge has been found, the high mutation rate is no longer required. Thus, mutation rates can be expected to rise only transiently by hitchhiking on the beneficial variants they produce, and to then be lowered again through selection for a reduction of deleterious mutations. However, this can be prevented by environmental fluctuations: a constantly changing environment necessitates constantly new adaptations, thus facilitating repeated hitchhiking events, and preventing transient mutation rate increases from being lost. And indeed, previous papers have succeeded in showing that fluctuating environments can allow the establishment of a stable, non-zero mutation rate, albeit under simplifying assumptions such as that of an infinitely large populations (Ishii et al. 1989, Carja et al. 2014). Furthermore, variable environments are crucial to the concept of evolvability itself, as higher evolvability, i.e. a greater ability for adaptive change, is intrinsically liked to environments selecting for frequent changes, thus further cementing the role that environmental variation plays in the context of mutation rates.

Considering that mutation rates may be an important target for evolvability selection, it is interesting to investigate the mechanism by which it can evolve. Mutation rates do not have a phenotype under direct selection: their phenotype only is expressed though their impact on the probability of a different phenotype changing. Selection on mutation rates can therefore only act indirectly, through selection on other loci. Mutation rates thus evolve through genetic "hitchhiking": the phenomenon that an allele changes in frequency due to being linked to a different allele under direct selection (Johnson 1999, Gerrish et al. 2007, Cobben et al. 2017). In the case of mutation rates, genetic hitchhiking translates to a mutator allele inducing beneficial mutations and then hitchhiking along in the beneficial mutation's selective sweep. This concept first arose under the name periodic selection (Atwood et al. 1951), and has since gained support from both theoretical and empirical studies (see for example Johnson 1999 and Shaver et al. 2002). Due to its dependence on genetic hitchhiking, the evolution of mutation rates is very sensitive to anything that might break the crucial linkage between mutator allele and beneficial mutation, such as high recombination rates or sexual reproduction. In fact, many studies find that the positive selection on mutation rates via hitchhiking should be absent or negligible in sexually reproducing populations (Sniegowski et al. 2000, Baer 2007, Martincorena & Luscombe 2013, Mujialli 2019). Other papers do succeed in showing the evolution of high mutation rates via hitchhiking under sexual reproduction in theoretical models, however only under additional assumptions and special

circumstances, such as presuming an explicit link between mutation rate and mate choice as well as strong sexual selection (Petrie & Roberts 2007), or considering only population expansions and allowing mutation rate to coevolve with dispersal speed (Cobben et al. 2017).

Overall, mutation rate evolution is an attractive approach to the concept of evolvability, which is supported by empirical data indicating that high mutation rates can be adaptive. Mutation rates can increase via hitchhiking with beneficial mutations, but these increases are only transient without repeatedly changing selection pressures in varying environments. Hitchhiking relies on linkage between mutator and beneficial mutation, casting doubt on the ability of mutation rates to evolve under recombination and sexual reproduction. Keeping this body of literature in mind, I here aim to see whether and how higher mutation rates evolve in a fluctuating environment, under explicit modelling of population dynamics (i.e. neither constant nor infinite population sizes). This topic is explored in three models, investigating a mutator invasion analysis (model 1), the evolution of mutation rate under asexual reproduction (model 2) and the evolution of mutation rate under sexual reproduction (model 3).

Model 1: Invasion of a mutator strain in a periodically changing environment

Model description

Model structure

The first model considers how the invasion of a strain in a resident population proceeds if the strain and resident population have different mutation rates, investigating whether the invading strain or resident population eventually outcompete the other. I analyzed this question using a continuous-time stochastic birth-death model written in R, tracking the number of individuals of each genotype through time. Each individual has two gene loci: one determining its ecological adaptation and one determining its mutation rate. The ecological adaptation is encoded at a single di-allelic locus, which can take the states A or B, thus determining whether the individual is adapted to resource A or B. The mutation locus determines whether the individual has a high or a low mutation rate. This results in four genotypes: a mutator adapted to resource A, a mutator adapted to resource B, a non-mutator adapted to resource A, and a non-mutator adapted to resource B. Either the resident or the invading strain act as mutator. During a mutation event, an individual's ecological adaptation allele changes from A to B or vice versa, with no lethal mutations. The individuals reproduce asexually with no recombination: during each birth event, the new individual inherits its parent's mutation rate, and its parent's ecological adaptation allele if there is no mutation event (or the opposite adaptation allele if there is a mutation event). Population growth is limited by carrying capacity: the birth rate is proportional to the fraction of the carrying capacity that is unoccupied, and decreases therefore with increasing population size. The carrying capacity for resource A and B are modelled separately. Both carrying capacities oscillate over time, and the oscillations are shifted such that the highest point of carrying capacity A coincides with the lowest point of carrying capacity B, and vice versa (see Figure 1). Thus the alleles A and B are selected in an alternating manner, as might be the case under, for example, seasonal selection pressures alternating between winter and summer, or wet and dry season.

I repeated the analysis with a second version of the model, in which the oscillating selection pressure is caused not by changes in the environment but by negative frequency dependent selection. This allowed me to investigate not only the effect of external, environmentally driven oscillations, but also the effect of oscillating selection pressures arising from evolutionary dynamics within the system. In

this model, an individual's birth rate depends inversely on the frequency of its ecological adaptation allele, rather than depending on ecological adaptation.

To analyse the results, I focused on the relative frequency of the high mutation genotype at the end of each simulation. Higher frequencies were interpreted as a faster and more complete invasion by the genotype.

Parameter values

I used the following parameter values: all genotypes have an initial abundance of 50 individuals (resulting in an initial total population size of 200), a death rate d = 0.1, and a the unmodified (maximal) birth rate b = 0.5. The carrying capacity for carriers of allele A is given by

$$\mathrm{K} = 250 * (\sin(\mathrm{s} * \mathrm{time})) + 300.$$

This results in oscillations in carrying capacity between 50 and 550 individuals, with the period length determined by the scaling factor s. I used two different values for s to investigate the effect of environmental oscillation speed: s = 0.05 (resulting in a period length of 125.66) and s = 0.1 (resulting in a period length of 62.83). The carrying capacity for carriers of allele B is given by the same formula but shifted by half the period length, allowing the two carrying capacities to fluctuate out of phase. The behaviour of the model was explored using different mutation rates between 0 and 0.1 for the resident mutation rate (M1) and the invading mutation rate (M2), with 200 replicates of 1500 time units for each parameter combination. To control for the influence of environmental oscillations, I implemented a simulation with a stable environment, in which the carrying capacities of both resources are stably at 300 individuals each. In the control simulation, I used M1 = 0 and M2 = 0.1, all other parameter values were left unchanged, including replicate number and simulation length.

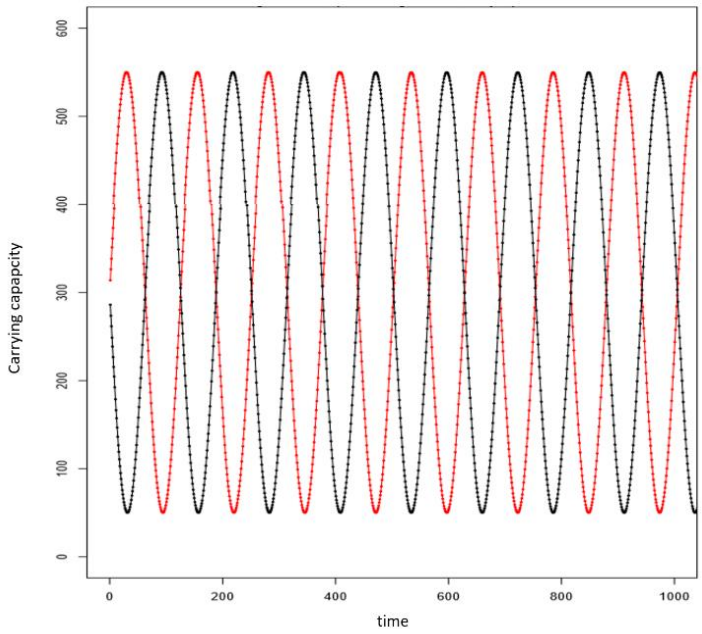

Figure 1: Changes in resource specific carrying capacities over time. This figure shows the oscillating carrying capacities of resources A (red) and B (black). The carrying capacities oscillate between 50 and 500 individuals, the period length is modified by s = 0.05 (i.e. period length = 125.66 time units). The carrying capacities for both resources oscillate out of phase so that the highest points of availability of resource A coincide with the lowest points of availability of resource B. This creates a variable, oscillating selection regime.

Simulation results

Oscillating resources

The four genotypes (mutator adapted to resource A, mutator adapted to resource B, non-mutator adapted to resource A, non-mutator adapted to resource B) start at equal frequencies and then evolve for 1500 time units. Figure 2 shows an example of how these four genotype frequencies change through time: in this case, the two non-mutator strains are lost early on, whereas the two mutator strains persist. Of the remaining strains, one is adapted to resource A and one is adapted to resource B, and their frequencies therefore show oscillations driven by the fluctuating carrying capacities of resources A and B.

I analyzed the effect of mutation rate on spread of the genotype. To do so, I considered the number of individuals of each genotype at the end of the simulation runtime, comparing the frequencies of the mutator and non-mutator strains. Overall, mutator strains are able to successfully outcompete non-mutator strains, with the mutator genotypes being strongly over-represented at the ends of most replicates (for an example, see Figure 3, for plots showing all invader/resident comparisons see Appendix Figure 1). The proportion of the invading mutation rate is significantly different form 0.5, this pattern holds regardless of the values chosen for M1 and M2, as long as one mutation rate is higher. If both mutation rates are of equal value, the any differences in the final numbers of individuals with either mutation rate are caused by drift, and are not statistically significant. For some parameter combinations, the invasion was complete by the end of the simulation, with the lower mutation rate gone from the population in all 200 replicates. This strong and consistent spread of high mutation rates cannot be observed in the control simulation, in which the population evolves in a stable environment (Figure 4). There seems to be a slight bias towards high mutation rates even amongst the mutation rate frequencies in the control simulation. Whilst this could potentially be explained by over-sampling, it necessitates further testing. Regardless of this, the difference between populations evolving in a fluctuating environment and those evolving in a stable control environment is striking, confirming that environmental oscillations are crucial in selecting for higher mutation rates.

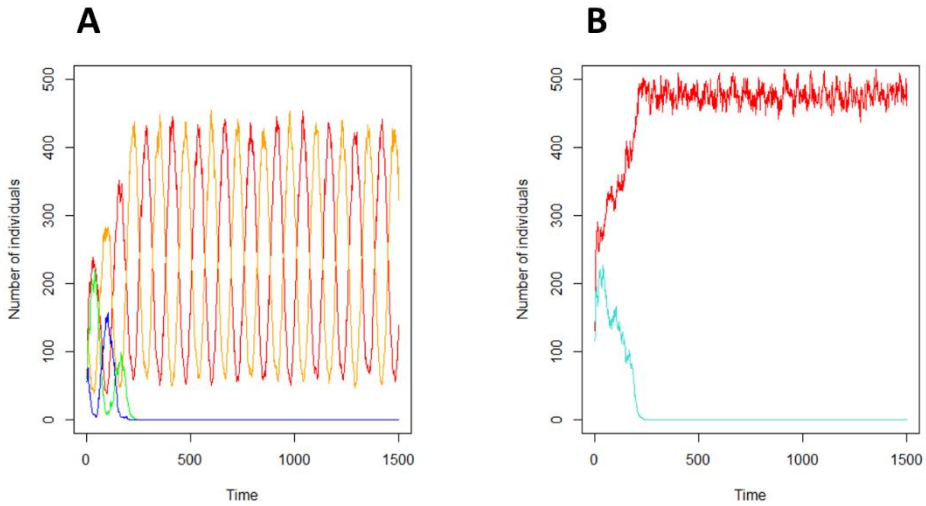

Figure 2: Representative simulations. A: Absolute numbers of individuals of each genotype over time: mutator adapted to resource A (red), non-mutator adapted to resource A (green), mutator adapted to resource B (yellow), and non-mutator adapted to resource B (blue). The non-mutator strains go extinct early on in the simulation, the abundance of the two mutator-strains fluctuates in accordance with the fluctuations of the resources that they are adapted to. B: Absolute number of mutators (red) and non-mutators (blue). The non-mutator strain genotype lost early on in the simulation.

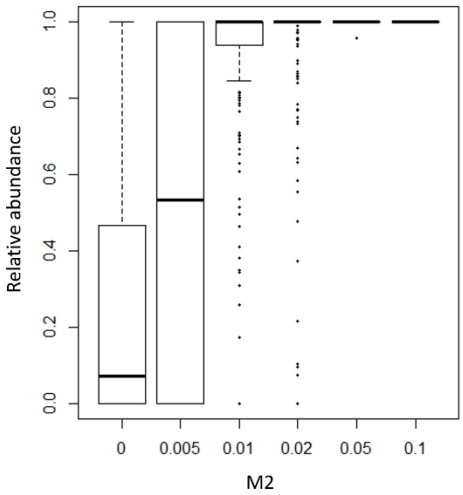

Figure 3: Frequencies of the invading mutation rate M2 at the simulation end points. This figure shows the frequencies of the invading mutation rate at the end of the simulation, for one resident mutation rate (M1 = 0.005), slow environmental oscillations (s = 0.05) and six different invading mutation rates, with 200 replicates each. For each value of M2, the mutation rate with the higher value, be it resident or invading, achieves a higher frequency at the end of the simulation. In cases where resident and invading mutation rate are of equal value (M1 = M2 = 0.005), the frequencies of the invading mutation rate are distributed evenly around 0.5.

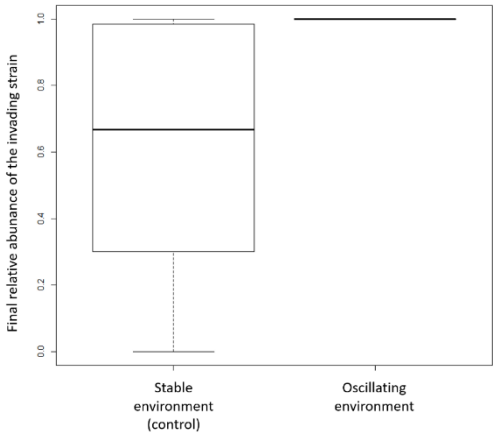

Figure 4: Frequencies of the invading mutation rate at the simulation end points. This figure shows the frequencies of the invading mutation rate at the end of the simulation, evolved under an oscillating environment and evolved under the stable environment of the control. The parameter values used here are: M1 = 0, M2 = 0.1, s = 0.05. The spread of the higher mutation rate is strongly reduced or absent in the control, demonstrating the crucial role that an variable environment plays in selecting for higher mutation rates.

Varying the parameters influences the speed at which higher mutation rates can invade. Firstly, a larger difference between the invading and resident mutation rates leads to a more complete invasion after 1500 time units (see Figure 5A). Secondly, slower oscillation speeds allow a faster and more complete invasion, both for each mutation rate difference (see Figure 5A) and overall (see Figure 5B).

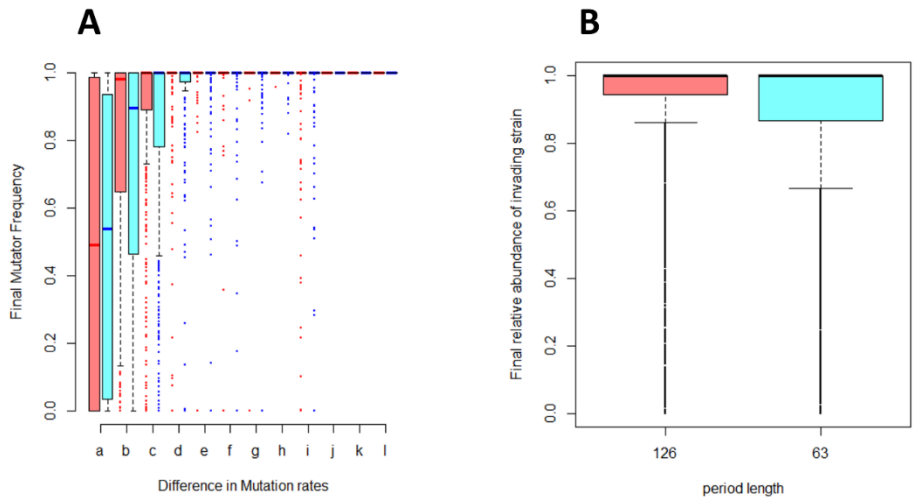

Figure 5: Effect of the difference between resident and invading mutation rate and of the environmental oscillation speed on mutation rate invasion. Figure 5A: This figure shows the frequency of the mutator strain at the end of the simulation, for varying differences between the invading and the resident mutation rate. The frequencies of the mutator are higher for larger differences in mutation rate, this pattern is more pronounced under slower environmental oscillations (s = 0.05, red) than under fast environmental oscillations (s = 0.1, blue).The following differences in mutation rate are displayed: a = 0, b = 0.005, c = 0.01, d = 0.015, e = 0.020, f = 0.030, g = 0.040, h = 0.045, i = 0.050, j = 0.090, k = 0.095, l = 0.10. Figure 5B: Combining the entire data set shows that overall, slower oscillations (s = 0.05, red) allow a faster and more complete invasion of the higher mutation rate than fast oscillations (s = 0.1, blue).

Negative frequency dependent selection

The four genotypes (mutator adapted to resource A, mutator adapted to resource B, non-mutator adapted to resource A, non-mutator adapted to resource B) start at equal frequencies and then evolve for 1500 time units. Figure 6 shows an example of how these four genotype frequencies change through time: in this case, the non-mutator strains decline fast and finally are lost towards the end of the simulation, whereas the mutator strains persist. Of the remaining strains, one is adapted to resource A and one is adapted to resource B, and their frequencies therefore show slight oscillations consistent with negative frequency dependent selection.

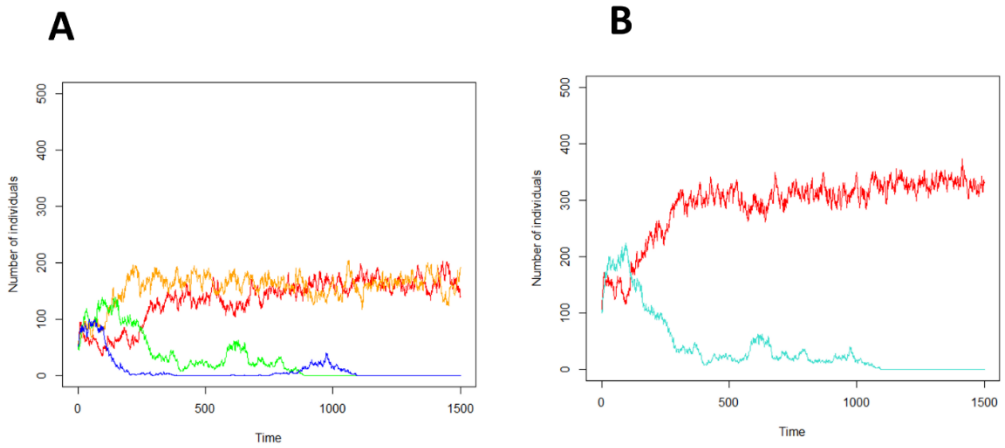

Figure 6: Exemplary simulation behaviour. Figure 6A: This figure shows the absolute numbers of individuals of each genotype over time: mutator adapted to resource A (red), non-mutator adapted to resource A (green), mutator adapted to resource B (yellow), and non-mutator adapted to resource B (blue). The abundance of the two non-mutator strains declines fast and they ultimately go extinct towards the end of the simulation, the abundance of the two mutator strains fluctuates in a manner consistent with negative frequency dependent selection. Figure 6B: This figure shows the absolute number of mutators (red) and non-mutators (blue). The non-mutator strains is lost towards the end of the simulation.

I analyzed the effect of mutation rate on spread of the genotype. To do so, I considered the number of individuals of each genotype at the end of the simulation runtime, comparing the frequency of the mutator strains and the non-mutator strains. Overall, the mutator strains are able to successfully outcompete the non-mutator strains, with the mutator-strains being over-represented at the ends of most replicates (for an example, see Figure 7, for plots showing all invader/resident comparisons see Appendix Figure 2). The proportion of the invading mutation rate is significantly different form 0.5, this pattern holds for most parameter combinations regardless of the values chosen for the M1 and M2, as long as one mutation rate is higher. A larger difference between M1 and M2 leads to a more complete invasion after 1500 time units (see Figure 8). If both mutation rates are of equal value, the differences in the final numbers of individuals with either mutation rate are caused by drift, and are not statistically significant. In this model, it was not possible to implement a control with stable, constant selection pressures: the fluctuating selection pressures result from an intrinsic negative frequency dependency rather than from extrinsically implemented fluctuations in resource abundance.

Overall, the spread of higher mutation rates under negative frequency dependent selection is clear, but less striking than under environmental oscillations.

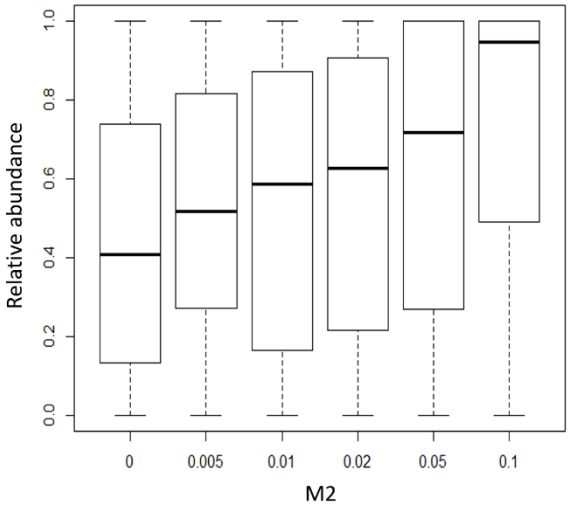

Figure 7: Frequencies of the invading mutation rate M2 at the simulation end points. This figure shows the frequencies of the invading mutation rate at the end of the simulation, for one resident mutation rate (M1 = 0.005) and six different invading mutation rates, with 200 replicates each. For each value of M2, the mutation rate with the higher value, be it resident or invading, achieves a higher frequency at the end of the simulation. In cases where resident and invading mutation rate are of equal value (M1 = M2 = 0.005), the frequencies of the invading mutation rate are distributed evenly around 0.5.

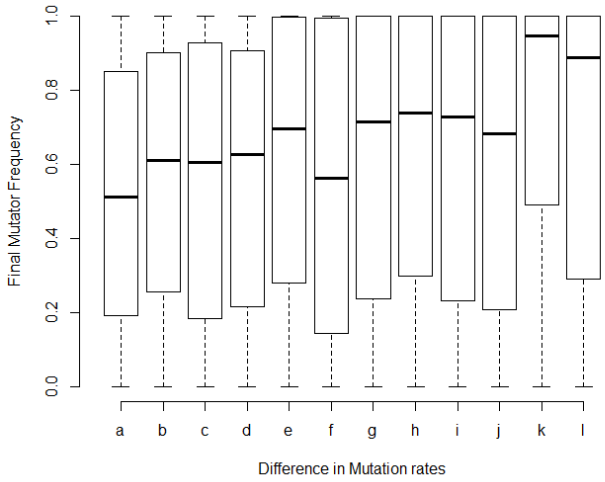

Figure 8: This figure shows the frequency of the mutator strain at the end of the simulation, for varying differences between the invading and the resident mutation rate. The frequencies of the mutator are generally higher for larger differences in mutation rate – however, this pattern is not as clear as it was for mutation rate evolution under environmental oscillations (see Figure 5A). The following differences in mutation rate are displayed: a = 0, b = 0.005, c = 0.01, d = 0.015, e = 0.020, f = 0.030, g = 0.040, h = 0.045, i = 0.050, j = 0.090, k = 0.095, l = 0.10.

Conclusions Model 1

Overall, Model 1 shows that under periodically alternating selection pressures arising from environmental oscillations, higher mutation rates can efficiently invade. Slower environmental oscillations, and larger differences between invading and resident mutation rates, expedite this invasion. A similar though less efficient invasion of higher mutation rates can also be found under alternating selection pressures arising from negative frequency depending selection.

Model 2: Evolution of mutation rate in a periodically changing environment

Model description

Model structure

In the second model I use a simulation written in C++ to explore how the mutation rate evolves under selection for evolvability, by tracking a population of individuals in a habitat which is modelled spatially explicit as a 50x50 grid with a total of 2500 cells. Each cell can be occupied by only one individual, the maximal population size is therefore 2500.

The second model bears several similarities to the one described in model 1: Again, each individual has two gene loci, one determining ecological adaptation either to resource A or to resource B, and one determining the mutation rate. The former locus operates in a similar manner to the ecological adaptation locus from model 1: depending on the allele at this locus the individual is either adapted to resource A or B, and has a birth rate proportional to the fitness payoff from the chosen resource. The fitness payoffs from the two resources fluctuate between zero and one, with the time-dependant fitness payoff from allele A given by

$$w = {sin(time * s * \frac{\pi}{2}) + 1}/{2}$$

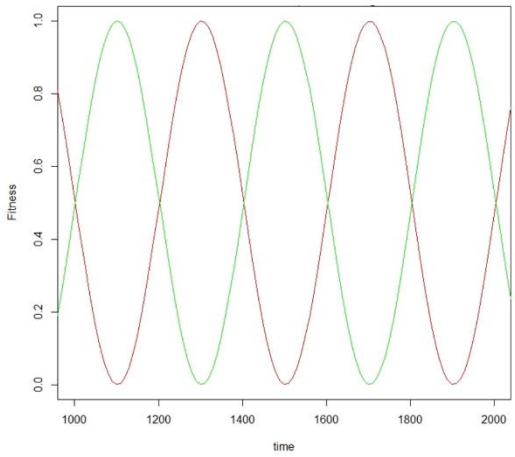

Figure 9: Changes in resource specific fitness over time. This figure shows the fitness pay-offs for individuals adapted to resource A (red) or B (green). These fitness pay-offs are normalized to lie between 0 and 1, and are here shown for environmental oscillations s = 0.01 (period length = 400 time steps). The fitness-payoffs are out of phase so that the highest pay-off from resource A coincides with the lowest pay-off from resource B. This creates a variable selection regime with alternating selection pressures.

The parameter s refers to the oscillation speed. The fitness payoff from allele B is given by the same formula but shifted by half the period length, causing the fitness payoffs to fluctuate out of phase (see Figure 9). As in model 1, the two ecological adaptation alleles are thus selected in an alternating manner, as might be the case under, for example, seasonal selection pressures alternating between winter and summer, or wet and dry season. Multiplying the current resource-dependent fitness payoff with the unmodified birth rate gives the realised birth rate. Unlike in the previous model, the spatially explicit implementation allows for the carrying capacity to be implemented independently of the oscillating resources: As each cell can only be occupied by one individual, a reproduction event can only be completed if there is an empty cell within the neighbourhood (i.e. surrounding cells, see Figure 10) of a reproducing focal individual. Effectively, this decreases the birth rate at high population sizes and in densely populated areas.

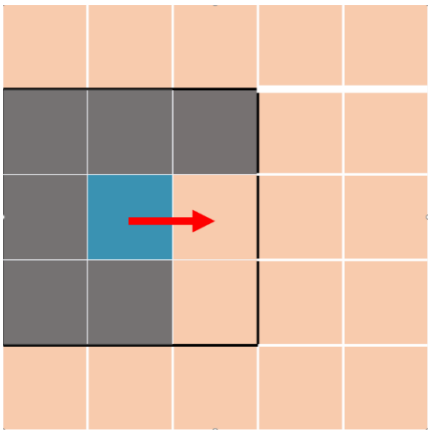

Figure 10: Implementation of the neighbourhood. This figure shows a habitat modelled as a grid of size 5 x 5. Occupied cells are marked in grey, the cell occupied by the focal individual is marked in blue. The focal individual's neighbourhood is indicated by a black square around it. Shown here is a neighbourhood with the neighbourhood size Nb = 1, i.e. the neighbourhood extends exactly one square above, below, to the right and to the left of the focal individual. The focal individual can only reproduce if there is at least one empty cell within the neighbourhood, this cell will then be occupied by the new offspring. In this case, the focal individual's offspring can occupy either of the two orange cells in the right bottom corner of the black neighbouhood square. The arrow indicates one possible cell that the focal individual might choose for reproduction.

The second locus encodes the mutation rate, that is the rate at which the ecological adaptation locus mutates from being adapted to resource A to being adapted to resource B, with no lethal mutations. The implementation of this mutation rate is the key difference between this model and the previous one: rather than investigating the invasion of a mutator-strain, the mutation rate is free to evolve any biologically sensible value (between 0 and 1). This mutation rate determines the likelihood of a mutation at the ecological adaptation locus only: the rate at which the mutation rate itself mutates is implemented separately and kept constant, irrespective of the value taken by the mutation rate. Mutation step sizes are randomly drawn from a normal distribution centered around zero with a standard deviation of 0.005. The mutation rate is capped at zero and at one.

As in the first model, the individuals reproduce asexually with no recombination: during each birth event, the new individual inherits its parent's mutation rate (or a slightly different mutation rate if there is a mutation of the mutation rate), and its parent's ecological adaptation allele if there is no mutation (or the opposite adaptation allele if there is a mutation of the ecological adaptation allele).

During each time step, an individual is randomly chosen to undergo a birth event (depending on the realized birthrate and availability of space) and subsequently, a death event (depending on the death rate). This is followed by a second round of death events, occurring with different probabilities in different simulations, allowing me to investigate the effect of death rate and thus population turnover. Any dead individuals are removed from the population, thus freeing up the cell they were occupying. Now, the current focal individual has undergone all potential events for this timestep, and a new focal individual is sampled from the population (without replacement), which undergoes the same events outlined above. New focal individuals are sampled until every individual in the population has acted as the focal individual exactly once. After this, time is advanced to the next timestep.

The results from this model with selection for evolvability were compared with those from a control simulation with no selection for evolvability: in the control scenario the sine curves, describing the oscillations in resource availability, are replaced with flat lines for both resource A and B. Without the environmental fluctuations, the payoffs of both resources are constant and equal to each other at all times – mutations from A to B or B to A therefore have no fitness consequences, and the mutation rate evolves under drift alone. The fitness pay-off for both resources is kept at 0.5 in the control, an average of the fitness values given by a sine curve between zero and one.

Parameter values

I used the following values for the parameters and properties of initial individuals: a death rate of 0.05 for the first round of death events, an unmodified birth rate of 0.9 (the realized birth rate is lower and depends on the availability of resource A or B and on the availability of space), an evolvable mutation rate initialized at a random value between 0.0010 and 0.0011, and a mutation rate specific mutation rate of 0.01 (the latter mutation rate is constant regardless of the value which the evolvable mutation rate takes). The population is initialized with hundred individuals placed in the center of the habitat in a 10 by 10 square, and is then allowed to evolve for 5000 timesteps.

I tested the influence of the following parameter values on model behaviour: the neighbourhood size Nb (Nb = 1; Nb = 50), the speed of environmental oscillations s (s = 0.01; s = 0.05), and the death rate during the second round of death events d (d = 0, d = 0.001, d = 0.05, d = 0.075). The neighbourhood size determines how many rows the neighbourhood extends from the focal individual in any direction (see Figure 10). At the maximal neighbourhood size Nb = 50, i.e. the size of the entire habitat, the population is homogeneously mixed and has no spatial structure.

Statistical analysis

I analysed the model behaviour statistically using a correlation analysis. The change in average mutation rate was calculated every 10 time steps as the average mutation rate after the next 10 time step minus the average mutation rate in the current time step, divided by the number of timesteps (i.e. 10). Data from the first 500 time steps were excluded, to avoid potential confounding effects resulting from initialising a small population which still needs to establish itself. Using data from timesteps 500 till 5000, I calculated the Pearson coefficient for the correlation between the change in average mutation rate and the number of individuals. For this analysis, I separated the data of individuals with different genotypes, i.e. I analysed the correlation between the number of individuals with genotype A and the change in mutation rate averaged over all individuals with genotype A, and analysed separately the correlation between the number of individuals with genotype B and the change in mutation rate averaged over all individuals with genotype B. I repeated this analysis for 30 replicates.

Simulation results

I tested 24 parameter combinations, 16 with environmental oscillations (the test simulations) and 8 without environmental oscillations (the control simulations). Under environmental oscillations, most parameter combinations show a marked increase in the population average of the mutation rate, as well as a striking pattern of step-wise growth (for an example see Figure 16).

Mutation rate evolution and population dynamics

Teasing apart the underlying population dynamics can explain this observed characteristic step-wise increase. Focusing on allele A shows that when the fitness gained from resource A is high, individuals with genotype A are abundant (Figure 11) and most new individuals born with allele A are the non-mutated offspring of parents with allele A (Figure 12). Conversely, when individuals with genotype A are scarce, most new individuals born with allele A are the mutated offspring of parents with allele B (Figure 12), and the mutation rate – averaged over all individuals with genotype A – increases (Figure 12). This holds for both alleles: when individuals with a particular genotype are abundant, the average mutation rate stays constant or decreases slightly, whereas when individuals with a particular genotype are rare, the average mutation rate shows a marked increase (Figure 13). Average mutation rate refers here to the mutation rate averaged across those individuals which have a particular genotype. The step-wise increase in average mutation rates for A and B occurs out of phase with each other (Figure 13) – when combining them, and averaging the mutation rate over the entire population, the population-wide step-wise increase seen in all simulations is recovered (Figure 14), with two steps per environmental oscillation period.

I performed a correlation analysis using Pearson's correlation coefficient to give statistical support to the observation that when a genotype is rare, the average mutation rate of individuals with this genotype increases. The correlation between the change in mutation rate (averaged across all individuals with a specific genotype) and the number of individuals with this genotype was highly significant and negative, for both the genotype A and B, and in all 30 replicates. This confirms that the dynamics explained above and in Figures 13 – 16 are indeed consistent across replicates. See Appendix Table 1 for a table of the p-values and correlation coefficients. I performed this analysis using the parameter values s = 0.01, Nb = 1, d = 0.

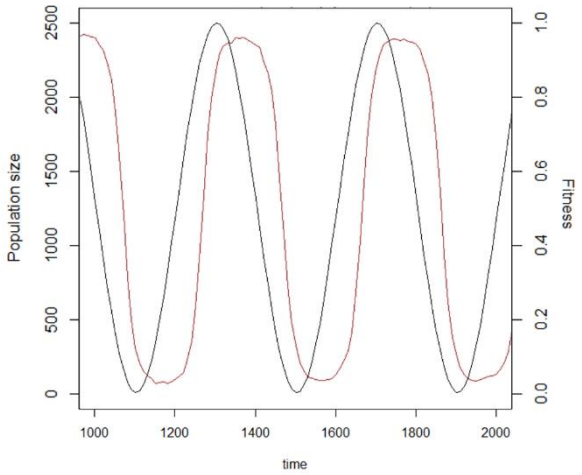

Figure 11: Linking allele fitness and abundance. This plot shows the change in allele fitness (black) and abundance (red, in absolute numbers) for allele A. The number of carriers of allele A tracks the fitness of allele A with a slight lag.

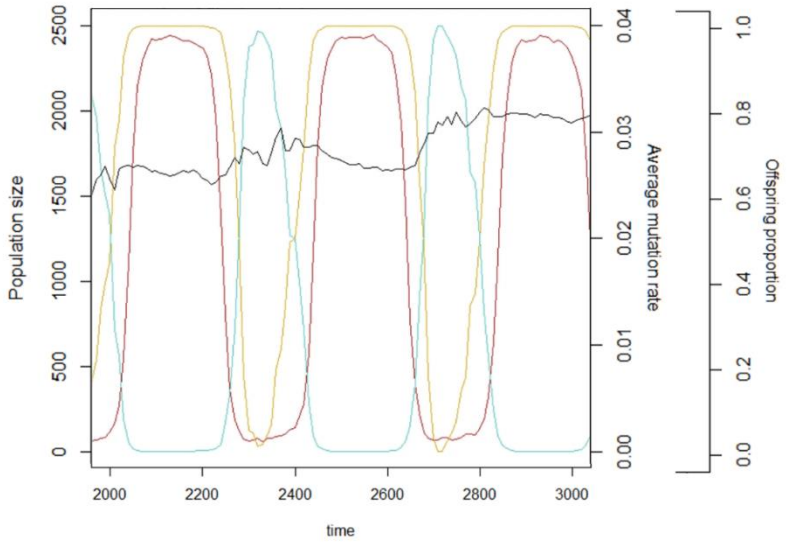

Figure 12: Linking allele carrier abundance, parental genotypes and mutation rate. This plot shows the abundance of carriers of allele A (red), the proportion of new A-carriers born from A-carriers (yellow), the proportion of new A-carriers born from B-carriers (blue), and the mutation rate averaged across all A-carriers (black). The average mutation rate increases when individuals carrying A are scarce and most newly-born A-carriers are the mutated offspring of B-carriers. When A-carriers are abundant and most newly born A-carriers are the offspring of other A-carriers, the mutation rate stagnates or falls.

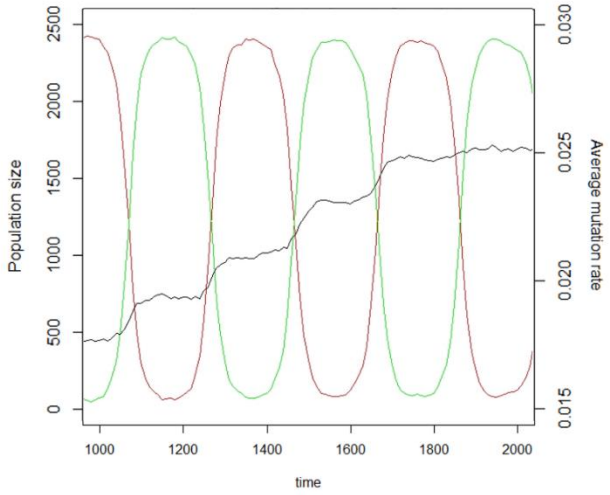

Figure 13: Linking allele carrier abundances and mutation rates. This plot shows the abundance of carriers of allele A (red), the abundance of carriers of allele B (green), the mutation rate averaged across all A-carriers (yellow), and the mutation rate averaged across all B-carriers (purple). The mutation rate amongst A-carriers increases when they are scarce and stagnates or falls when they are abundant. The same pattern can be seen for the mutation rate amongst, and abundance of, B-carriers. The mutation rates averaged across A-carries and B-carriers fluctuate out of phase.

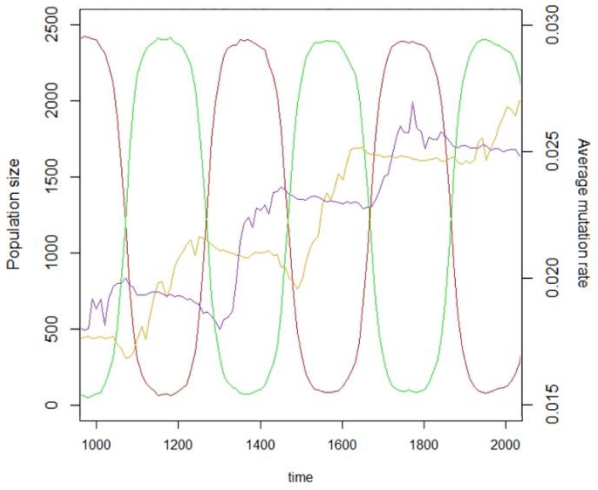

Figure 14: Linking allele abundances and population-wide mutation rate averages. This plot shows the abundance of carriers of allele A (red), the abundance of carriers of allele B (green), the mutation rate averaged across the entire population (black). When either abundance decreases, the mutation rate increases in a step wise manner.

Dependence on parameters

The increase of average mutation rate is influenced by various parameters. The increase is steeper for a higher death rate (see Figure 15) and for slower environmental oscillations (see Figure 17). These two parameters interact: a decreased death rate slows down mutation rate evolution under fast environmental oscillations (see Figure 15) but not under slow environmental oscillations (see Figure 16), indicating that these parameters can compensate for one another to some extent. Neighbourhood size does not affect mutation rate evolution (see Figure 18). The figures shown in this section are examples illustrating the influence of the tested parameters, not all explored parameter combinations are displayed. Plots of all parameter combinations and controls are available in the data archive. Overall, 12 out of 16 parameter combinations show a clear increase in mutation rate as well as behavior evidently different from the control simulations. No increase in mutation rate can be observed only when fast oscillations are combined with low death rates. Finally, the controls did not

show an increase in mutation rate, beyond that which is consistent with drift and a bounded random walk (see Appendix Figure 3).

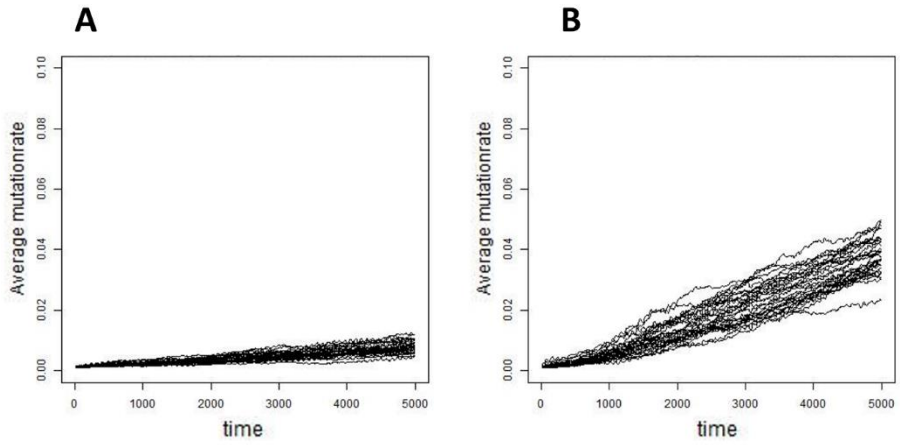

Figure 15: The effect of death rate on mutation rate evolution under fast environmental oscillations. This figure shows the increase in the population average of the mutation rate over the course of the simulation, for two different death rates with 30 replicates each. A low death rate (Figure 15A, d = 0) does not allow the mutation rate to rise as far as a high death rate (Figure 15B, d = 0.05). Parameters other than death rate are kept constant (s = 0.05, Nb = 1).

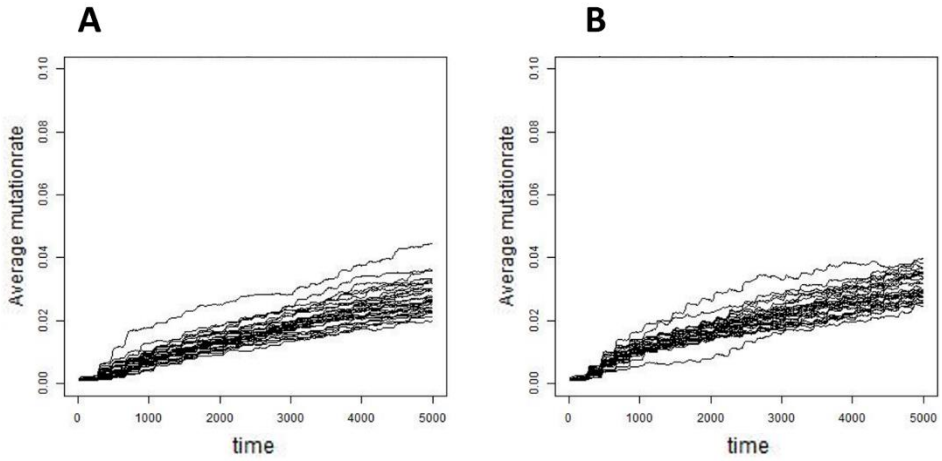

Figure 16: The effect of death rate on mutation rate evolution under slow environmental oscillations. The figure shows the increase in the population average of the mutation rate over the course of the simulation, for two different death rates with 30 replicates each. In contrast to figure 15, there is no observable difference in the pattern of mutation rate evolution under a lower death rate (Figure 16A, d = 0) or a high death rate (Figure 16B, d = 0.05). Parameters other than death rate are kept constant (s = 0.01, Nb = 1).

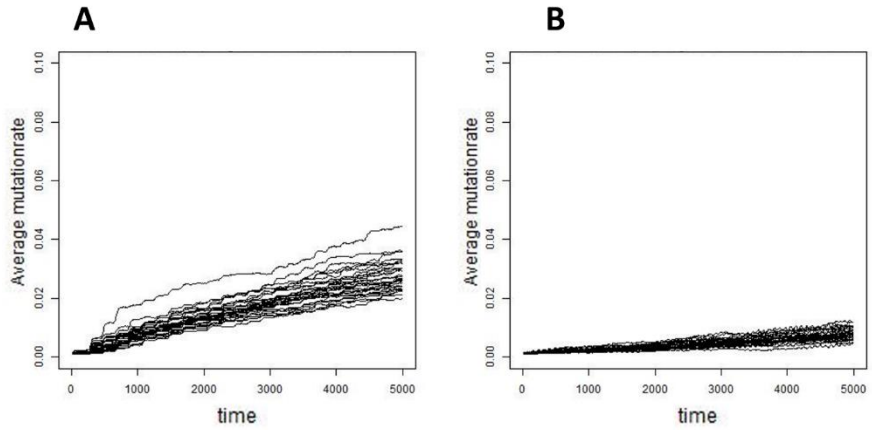

Figure 17: The effect of environmental oscillation speed on mutation rate evolution. This figure shows the increase in the population average of the mutation rate over the course of the simulation, for two different environmental oscillation speeds with 30 replicates each. Slow environmental oscillations (Figure 17A, s = 0.01) allow the mutation rate to rise further than fast environmental oscillations (Figure 17B, s = 0.05). Parameters other than the speed of environmental oscillations are kept constant (d = 0, Nb = 1).

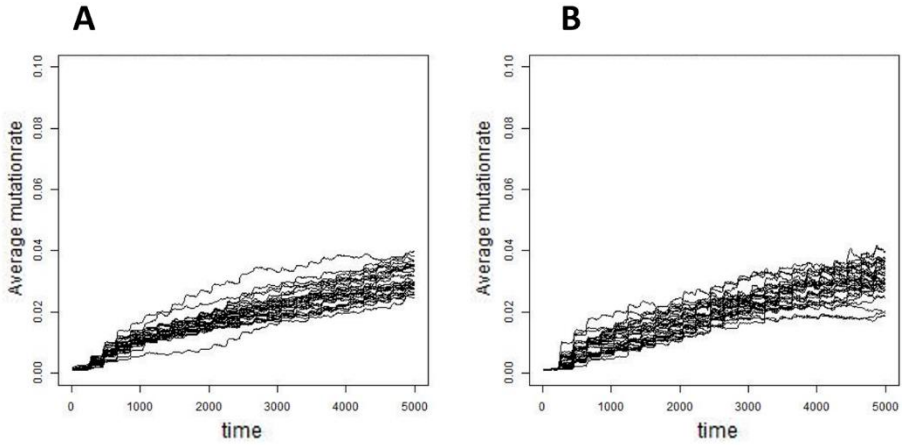

Figure 18: The effect of neighbourhood size on mutation rate evolution. This figure shows the increase in the population average of the mutation rate over the course of the simulation, for two different neighbourhood sizes with 30 replicates each. There is no observable difference in the pattern of mutation rate evolution under small neighbourhood sizes (Figure 18A, Nb = 1) or large neighbourhood sizes (Figure 18B, Nb = 50). Parameters other than the neighbourhood size are kept constant (s = 0.01, d = 0.05).

Conclusions Model 2

Overall, Model 2 shows that under periodically alternating selection pressures arising from environmental oscillations, the population average of the mutation rate increases, in a striking step-wise pattern. The mutation rate, averaged across all carriers of a specific allele, increases when the allele is selected against. Slower environmental oscillations and a higher death rate, allowing a higher population turnover, allow steeper mutation rate growth.

Model 3: Evolution of mutation rate in a complex, variable environment

Model description

The third model, also written in C++, explores how the mutation rate evolves under selection for evolvability, similar to the second model. However, this model considers the evolution of mutation rate under the more realistic conditions of recombination, multi-locus genetics, non-linear genotype-phenotype mapping and random environmental fluctuations, as well as implementing sexual reproduction.

Model structure

Some aspects of the second model were retained in the third: The habitat is again modelled spatially explicit as a 50x50 grid with a total of 2500 cells with maximally one individual per cell (resulting again in a maximal population size of 2500). Most traits and parameters are implemented in the same way: again, there is a globally fixed birth rate, two population-wide fixed death rates determining the chance of dying during the two rounds of death events, an individually evolvable mutation rate, and a globally fixed rate at which the mutation rate mutates. However, this model has three crucial differences: the non-linear genotype-phenotype map, the fitness seascape undergoing random oscillations, and sexual reproduction with recombination.

Genotype-phenotype mapping

There are 20 loci involved in determining the ecological adaptation, each of which has two alleles (1 and 0). The ecological adaptation loci are represented by a bit string with 20 positions. Each position is a locus, with the possible binary alleles 1 and 0. This 20 bit long string is interpreted as a binary number and translated into one single decimal number, which represents the individual's phenotype. For example, a genotype of 00110011101011101011 is translated into the phenotype of 211691. This way, 1048576 different phenotypes can be encoded (1048575 being the largest decimal number that can be represented as a bit string with 20 positions). The genotype to phenotype mapping is therefore non-linear and allows for a variety of mutational step sizes: a mutation at the 20^{th} position of the bit string is far less influential than a mutation at the first position. This was a deliberate choice: this genotype-phenotype map not only reflects natural complexity and variation in mutation step sizes, but is also an important factor favoring the evolution of evolvability. The link between various types of genotype-phenotype mapping, evolvability and the evolution of evolvability is well established (see for example Crombach & Hogeweg 2008 and Pigliucci 2010). A mixture of small and large mutation step sizes allow mutator strains to explore their immediate surroundings in the fitness landscape, as well as far away peaks and valleys. This is especially important under random environmental fluctuations: At times when the population is strongly maladapted, large mutation steps are required to find the optimal phenotype, and the fitness difference between mutators and non-mutators is especially stark – it is in such situations that the selection for higher mutation rates is especially strong, and a non-linear genotype-phenotype map which facilitates these situations is therefore an important aspect of the evolution of mutation rates.

Phenotype-fitness mapping and the fitness seascape

To implement selection for evolvability, I modelled random environmental fluctuations, using a fitness seascape that changes in the following way: one peak and one valley slowly appear and grow until they reach their maximal size, then they slowly shrink again until the fitness landscape is flat. After this, a new peak and a new valley appear in a new location. The fitness seascape at any given point is represented by a sine curve, mapping a fitness value between zero and one to each phenotype:

$$w = \sin((\frac{2 * \pi}{1048575} * \text{phenotype} + \alpha) * G(t)) + 1/2$$

The term $\frac{2*\pi}{1048575}$ normalizes the sine curve to give fit exactly one sine period length over all phenotypes (1048575 being the maximal phenotype), resulting in a fitness seascape with maximally one optimum (see Figure 19A).

The parameters G(t) and α allow this fitness seascape to change randomly with time. G(t) depends on a "time factor", a variable between zero and four, which grows by a fixed small amount *s* in each time step. For example, for s = 0.01, the time factor would grow starting from time factor = 0 in timestep zero, to time factor = 4 in timestep 400, and would be re-set to time factor = 0 in timestep 401. The parameter s therefore gives the oscillation speed. For 0 < time factor < 2:

$$G(t) = 1 - time\ factor$$

causing the sine curve to change gradually first to a flat line and then to – sin(x). For 2 < time factor < 4:

$$G(t) = time\ factor - 3$$

resulting in the opposite movement, a curve changing gradually from – sin(x) first to a flat line and then to sin(x). This way, the fitness function encodes a fitness seascape with slowly appearing and disappearing peaks and valleys. Furthermore, the locations of each new peak and valley are randomized, in order to select for different phenotypes at different time points. To do so, the parameter α is randomly re-assigned a new value each time the fitness seascape is flat (i.e. when G(t) = 0). This way, each new growing fitness peak is shifted to a new phenotype.

Finally, the fitness value given by the function w is modified depending on the selection strength k, a parameter determining the gradient of the fitness peak. Under strong selection (k = 3), the sine curve is taken to the power of 3, creating a steeper peak (see Figure 19B). The selection is thus more rigorous, and even small deviations from the optimal phenotype are penalized with greater severity. Under weak selection (k = 1), the sine curve is not modified, i.e. taken to the power of 1.

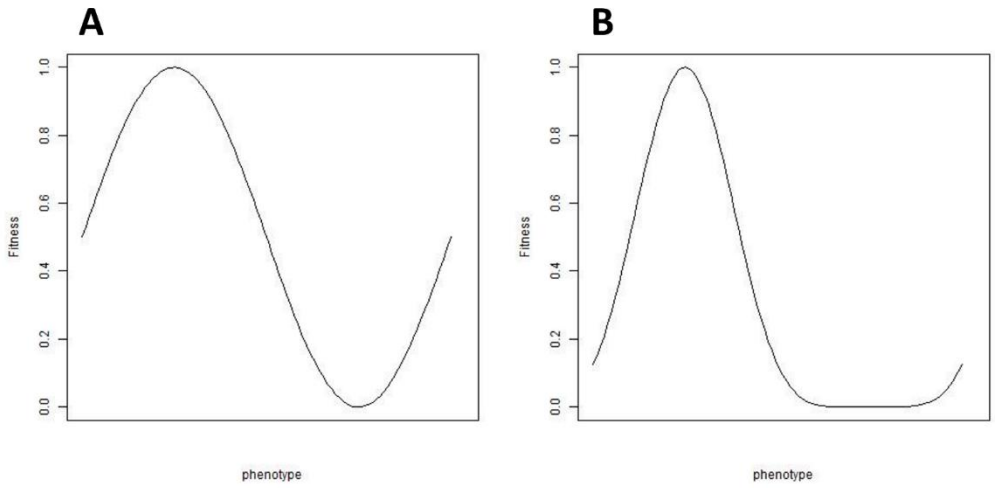

Figure 19: The phenotype-fitness map. This figure shows a snapshot of the phenotype-fitness map, or fitness seascape, at one particular point in time. A sine function is used to map a fitness value to each phenotype. At this particular time, the fitness seascape has one high hill and one deep valley. Over time, the hill and valley will disappear and be replaced with new hills and valleys. Figure 19A shows the phenotype-fitness map under weak selection, i.e. k = 1. Figure 19B shows the phenotype-fitness map under strong selection, i.e. k = 3: the sine curve is taken to the power of 3 resulting in a steeper fitness gradient. The selection is thus more rigorous, and even small deviations from the optimal phenotype are penalized with greater severity.

Sexual reproduction

As in model 2, birth events occur at a rate determined by the realized birth rate (i.e. the globally fixed birth rate multiplied by the individual's fitness), with the added condition of an empty cell in the neighbourhood for the offspring to occupy. Under sexual reproduction, there is a second condition: the presence of another individual within the neighbourhood, to act as the focal individual's mate. I did not implement discrete sexes or minimal ages of sexual maturity, hence any individual that has not undergone a death event may act as a mate. If there is more than one empty cell or more than one potential mate, then one of them is chosen randomly as location of the offspring, or as mate, respectively.

Once a mate has been found, the resulting offspring's traits need to be determined. First, the offspring's ecological adaptation loci are assigned, by randomly choosing one of the two parental sets of ecological adaptation loci for the offspring. With a probability equal to the recombination rate r, a recombination event occurs prior to inheritance. During recombination, a short sequence is exchanged between the parental bit strings. The start position is randomly drawn, the length of the sequence is randomly drawn between one and five. Furthermore, with a probability equal to the focal individual's mutation rate, the offspring's ecological adaptation loci contain a mutation. During mutation, one position of the bit string is randomly selected and flipped (i.e. a 0 becomes a 1 and a 1 becomes a 0), with no lethal mutations. The mutation rate is per "genome", not per "gene" (i.e. in each timestep there is the potential for one mutation event per individual per bit string, not per bit string position). Next, the offspring's mutation rate is determined. As in model 2, "mutation rate" refers to the ecological-adaptation-loci specific mutation rate, an evolvable property of the individual which is implemented separately from the rate at which the mutation rate itself mutates. The offspring inherits the mutation rate and ecological adaptation loci from the same parent, unless a second recombination event occurs, in which case the mutation rate and ecological adaptation loci are inherited from different parents. After inheritance, the ecological-adaptation-loci specific mutation rate can itself mutate, with a probability equal to the mutation-rate-specific mutation rate. Mutation step sizes are randomly drawn from a normal distribution centered around zero with a standard deviation of 0.005, and the mutation rate is capped at zero and at one. The offspring is now ready to be added to the population.

Parameter values

All fixed parameter settings from model 2 were kept unchanged, except for the simulation run time which increased to 10000 timesteps. In addition to this, the ecological adaptation loci in the initial population are set by randomly choosing a decimal number for each individual between 0 and 1048575 (the maximal decimal number that can be represented as a bit string with 20 positions), and translating it into a bit string. Individuals at later stages obtain their ecological adaptation loci from their parents, they therefore do not need to be set externally.

The results from this model with selection for evolvability were compared with those from a control with no selection for evolvability: similar to the control from model 2, the fitness seascape is replaced with a flat fitness landscape without environmental fluctuations. All phenotypes receive equal fitness, mutations have no fitness consequences, and the mutation rate evolves under drift alone. The fitness pay-off was kept at 0.5 at all times, an average of the fitness values given by a sine curve between zero and one.

As in model 2, I studied how model behaviour is influenced by the neighbourhood size Nb, the speed of environmental oscillations s, and the death rate d, although the parameter s is encoded slightly differently in this model, due to the more complex phenotype-fitness mapping. In addition to these

parameters, I studied the effect of the selection strength k and the recombination rate r. I used the following values: Nb = 1, Nb = 10, Nb = 50; s = 0.01, s = 0.05, s = 0.1; d = 0, d = 0.001, d = 0.05; k = 1, k = 3; and r = 0, r = 0.01, r = 0.3. At the maximal neighbourhood size Nb = 50 the population is homogeneously mixed and has no spatial structure. Overall, I tested 90 different combinations of parameter values. Furthermore, I tested these parameter combinations in several different variations of this model – with and without selfing, with bitstrings with 10 or 20 positions as ecological adaptation loci, and with different mutation mechanisms. Since I could not identify a qualitative difference between those scenarios, I selected one scenario (no selfing, 20 positions, capped mutation rates) for a more in-depth quantitative analysis.

Statistical analysis

I analysed the model behaviour statistically using a correlation analysis, similar to the analysis carried out in model 2. The change in average mutation rate was calculated every 10 time steps as the average mutation rate after the next 10 time step minus the average mutation rate in the current time step, divided by the number of timesteps (i.e. 10). Data from the first 500 time steps were excluded, to avoid potential confounding effects resulting from initialising a small population which still needs to establish itself. Using data from timesteps 500 till 10000, I calculated the Pearson coefficient for the correlation between the change in average mutation rate and the population size. I repeated this analysis for 30 replicates. Furthermore, I analysed how those parameters that were varied between simulations affected the outcome. To do this, I carried out an anova on the mutation rates obtained at the end of the simulations for different parameter values, with 30 replicates for each parameter value. I used Tukey's range test for post-hoc analysis.

Simulation results

Overall, the mutation rate showed a striking increase over the course of the simulation for most parameter combinations. This increase was not seen in the control simulations without environmental oscillations, which instead showed behavior consistent with drift and a bounded random walk (see Appendix Figure 4).

I performed a correlation analysis using Pearson's correlation coefficient, to test the correlation between the change in mutation rate (averaged across the entire population) and the population size. In 26 out of 30 replicates, this correlation was significantly negative, showing that the mutation rate increases at times of low population sizes, i.e. during population crashes. There were no significantly positive correlations. For a table showing correlation coefficients and p-values see Appendix Table 2. I performed this analysis using the parameter values s = 0.01, Nb = 1, d = 0.05, k = 1, and r = 0.

To what extent the mutation rate increases depends on the parameter values. Slower oscillations (Figure 20 & Figure 21), higher death rates (Figure 22 & Figure 23), stronger selection (Figure 24 & Figure 25), and lower recombination rates (Figure 26 & Figure 27) all lead to higher mutation rates at the end of the simulation run time. No effect can observed for neighbouhood size (Figure 28 & Figure 29). Surprisingly, even under non-zero recombination rates, the mutation rates obtained at the end of the simulation are significantly higher than in the control (Figure 30). I tested different death rates by varying the death rate in the second round of death events whilst keeping the death rate in the first round constant, all values for death rates given in this section therefore refer to the death rate during the second round of death events only. Certain parameter combinations caused the populations to go extinct, in which case the final average mutation rate was not taken into consideration, as such averages are liable to be misleading. When strong selection was combined either with slow oscillations

or with high death rates, most or all replicates went extinct relatively early in the simulation. The figures shown in this section are example plots illustrating the influence of the tested parameters, not all explored parameter combinations are shown here. Plots of all parameter combinations and controls are available in the data archive.

Regardless of the parameter values, the increase in mutation rate does not proceed along a linear or repeatable trajectory, but shows instead irregular fluctuations, periods of stasis interspersed with brief bursts of steep increases or decreases (see Figure 31). This is also evident in the "fuzzy" appearance of those plots showing the average mutation rate over time for 30 replicates (Figures 20, 22, 24, 26, 28). This pattern of irregular fluctuations and burst-like evolution is typical for genes evolving via genetic hitchhiking.

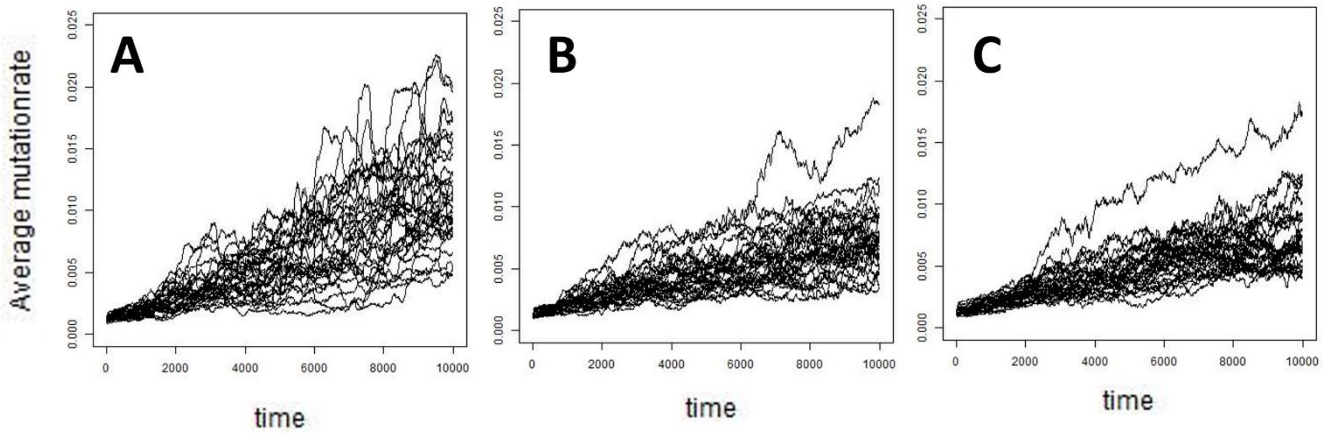

Figure 20: The effect of environmental oscillation speed on mutation rate evolution. This figure shows the increase in the population average of the mutation rate over the course of the simulation, for three different environmental oscillation speeds with 30 replicates each. Slow environmental oscillations (Figure 20A, s = 0.01) allow the mutation rate to rise further than medium (Figure 20B, s = 0.05) or fast environmental oscillations (Figure 20C, s = 0.1). Parameters other than the speed of environmental oscillations are kept constant (d = 0, Nb = 1, k = 1, r = 0).

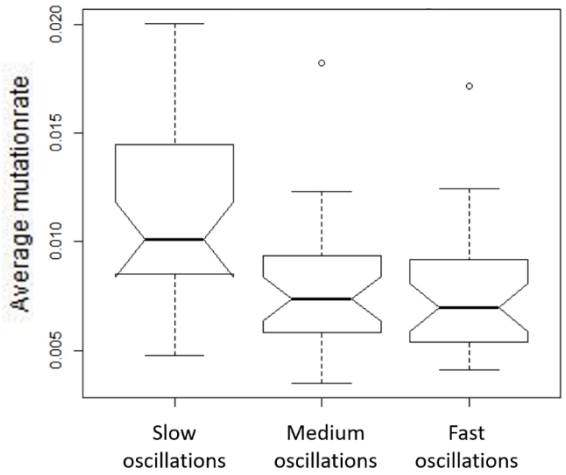

Figure 21: The effect of environmental oscillation speed, statistical analysis. This figure shows the population average of the mutation rate at the end of the simulation, for three different environmental oscillation speeds with 30 replicates each. Slow environmental oscillations (s = 0.01) allow the mutation rate to rise further than medium (s = 0.05) or fast environmental oscillations (s = 0.1). The average mutation rates in the final time step are significantly larger for s = 0.01 compared with s = 0.05 (p = 0.0003) or compared with s = 0.1 (p = 0.0004), there is no significant difference between the final average mutation rate for s = 0.05 and s = 0.1 (p = 0.9932). Parameters other than oscillation speed are kept constant (k = 1, Nb = 1, d = 0, r = 0).

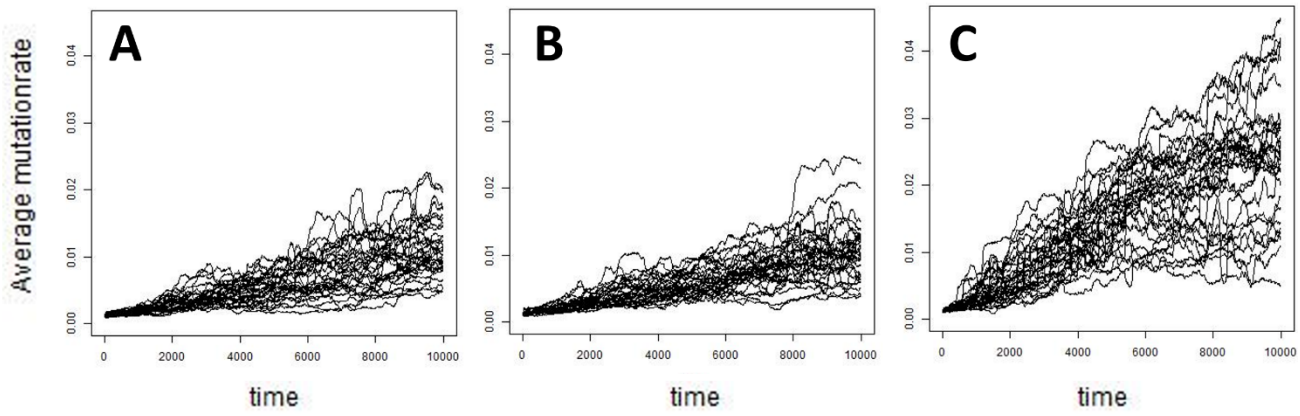

Figure 22: The effect of death rate on mutation rate evolution. This figure shows the increase in the population average of the mutation rate over the course of the simulation, for three different death rates with 30 replicates each. Low death rates (Figure 22A, d = 0) or medium death rates (Figure 22B, d = 0.001) do not allow the mutation rate to rise as far as high death rates (Figure 22C, d = 0.05). Parameters other than the death rate are kept constant (s = 0.01, Nb = 1, k = 1, r = 0).

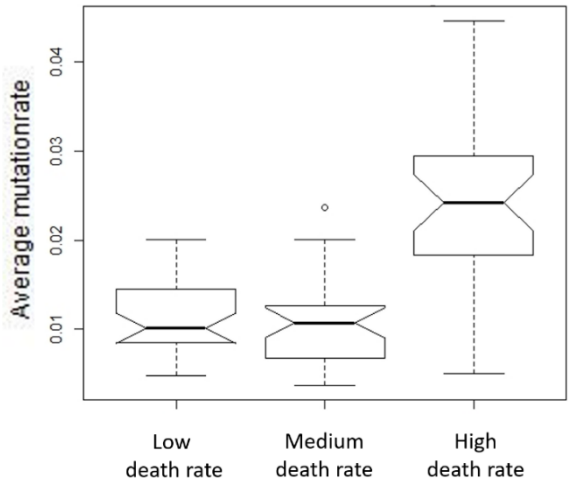

Figure 23: The effect of death rate, statistical analysis. This figure shows the population average of the mutation rate at the end of the simulation, for three different death rates with 30 replicates each. Low death rates (d = 0) or medium death rates (d = 0.001) do not allow the mutation rate to rise as far as high death rates (d = 0.05). The average mutation rates in the final time step are significantly larger for d = 0.05 compared with d = 0 ($p < 0.0001$) or compared with d = 0.001 ($p < 0.0001$), there is no significant difference between the final average mutation rate for d = 0 and d = 0.001 ($p = 0.9106$). Parameters other than death rate are kept constant (s = 0.01, Nb = 1, k = 1, r = 0).

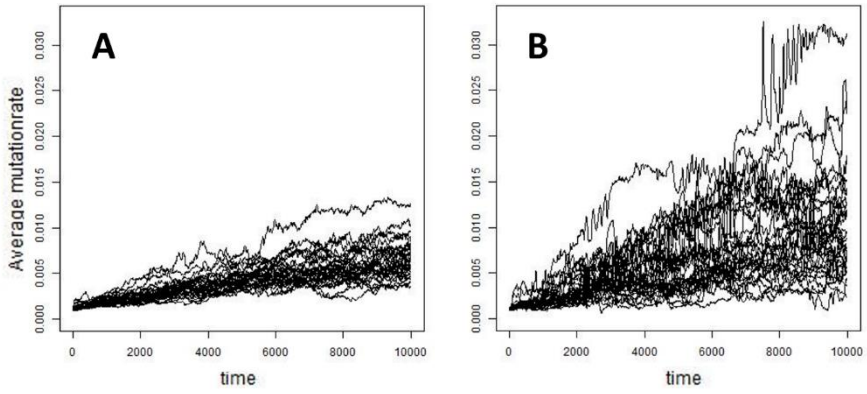

Figure 24: The effect of selection strength on mutation rate evolution. This figure shows the increase in the population average of the mutation rate over the course of the simulation, for two different selection strengths with 30 replicates each. Weak selection (Figure 24A, k = 1) does not allow the mutation rate to rise as far as strong selection (Figure 24B, k = 3). Parameters other than the death rate are kept constant (s = 0.05, Nb = 1, d = 0.001, r = 0.01).

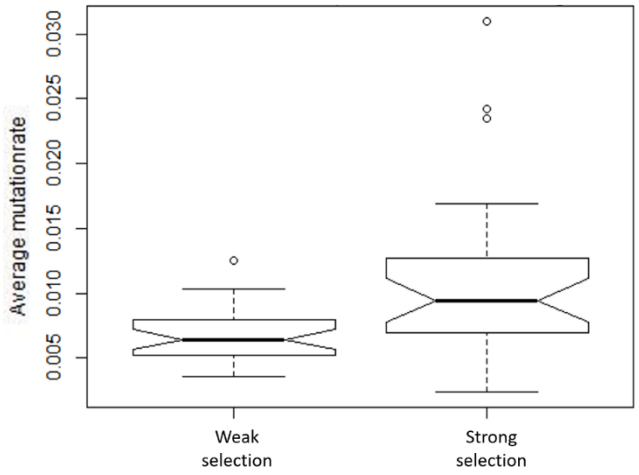

Figure 25: The effect of selection strength, statistical analysis. This figure shows the population average of the mutation rate at the end of the simulation, for two different selection strengths with 30 replicates each. Weak selection (k = 1) does not allow the mutation rate to rise as far strong selection (k = 3) (p = 0. 0013). Parameters other than selection strength are kept constant (s = 0.05, Nb = 1, d = 0.001, r = 0.01).

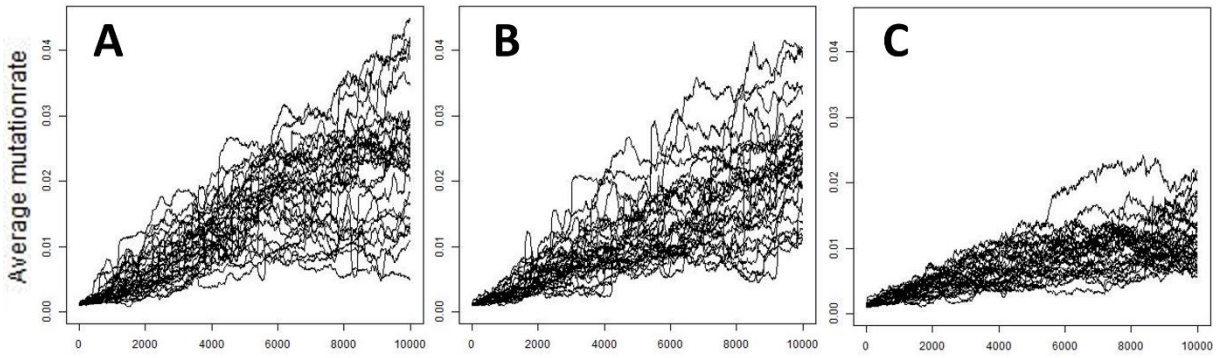

Figure 26: The effect of recombination rate on mutation rate evolution. This figure shows the increase in the population average of the mutation rate over the course of the simulation, for three different recombination rates with 30 replicates each. Low recombination rates (Figure 26A, r = 0) and medium recombination rates (Figure 26B, r = 0.01) allow the mutation rate to rise further than high recombination rates (Figure 26C, r = 0.3). Parameters other than the recombination rate are kept constant (s = 0.01, d = 0.05, Nb = 1, k = 1).

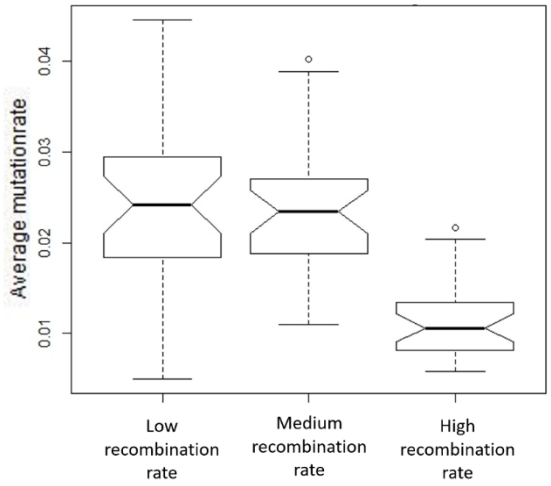

Figure 27: The effect of recombination rate, statistical analysis. This figure shows the population average of the mutation rate at the end of the simulation, for three different recombination rates with 30 replicates each. Low recombination rates (r = 0) or medium recombination rates (r = 0.01) allow the mutation rate to rise further than high recombination rates (r = 0.3). The average mutation rates in the final time step are significantly larger for r = 0 or r = 0.01 compared with r = 0.3 ($p < 0.0001$ in both cases), there is no significant difference between the final average mutation rate for r = 0 and r = 0.01 ($p = 0.7927$). Parameters other than recombination rate are kept constant (s = 0.01, Nb = 1, d = 0.05, k = 1).

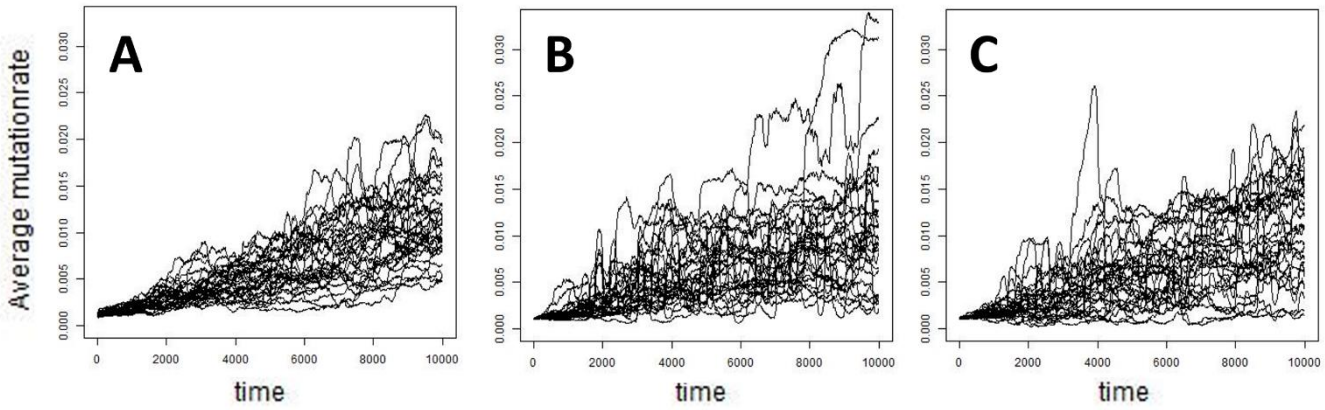

Figure 28: The effect of neighbourhood size on mutation rate evolution. This figure shows the increase in the population average of the mutation rate over the course of the simulation, for three different neighbourhood sizes with 30 replicates each (Figure 28A: Nb = 1, Figure 28B: Nb = 10, Figure 28C: Nb = 50). The increase in mutation rate is not affected by neighbourhood size. Parameters other than the neighbourhood size are kept constant (s = 0.01, d = 0, r = 0, k = 1).

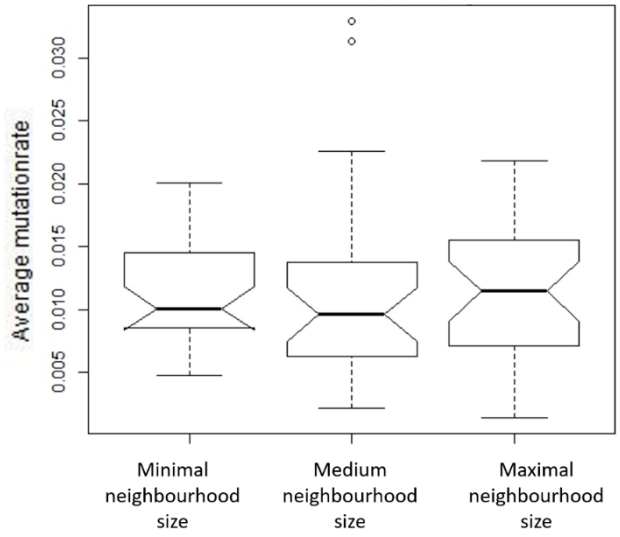

Figure 29: The effect of neigbourhood size, statistical analysis. This figure shows the population average of the mutation rate at the end of the simulation, for three different neigbourhood sizes with 30 replicates each. There is no significant difference between the final average mutation rate for Nb = 1, Nb = 10 and Nb = 50 (p = 0.9606), showing that neighbourhood size does not affect mutation rate evolution. Parameters other than neighbourhood size are kept constant (s = 0.01, d = 0, k = 1, r = 0).

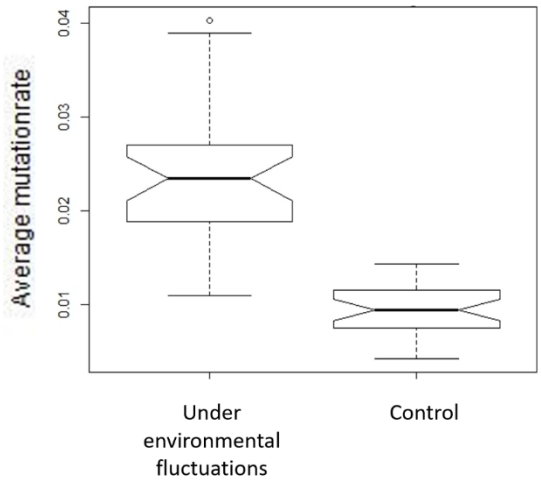

Figure 30: Evolution of high mutation rates even under sexual reproduction and moderate non-zero recombination rates. The mutation rates observed under environmental fluctuations are significantly higher than those observed in the control scenario under constant environmental conditions ($p < 0.0001$), even under sexual reproduction and a non-zero recombination rate of $r = 0.01$. All parameters are kept constant ($s = 0.01$, $d = 0.05$, $k = 1$, $Nb = 1$, $r = 0.01$).

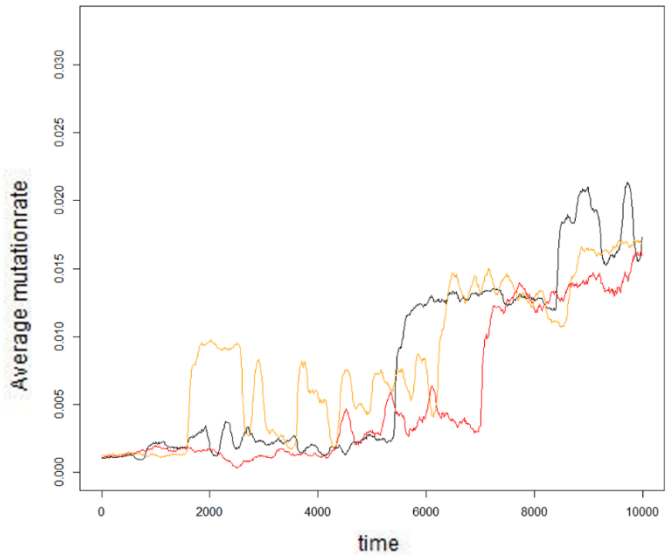

Figure 31: Mutation rate evolution under random environmental fluctuations - example trajectories. This plot shows the trajectories of mutation rate evolution of three independent trajectories, evolving under identical conditions and parameter values ($s = 0.01$, $d = 0$, $k = 1$, $Nb = 50$, $r = 0$). The trajectories differ strongly, and show irregular fluctuations and unpredictable bursts of sudden steep increases or decreases, patterns commonly observed in genes evolving via genetic hitchhiking.

Conclusions Model 3

Overall, Model 3 shows that under fluctuating selection pressures arising from random environmental oscillations, the population average of the mutation rate rises. Increases in mutation rate correlate with lower population sizes. Slower environmental oscillations, higher death rates (allowing a higher population turnover), stronger selection and reduced recombination rates all allow steeper mutation rate growth.

Discussion

In all three models, I found a consistent increase over time in the population average mutation rate under environmental oscillations. A weaker effect was also observed under fluctuating selection pressures caused by negative frequency dependence. This increase in mutation rate is favoured by high death rates, strong selection, slow environmental oscillations and low recombination rates, and is not influenced by spatial processes. Correlation analyses between population sizes and changes in the average mutation rate reveal the link between mutation rate evolution and population dynamics. The results were repeatable across vastly differing scenarios, including sexual and asexual reproduction, single locus and multi-locus systems, spatially structured populations and spatially homogenous populations, and predictable and random environmental fluctuations.

Population dynamics as driving force in mutation rate evolution

I here propose environmentally induced repeated founder effects as a novel mechanism for mutation rate evolution. This mechanism can be best understood by considering first a scenario in which the environmental adaptation is encoded at a single di-allelic locus, as simulated in models 1 and 2. Whenever allele A has higher fitness gains, individuals with allele A are more common than those with allele B, and have a larger absolute number of births: most new individuals will be born to parents with allele A, since they are more abundant and have a higher, fitness-dependent per capita birth rate. Let us consider now that individuals born with allele B can come from two sources: they can be the un-mutated offspring of parents with allele B, or they can be the mutated offspring of parents with allele A. Since individuals with allele A have a higher absolute number of births, most individuals born with allele B are likely to be the mutated offspring of a parent with allele A rather than the un-mutated offspring of a parent with allele B. Parents with higher mutation rates are more likely to have mutated offspring. Therefore, the majority of newly born individuals with allele B are not only likely to have had parents with allele A, but will also inherit their higher-than-average mutation rate. This creates a situation in which there are few, maladapted individuals with allele B which mostly have high mutation rates. Next, due to the implemented oscillations in the environment, allele B becomes the new adaptive allele. The number of individuals with allele A dwindles, and the population is taken over by a growing number of individuals with allele B, all seeded from the small pool of previously maladapted B carriers with high mutation rates. The allele for higher mutation rate, overrepresented amongst carriers of B, spreads together with the now adaptive B allele. Now are back in the starting conditions, as the adaptive allele is more common, and the process repeats. For a schematic visualization of the process, see Figure 32. Continuous environmental oscillations ensure repeated events in which the population is “re-founded” from a small pool of individuals with high mutation rates, leading to the step-wise spread of higher mutation rates seen in model 2. Support for this mechanism comes from the correlation analysis carried out in the second model. As predicted, the majority of individuals born with the rare allele are the mutated offspring of individuals with the common allele (see Figure 12), and the average mutation rate of individuals with a particular allele increases when the allele is rare (see Figure 13).

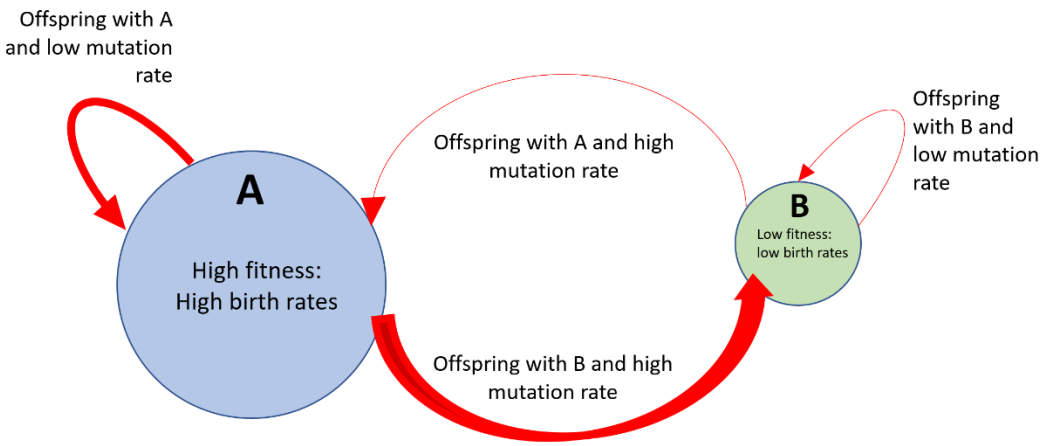

Figure 32: Evolution of increased mutation rates through population dynamics: a schematic diagram. Individuals carrying allele A are displayed in blue, individuals carrying allele B are displayed in green, new births are displayed in red. The size of the circle or arrow indicates the number of individuals. Let us assume that allele A is favoured by selection in the current environment. In this case, there are many carriers of allele A (large blue circle) and few carriers of allele B (small green circle). Most newly born individuals are the offspring of A carriers rather than of B carriers, as A carriers are both more abundant and have, due to their higher fitness, higher birth rates. This is shown graphically by the red arrows departing from A, which are broader than the arrows departing from B. The majority of newly born B-carriers (indicated by the red arrows entering the B circle) are therefore not offspring of B-carriers (thin arrow entering the B circle) but rather mutated offspring of A-carriers (thick arrow going from the A circle to the B circle). Most newly born B-carriers are therefore likely to have a higher mutation rate, since individuals with high mutation rates are more likely to be mutated than individuals with low mutation rates. This way, higher mutation rates accumulate amongst the B-carriers. Once the environment changes, the B allele will be selected for, and the few B-carriers (the green circle) will act as founder population.

This mechanism of repeated founder effects can be extended to multi-locus systems, as explored in model 3. Unlike models 1 and 2, model 3 does not ensure that all mutated individuals are likely to be optimally adapted under the next environmental state, since there are many more than two possible phenotypes. However, one crucial aspect remains the same: after an environmental change, the population is re-founded by a mutated individual which is likely to carry an allele for a high mutation rate in addition to the mutated allele adapting it to the new environment. In other words: the rugged and periodically changing fitness seascape created by varying environmental conditions generates selection for evolvability which leads to an increase in the mutation rate. This is driven by population dynamics: lack of environmental adaptation causes population crashes, after which the population is founded anew from mutated and thus better adapted individuals. Those individuals are likely to have higher mutation rates, and during the population recovery phase, the higher mutation rate hitchhikes along with the environmental adaptation allele. Thus a higher mutation rate evolves through a series of population crashes and founder events. Support for this mechanism comes from the correlation analysis carried out in the third model. As predicted, the change in population average mutation rate is inversely related to population size, i.e. the population average mutation rate increases during population crashes.

To some extent, these models explore an already canonical concept: In all three models, the evolution or spread of the mutation rate proceeds via hitchhiking with beneficial mutations, for a detailed explanation of this mechanism see the introduction. Hitchhiking itself is well-established, and widely regarded as the standard mechanism underlying mutation rate evolution. What is new is the emphasis on population dynamics: the consistent and rapid spread of higher mutation rates arises through serial founder events facilitated purely by population dynamics.

Considering the significance of population crashes in these models, it is worth examining them more closely. Interestingly, the population crashes observed in the single-locus models (models 1 and 2) are starkly different from those observed in the multi-locus model (model 3), see Figure 33. In the multi-locus model, irregular severe reductions of population size occur: when the environment changes and the population is strongly maladapted, then the birthrates are low enough to cause a crash. This does not occur at regular intervals, as some new environments are sufficiently close to the old one to prevent a full crash (Figure 33C). The situation in single-locus models is quite different: here, there are no population crashes in the conservative sense, as there are no strong fluctuations in total population size (Figure 33A and 33B). Nevertheless, there are functional population crashes. When the environment changes, it is merely a small minority which reproduces: since reproductive rates are fitness dependent, the majority of the population is functionally gone –they do not reproduce and are therefore invisible to selection. They merely linger for the rest of their lifespan, and their replacement proceeds too fast to allow a drop in total population size. Yet, if one considers the carries of the allele A and the carries of allele B as two subpopulations, these subpopulations undergo separate crashes, with peaks in population sizes of one subpopulation masking the crash of the other (see Figure 13 or Figure 14). Overall, both types of population crashes, overt ones and masked ones, lead to the same population dynamics of repeated founder events and an associated increase in the mutation rates.

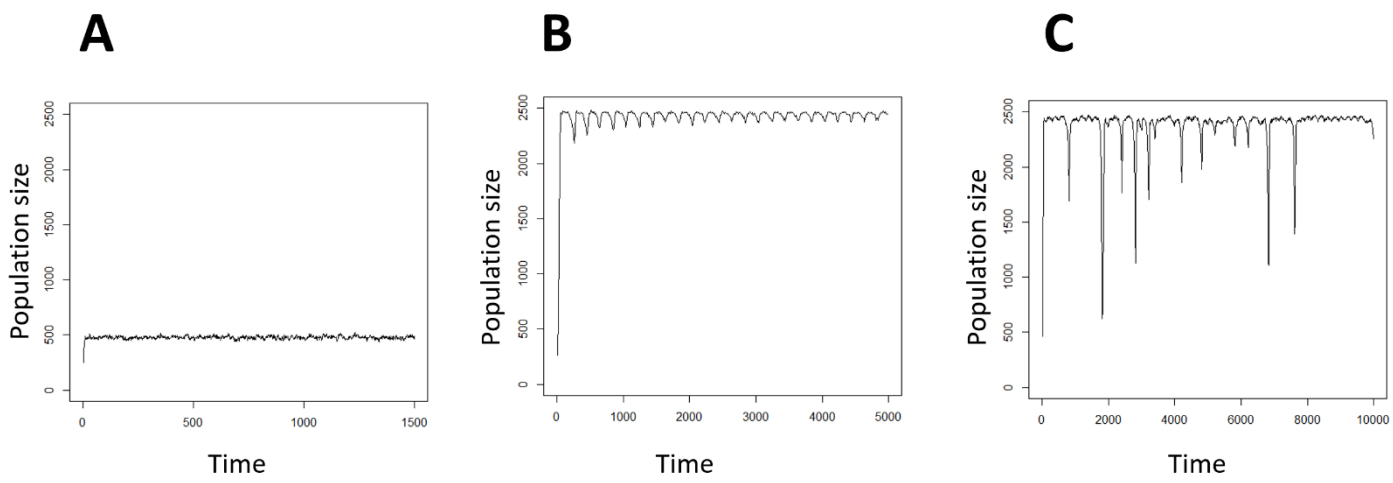

Figure 33: Population dynamics in different models. These plots show examples of the population size changes over time for model 1 (Figure 33 A), model 2 (Figure 33 B) and model 3 (Figure 33 C). Model 3 shows severe population crashes, whereas in model 1 and 2 the population size stays approximately constant. The x-axes of the plots differ due to the differences in simulation run time (1500 time steps in model 1, 5000 time steps in model 2 and 10000 time steps in model 3).

Mutation rate evolution under sexual reproduction

This reliance on population dynamics plays a decisive role in the topic of mutation rate evolution under sexual reproduction. Much of the literature on mutation rate evolution considers the evolution of high adaptive mutation rates under sexual reproduction unlikely (Sniegowski et al. 2000, Baer 2007, Martincorena & Luscombe 2013). Genetic hitchhiking relies, by definition, on linkage between the mutator allele and the beneficial mutation, and recombination during sexual reproduction breaks this linkage (for a more in depth discussion of the literature, see the introduction). However, in my model extensive evolution via genetic hitchhiking can be observed despite sexual reproduction (model 3). I propose that this results from the population dynamics: the mutator allele spreads via a rapid genetic sweep after a population crash. It is a well-documented phenomenon in population genetics that such

population bottlenecks and inbreeding often go hand in hand (see for example Frankham 1998, Triest & Tim 2015) – and inbreeding reduces the effective amount of recombination caused by sexual reproduction, allowing the linkage to be maintained. Thus mutation rate evolution via repeated founder events might circumvent the challenges posed by sexual reproduction. Further to this, a second type of linkage can be observed in those models in which the environmental adaptation is encoded in a single di-allelic locus. As explained above, most instances of mutation are from the fit to the unfit environmental adaptation allele, since carriers of the fit allele are more common. This gives rise to a statistical linkage between the mutator allele and the currently detrimental mutation – however, due to the alternating selection pressure, this mutation will soon be selected for, producing a statistical linkage between the mutator and the now beneficial mutation. This linkage crucially does not arise from a physical connection, and thus cannot be broken by recombination or sexual reproduction. Whilst I did not test the combination of a single-locus di-allelic environmental adaptation gene and sexual reproduction, I predict based on my results that such a statistical linkage may be able to compensate for recombination during sexual reproduction. A similar statistical linkage between mutator allele and beneficial mutation in connection with repeated founder events was already found in a previous study, but arose there through a different mechanism which operates through spatial sorting in an expanding population (Cobben et al. 2017). Overall, these connections between population dynamics and linkage demonstrate that population dynamics do indeed play a crucial role in mutation rate evolution.

The influence of model parameters

The steepness of mutation rate increase, and the mutation rate values obtained at the end of the simulation, are influenced by several different parameters. In fact, all parameters other than neighbourhood size had statistically significant impacts on the mutation rate values at the end of the simulation. Furthermore, the impact of the tested parameters on mutation rate evolution was repeatable across model types, despite the differences between models, increasing the robustness of these results. This does not apply to selection strength and recombination rate, which were tested in the third model only. First I want to consider the effect of recombination rate: high recombination rates break the linkage between mutator and beneficial mutation, and thus prevent the spread of higher mutation rates. This link between recombination and mutation rate evolution is already well-established (see for example Gerrish et al. 2007). However, varying recombination rates in this model demonstrates that the population dynamics discussed above allow high mutation rates to evolve, despite non-zero recombination rates and sexual selection, something that is traditionally thought to be unlikely (Sniegowski et al. 2000, Martincorena & Luscombe 2013).

The next parameter of interest is the speed at which the environment changes. Generally, the mutation rate increases less under faster environmental oscillations. I suggest that under fast oscillations, the population is no longer forced to adapt to the changed environment via mutation, since soon enough a more favourable environment will return: instead of genetic adaptation, the population can simply bide their time unchanged. Environmental oscillations therefore need to be slow enough compared to life span to encourage genetic adaptation. This also explains the interaction between oscillation speed and the population-wide death rate observed in model 2: environmental oscillations occurring at the same speed did not lead to a mutation rate increase in populations with low death rates, and to a strong increase in populations with high death rates. This shows that high death rates prevent the ability of a population to “bide their time unchanged”, environmental oscillation speeds should therefore be considered in relation to the death rate. Indeed, in models with especially high death rates, mutation rates actually increased more under fast oscillations – presumably due to the larger total number of founder effects: doubling the oscillation speed doubles the number of founder events

occurring in the simulation time span. The interaction of lifespan and speed of environmental oscillations is therefore crucial to understanding how a changing environment impacts mutation rate evolution.

In addition to their interaction with oscillation speed, higher population-wide death rates favour higher mutation rates through a second mechanism: increased death rates cause increased population turnover, thus exacerbating the severity of population crashes and allowing high-mutator strains to replace the previous population more efficiently during the founder event following the population crash. This is mirrored by the effect of selection strength: in the third model, increasing selection strength lead to steeper mutation rate increases. As with higher death rates, stronger selection, i.e. a steeper fitness gradient, deepens population crashes and thus helps the mutation rate to spread. Both high death rates and strong selection thus favour those population dynamics that, as discussed above, lie at the very heart of mutation rate evolution in these models.

The role that death rate plays in the evolution of mutation rates is especially interesting considering the large body of literature linking stressful environmental conditions and the evolution of mutation rates. A range of experimental studies found that in both prokaryotes and eukaryotes, stressful environmental conditions are associated with high mutation rates (Hoffman & Hercus 2000, Lamb et al. 2008, Matic 2013, Swings et al. 2017). This result is often explained either as a maladaptive failure of replication faithfulness assurance mechanisms, or as an adaptive response allowing organisms to adapt to a new environment. In addition to those explanations, my study proposes population dynamics as a novel link between stressful environments and mutation rate evolution: in stressful environmental conditions, death rates are higher. This allows population dynamics with frequent and severe population crashes to emerge, which in turn favours the evolution of higher mutation rates. Interestingly, one empirical paper on the link between stressful environmental conditions and mutation rates did indeed find an association between hypermutator strains and increased death rates (Swings et al. 2017). However, the authors suggest that this is due to the accumulation of deleterious mutations caused by the high mutation rate, casting high death rates as a cost of mutation selecting against higher mutation rates. This is in contrast to my models where high death rates, induced environmentally rather than through genetic load, favour higher mutation rates.

Implications for other fields of study

The findings of this study have exciting implications for other fields and topics. For example, it is an established result in population genetics that bottlenecks cause inbreeding and thus a loss of diversity. However, based on my results showing that both bottlenecks and inbreeding allow the evolution of larger mutation rates, one might speculate that after the population bottleneck, elevated mutation rates might generate increased amounts of novel variation, which could compensate for this loss of diversity. On the other hand, increased mutational loads caused by the elevated post-bottleneck mutation rates might compound problems in populations already suffering from inbreeding depression. Considering the link between population dynamics and mutation rate evolution in the context of population genetics of endangered species might thus open up new avenues for conservation related research. Furthermore, this study shows a link between mutation rate evolution and population dynamics, especially population crashes. Therefore, anything that buffers population crashes might reduce evolvability, such as plasticity (Nicotra et al. 2010), life-history strategies (Ferrer et al. 2004), or buffer effects created by environmental conditions like variation in territory quality (Gill et al. 2001). Conversely, anything that might induce population crashes, or at least oscillations in population size, might increase evolvability, including patterns of synchronized dispersal or predation (Huitu et al. 2005), environmental, biophysical or climatic fluctuations (Robinson 2012), and chaotic population dynamics (Tilman & Wedin 1991).

Conclusion

Overall, this study shows the evolution of elevated mutation rates under environmental fluctuations, facilitated by population dynamics displaying repeated population crashes followed by repeated founding events. Parameter combinations that heighten these population dynamics further mutation rate evolution, including slower environmental oscillations, higher population wide death rates, steeper selection gradients and lower recombination rates. Emphasizing the role of population dynamics in mutation rate evolution opens up new avenues for further research, including new facets of how linkage may affect genetic hitchhiking under sexual reproduction, and implications for population genetics and conservation studies of species undergoing population bottlenecks.

References

Atwood, K. C., Schneider, L. K., & Rryan, F. J. (1951). Selective mechanisms in bacteria. *Cold Spring Harbor Symposia on Quantitative Biology* (16), 345-355.

Baer, C. F., Miyamoto, M. M., & Denver, D. R. (2007). Mutation rate variation in multicellular eukaryotes: causes and consequences. *Nature Reviews Genetics*, 8(8), 619-631.

Belshaw, R., Gardner, A., Rambaut, A., & Pybus, O. G. (2008). Pacing a small cage: mutation and RNA viruses. *Trends in ecology & evolution*, 23(4), 188-193.

Carja, O., Liberman, U., & Feldman, M. W. (2014). Evolution in changing environments: Modifiers of mutation, recombination, and migration. *Proceedings of the National Academy of Sciences*, 111(50), 17935-17940.

Cobben, M. M., Mitesser, O., & Kubisch, A. (2017). Evolving mutation rate advances the invasion speed of a sexual species. *BMC evolutionary biology*, 17(1), 150.

Crombach, A., & Hogeweg, P. (2008). Evolution of evolvability in gene regulatory networks. *PLoS Computational Biology*, 4(7), e1000112.

Desai, M. M., Fisher, D. S., & Murray, A. W. (2007). The speed of evolution and maintenance of variation in asexual populations. *Current biology*, 17(5), 385-394.

Ferrer, M., Otalora, F., & GarcÍa-Ruiz, J. M. (2004). Density-dependent age of first reproduction as a buffer affecting persistence of small populations. *Ecological Applications*, 14(2), 616-624.

Foster, P.L. (2007). Stress-induced mutagenesis in bacteria. *Critical Reviews in Biochemistry and Molecular Biology*, 42, 373–397.

Frankham, R. (1998). Inbreeding and extinction: island populations. *Conservation biology*, 12, 665-675.

Furió, V., Moya, A., & Sanjuán, R. (2005). The cost of replication fidelity in an RNA virus. *Proceedings of the National Academy of Sciences*, 102(29), 10233-10237.

Gerrish, P. J., Colato, A., Perelson, A. S., & Sniegowski, P. D. (2007). Complete genetic linkage can subvert natural selection. *Proceedings of the National Academy of Sciences*, 104(15), 6266-6271.

Gill, J. A., Norris, K., Potts, P. M., Gunnarsson, T. G., Atkinson, P. W., & Sutherland, W. J. (2001). The buffer effect and large-scale population regulation in migratory birds. *Nature*, 412(6845), 436.

Hanson, E. D. (1966). Evolution of the cell from primordial living systems. *The Quarterly Review of Biology*, 41(1), 1-12.

Hoffmann, A. A., & Hercus, M. J. (2000). Environmental stress as an evolutionary force. *Bioscience*, 50(3), 217-226.

Houle, D. (1992). Comparing evolvability and variability of quantitative traits. *Genetics*, 130, 195-204.

Huitu, O., Laaksonen, J., Norrdahl, K., & Korpimäki, E. (2005). Spatial synchrony in vole population fluctuations–a field experiment. *Oikos*, 109(3), 583-593.

Ishii, K., Matsuda, H., Iwasa, Y., & Sasaki, A. (1989). Evolutionarily stable mutation rate in a periodically changing environment. *Genetics*, 121(1), 163-174.

Jeffreys, A. J., Royle, N. J., Wilson, V., & Wong, Z. (1988). Spontaneous mutation rates to new length alleles at tandem-repetitive hypervariable loci in human DNA. *Nature*, 332(6161), 278.

Johnson, T. (1999). Beneficial mutations, hitchhiking and the evolution of mutation rates in sexual populations. *Genetics*, 151(4), 1621-1631.

Kauffman, S. A. (1990). Requirements for evolvability in complex systems: orderly dynamics and frozen components. *Physica D: Nonlinear Phenomena*, 42(1-3), 135-152.

Laland, K. N., Uller, T., Feldman, M. W., Sterelny, K., Müller, G. B., Moczek, A., ... & Odling-Smee, J. (2015). The extended evolutionary synthesis: its structure, assumptions and predictions. *Proceedings of the Royal Society B: Biological Sciences*, 282(1813), 20151019.

Lamb, B. C., Mandaokar, S., Bahsoun, B., Grishkan, I., & Nevo, E. (2008). Differences in spontaneous mutation frequencies as a function of environmental stress in soil fungi at “Evolution Canyon,” Israel. *Proceedings of the National Academy of Sciences*, 105(15), 5792-5796.

MacLean, R. C., Torres-Barceló, C., & Moxon, R. (2013). Evaluating evolutionary models of stress-induced mutagenesis in bacteria. *Nature Reviews Genetics*, 14(3), 221.

Maharjan, R. P., Liu, B., Li, Y., Reeves, P. R., Wang, L., & Ferenci, T. (2013). Mutation accumulation and fitness in mutator subpopulations of Escherichia coli. *Biology letters*, 9(1), 20120961.

Martincorena, I., & Luscombe, N. M. (2013). Non-random mutation: the evolution of targeted hypermutation and hypomutation. *BioEssays*, 35(2), 123-130.

Matic, I. (2013). Stress-induced mutagenesis in bacteria. *Stress-induced mutagenesis* (pp. 1-19). Springer, New York, NY.

Mujalli, D. R. (2019). Ecological Modeling of Adaptive Evolutionary Responses to Rapid Climate Change. PhD thesis, University of Potsdam, Germany.

Nicotra, A. B., Atkin, O. K., Bonser, S. P., Davidson, A. M., Finnegan, E. J., Mathesius, U., ... & van Kleunen, M. (2010). Plant phenotypic plasticity in a changing climate. *Trends in plant science*, 15(12), 684-692.

Nuño de la Rosa, L. (2017). Computing the extended synthesis: Mapping the dynamics and conceptual structure of the evolvability research front. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*, 328(5), 395-411.

Parsons, K. J., Son, Y. H., Crespel, A., Thambithurai, D., Killen, S., Harris, M. P., & Albertson, R. C. (2018). Conserved but flexible modularity in the zebrafish skull: implications for craniofacial evolvability. *Proceedings of the Royal Society B: Biological Sciences*, 285(1877), 20172671.

Petrie, M., & Roberts, G. (2007). Sexual selection and the evolution of evolvability. *Heredity*, 98, 198.

Regoes, R. R., Hamblin, S., & Tanaka, M. M. (2013). Viral mutation rates: modelling the roles of within-host viral dynamics and the trade-off between replication fidelity and speed. *Proceedings of the Royal Society B: Biological Sciences*, 280(1750), 20122047.

Pigliucci, M. (2010). Genotype–phenotype mapping and the end of the 'genes as blueprint' metaphor. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1540), 557-566.

Robinson, K. L. (2012). Climate drives local to global variations of coastal gelatinous zooplankton. PhD thesis, University of South Alabama.

Rowe, David, John Leaney, and David Lowe. "Defining systems evolvability-a taxonomy of change." *Change* 94 (1994): 541-545.

Rutherford, S. L. (2003). Between genotype and phenotype: protein chaperones and evolvability. *Nature Reviews Genetics*, 4(4), 263.

Schwabauer, R. (1976). Enzymatic neurons. *Journal of theoretical biology*, 59(1), 223-230.

Shaver, A. C., Dombrowski, P. G., Sweeney, J. Y., Treis, T., Zappala, R. M., & Sniegowski, P. D. (2002). Fitness evolution and the rise of mutator alleles in experimental Escherichia coli populations. *Genetics*, 162(2), 557-566.

Sniegowski, P. D., Gerrish, P. J., Johnson, T., & Shaver, A. (2000). The evolution of mutation rates: separating causes from consequences. *Bioessays*, 22(12), 1057-1066.

Sprouffske, K., Aguilar-Rodríguez, J., Sniegowski, P., & Wagner, A. (2018). High mutation rates limit evolutionary adaptation in Escherichia coli. *PLoS genetics*, 14(4), e1007324.

Stojanova, B., Koláříková, V., Šurinová, M., Klápště, J., Hadincová, V., & Münzbergová, Z. (2019). Evolutionary potential of a widespread clonal grass under changing climate. *Journal of evolutionary biology*, 32(10), 1057-1068.

Swartout, W. R., & Neches, R. (1986). The Shifting Terminological Space: An Impediment to Evolvability. *Association for the Advancement of Artificial Intelligence,* 936-941.

Swings, T., Van den Bergh, B., Wuyts, S., Oeyen, E., Voordeckers, K., Verstrepen, K. J., ... & Michiels, J. (2017). Adaptive tuning of mutation rates allows fast response to lethal stress in Escherichia coli. *Elife*, 6, e22939.

Tenaillon, O., Barrick, J. E., Ribeck, N., Deatherage, D. E., Blanchard, J. L., Dasgupta, A., ... & Schneider, D. (2016). Tempo and mode of genome evolution in a 50,000-generation experiment. *Nature*, 536(7615), 165.

Tilman, D., & Wedin, D. (1991). Oscillations and chaos in the dynamics of a perennial grass. *Nature*, 353(6345), 653.

Triest, L., & Tim, S. (2015). Strong bottlenecks, inbreeding and multiple hybridization of threatened European Ruppia maritima populations. *Aquatic botany*, 125, 31-43.

Van Gestel, J., & Weissing, F. J. (2016). Regulatory mechanisms link phenotypic plasticity to evolvability. *Scientific reports*, 6, 24524.

Verd, B., Monk, N. A., & Jaeger, J. (2019). Modularity, criticality, and evolvability of a developmental gene regulatory network. *eLife*, 8, e42832.

Wagner, G. P., & Altenberg, L. (1996). Perspective: complex adaptations and the evolution of evolvability. *Evolution*, 50(3), 967-976.

Wagner, G. P., & Zhang, J. (2011). The pleiotropic structure of the genotype–phenotype map: the evolvability of complex organisms. *Nature Reviews Genetics*, 12(3), 204.

Wright, R. M., Mera, H., Kenkel, C. D., Nayfa, M., Bay, L. K., & Matz, M. V. (2019). Positive genetic associations among fitness traits support evolvability of a reef-building coral under multiple stressors. *bioRxiv*, 572321.

Evolution of evolvability

Part II: Selection for enhanced mortality?

Master thesis of Jana M. Riederer

s3507130

February 2019 – December 2019

(50 ECTS)

Supervised by prof. dr. F.J. Weissing & T.J.B van Eldijk

MSc Ecology & Evolution – Track Evolutionary Biology

Theoretical Research in Evolutionary Life Sciences

Rijksuniversiteit Groningen

rijksuniversiteit groningen

MARM GROUP

Evolution of evolvability II: Selection for enhanced mortality?

Introduction

Delaying death, or even attaining immortality, has attracted the fascination of scientists and artists alike throughout human history. Researching the evolutionary origins of ageing and the mechanisms determining lifespan is also of great medical relevance: a greater understanding of ageing might help us to potentially reverse some of the effects, giving us a new approach for age-related deteriorations in health. These implications for the medical field and human well-being has led to the science of ageing receiving increasing attention (Goldsmith 2008, Lopez-Otin et al. 2013, Grey 2015, Rando & Chang 2012).

Ageing can be defined as an age-related increase in mortality. I will here consider literature not only on ageing, but on mortality, lifespan and death rates as well, since despite subtle differences they are essentially closely related ways of approaching the same concept. I will also focus on intrinsic mortality, distinct from extrinsic mortality, i.e. distinct from mortality due to outside influences such as predation or accidents. From an evolutionary point of view, intrinsic mortality poses a puzzling question: considering that immortal beings could, theoretically, live and reproduce forever, why do we see intrinsic mortality at all? The answers to this question are manifold, but can be split into two opposing sides: explanations considering intrinsic death maladaptive (also referred to as non-programmed death), and explanations considering intrinsic death as adaptive (also referred to as programmed death).

The former explanations are well established in the literature on death rates and ageing, and are commonly categorized into three classical theories of ageing: the mutation accumulation theory, the antagonistic pleiotropy theory and the disposable soma theory. The first of these, the mutation accumulation theory, suggests that mutations acting late in life cannot be removed efficiently by selection and thus accumulate: most individuals die by a certain age due to external factors, and mutations acting after this age are shielded from selection by virtue of not being expressed sufficiently often (Medawar 1952, Medawar 1946). In other words, late-acting mutations lie in the selection shadow of extrinsic mortality. The second of the classical theories suggests an antagonistic pleiotropy between traits which are beneficial at early life stages but disadvantageous at later stages. Such genes might still be favoured by evolution, as early life stages are critical to overall fitness, even if it comes a survival cost in later life stages (Williams 1957). Finally, the disposable soma theory proposes a similar trade-off, with the role of the early life benefit played by an increase in reproductive output. The idea here is that resources can be invested either in reproduction or repair, and any increase in fertility is paid for with a reduction in lifespan (Kirkwood 1997). From a gene's point of view, so to speak, the germline is crucial for evolutionary success, whereas the survival of the individual, i.e. the soma, is secondary – reproduction may be valued over repair, leading to the term "disposable soma". The common theme throughout these three theories is that any intrinsic death is necessarily maladaptive and would be selected against, were selection not bound in its efficiency by selection shadows or trade-offs.

However, there is also some indication that intrinsic mortality rates need not be determined by the maladaptive but inevitable processes assumed in the classical ageing theories. First evidence comes from the various species which do not show the classical intrinsic death rates or age induced increases in mortality observed in most other animal species, such as famously the "immortal jellyfish" *Turritopsis dohrnii*, as well as various species of rockfish, sturgeon, or bivalves (Finch & Austad 2001).

Indeed, some species even show negative senescence: an age-related decrease in mortality (Vaupel et al. 2004, Jones et al. 2014). These species thus offer support to the idea that intrinsic immortality might not be as fundamentally impossible as often assumed, and ageing might not be an universal unavoidable constraint but rather a unique feature of a limited number of species. Furthermore, even species which at first glance seem to fit the model of inevitably and maladaptively short lifespans, such as mice, rats or fruit flies, can experience greatly extended lifespans through caloric restrictions (Masoro 2003) or through specific mutations (Bartke 2001, Lithgow & Gill 2003) – demonstrating that observed constraints on lifespan are not necessarily unavoidable, and therefore require an explanation beyond the ones given in the classical theories of ageing. Some studies also tested the underlying assumptions of the classical theories of ageing, and found them not to be universal: for example, both intra- and interspecies comparisons show that there is no general inverse relationship between fertility and lifespan, confirming that the repair-reproduction trade-off is not compelling (Ricklefs & Cadena 2007, Mitteldorf 2010, Tarín et al. 2014). Finally, some indication for pathways by which such adaptive, programmed lifespan restrictions might be realised come from studies into cellular senescence (Rhinn et al. 2019), conserved mechanisms of ageing in protists (Clark 2004) as well as from research on epigenetic clocks (Lowe et al. 2016).

These studies, attesting to the discrepancies between purely maladaptive death theories and empirical data, give rise to an intriguing question: if intrinsic mortality is indeed not as unavoidable as postulated in classical ageing theories, then why do we observe any intrinsic mortality or programmed lifespan limitations in the first place? In answer to this question various adaptive ageing hypotheses have been proposed, which can be classified broadly into three categories: evolution through group selection, through kin selection, or through selection for evolvability (see amongst others Goldsmith 2014). Group selection arguments suggest a role for ageing in preventing resource depletion and group extinction and ultimately in stabilizing population dynamics (Mitteldorf 2006, Werfel et al. 2015). Kin selection arguments were first brought forward by Wallace in the 1860s, suggesting that intrinsic mortality may prevent parents from competing with their offspring (Wallace 1889). This and similar ideas were tested in various theoretical papers (see for example Travis 2004), however under strongly simplifying assumptions such as young offspring being intrinsically more valuable than older parents, invoking, for example, age-related declines in reproductive fitness. This is problematic as fundamental aspects of ageing (i.e. age-related declines in fitness) must be presumed in order to explain the evolution of ageing. Finally, there have been to date several attempts to explain adaptive ageing using the concept of evolvability. In the 1880s, before the concept of evolvability was first established, Weisman first proposed that aging might allow a faster spread of new alleles, and be favoured by selection on this basis (Weisman 1889). Evolvability has since frequently been mentioned or discussed in various reviews on death rate evolution (for example Goldsmith 2014), but surprisingly only few, partially not peer-reviewed papers aimed to investigate this link between evolvability and death rate (Libertini 1983, Libertini et al. 2017). These papers addressed important parts of the question about how evolvability and death rates are linked, such as the extent to which favourable alleles spread faster under higher death rates. However, to the best of my knowledge, no paper so far has explicitly tested whether selection for evolvability can cause elevated death rates.

I here want to investigate the hypothesis that the joint action of selection for evolvability and kin selection lead to the evolution of elevated death rates in kin structured and resource limited environments. The underlying mechanism that I propose is the following: In a resource limited and kin structured population, related individuals share and compete over resources. The death of one individuals thus frees up resources for its relatives, increasing the indirect fitness of the dead individual – kin selection hence alleviates the fitness costs of dying. On its own, this would not be enough to give rise to elevated death rates, as replacing oneself with a non-identical relative generally decreases

fitness. However, in a variable environment selecting for evolvability such non-identical kin are of great value: each extra relative represents another opportunity to find a novel solution to the constantly changing ecological challenges. These kin, in turn, will generate new variation by reproducing and freeing up resources for the future generations through their own early death. In other words: early-death lines experience greater population turnover, allowing them to cycle through more genetically different sets of offspring than longer-lived lines. They thus experience a greater deal of genetic variation and are more likely to chance across novel solutions to ecological challenges: higher death rates increase their ability to keep up with the changing selection pressures of a changing environment, ultimately increasing evolvability. From a gene's point of view, a novel allele for ecological success in the current environment is more likely to arise in a genome containing an allele for earlier death than in a genome containing an allele for delayed death, and the earlier death allele can then hitchhike along with the selective sweep of the novel ecological adaptation allele. To sum up: I here test the hypothesis that in a kin-structured population in a resource-limited environment, fluctuating environmental conditions select for increased evolvability which is achieved through a decreased lifespan.

Model description

Model structure

T The fourth model, exploring death rate evolution, is largely identical to the third model, showing mutation rate evolution under sexual selection. Thus, this model too features sexual reproduction, multi-locus genetics, non-linear genotype-phenotype mapping and random environmental fluctuations – with one difference: the mutation rate is set at a constant value, and instead the death rate (specifically, the death rate during the second round of death events) is allowed to evolve.

Parameter values

All fixed parameter settings from model 3 were kept unchanged, including the simulation run time of 10000 timesteps. As in model 3, I studied how model behaviour is influenced by the neighbourhood size Nb, the speed of environmental oscillations s, the selection strength k, and the fixed death rate during the second round of death events d. Separating the evolvable death rate and the fixed death rate into two separate rounds of death events allows me to implement different levels of maximal lifespans, without interfering with the evolution of the evolvable death rate in the first rounds of death events. This way, the evolution of the evolvable death rate remains comparable between simulations with different parameter settings. Unless specified, "death rate" refers to the evolvable death rate in the rest of the report, and not to the population-wide fixed death rate. Varying neighbourhood size allows me to manipulate the strength of kin selection: smaller neighbourhoods result in increased kin structure, and thus more powerful kin selection. At maximal neighbourhood size Nb = 50 the population is homogeneously mixed and has no spatial structure, kin selection is hence absent. Overall, I tested 180 different combinations of parameter values.

Furthermore, I explored different variations of this model – with and without selfing, with bitstrings with 10 or 20 positions as ecological adaptation loci, with different mutation mechanisms, and with and without a trade-off between reproduction and life span. Since I could not identify a qualitative difference between those scenarios, I selected one scenario (no selfing, 20 positions, capped death rates, no trade-off between reproduction and life span) for a more in-depth quantitative analysis.

Control simulations

The results from this model were compared with those from two different controls: firstly, a simulation without selection for evolvability, and secondly a simulation without kin selection. In the first control, the oscillating fitness seascape is replaced with a flat fitness landscape, and the environment is thus stable. The fitness pay-off is kept at 0.5 at all times, an average of the fitness values given by a sine curve between zero and one. In the second control, a homogeneously mixed population is achieved by setting the neighbourhood size to the maximum value of 50. This absence of kin structure does not allow for kin selection to act.

Differences in death rate between the model and the first control show the effect of an oscillating environment selecting for evolvability. In the case of selection for evolvability overpowering selection for immortality, leading to increasing death rates, this is a decisive control. However, in the case of decreasing death rates, the situation is not entirely clear yet: a less steep decline in death rates could be due to selection for higher death rates caused by selection for evolvability – or it could be due to repeated founder events in oscillating environments increasing the strength of drift and hence reducing the efficiency of selection for immortality. The second control is then used to separate the action selection for evolvability from the action of drift. As kin selection is not acting in the second control, selection for evolvability cannot act on death rates. Any differences between simulations evolving under stable and under oscillating environments is then merely due to the differences in the strength of selection for immortality. To summarise, the hypothesis is accepted if there is a difference in model behaviour between the model and control 1, and if this difference is not evident in the comparison of control 2 and control 1.

Data analysis

The most obvious data analysis – comparing the evolved values of the death rate at the end of the simulation – was not feasible in this model. Overall, all simulations lead to the evolution of very low death rates. Once death rates have dropped below a certain value, they are unlikely to respond to any selection pressure: elongating lifespans follows the law of diminishing returns (Olshansky 2018), weakening the selection for longer lifespans. Furthermore, once the duration of the lifespan approaches the remaining duration of the simulation, the lifespan phenotype is not expressed sufficiently often to be acted upon by selection. In fact, in most simulations, the death rates rapidly achieve a value that translates to a lifespan close to the remaining simulation run time. The remaining analysis of the simulation results therefore focuses not on the death rate obtained at the end of the simulation run time, since these might be subject more to drift than selection and thus uninformative, but rather on a two-pronged approach: the slope of the death rate trajectory, and the change in death rate between timesteps.

The slopes of the death rate trajectory are calculated in the following way: I used the lm() function in R to fit a linear regression model to the data from each replicate separately. The slope of this model can be considered to be the slope of the death rate trajectory, resulting in 30 trajectory slopes for the 30 replicates from each parameter combination. For this analysis, the data on death rates below 0.01 were excluded, allowing me to avoid regions of low selection strength and consider only the linear section of the death rate trajectory. I confirmed the fit of the linear model slope to the corresponding trajectory through visual inspection for a sample of the data. Certain parameter combinations caused the populations to go extinct, those replicates were excluded from the analysis, as such averages are liable to be misleading. In cases where more than 5 replicates went extinct, the entire parameter combination was excluded from the analysis. The thus obtained data on the death rate trajectory slopes was analysed using the Wilcoxon test for pair-wise comparisons and ANOVA with post-hoc Tukey tests for comparisons of more than two parameter combinations.

The second approach to this project focuses on the change in death rate between timesteps, relating it to the population dynamics. The question here is how death rate changes during population crashes: does it increase (Figure 1 A) or decrease (Figure 1 B)? For this analysis I combined the data of all 30 replicates for each parameter combination tested, but excluded data from the first 500 time steps. This ensures that only established population dynamics are included, i.e. it prevents artefacts caused by the initial population size which is small not due to the underlying population dynamics but due to the initialization process. For the remaining simulation length, I calculated every ten time points: the difference in death rate $\Delta d = {d(t+10) - d(t)}/{10}$, the absolute population size N(t), and the difference in population size $\Delta N = {N(t+10) - N(t)}/{10}$. Next, the data is divided according to the population dynamics, using two different classification methods. In the first method, I separate stable phases and crash phases according to the population size, using a threshold $x = p * mN$, where mN is the median population size and p is either 0.8 or 0.9. If the number of individuals is above the threshold, the population is stable (N(t) > x), elsewise it is undergoing a population crash (N(t) < x). Varying the value of p did not affect the results. In the second method, I separated stable phases, phases of population decline and phases of population recovery according to the changes in population size rather than according to the absolute population size: stable phases are defined as $-10 < \Delta N < 10$, phases of population decline as $\Delta N < -10$, and phases of population recovery as $\Delta N > 10$. I performed a t-test on the changes in death rate Δd from each of the two (classification by population size) or three (classification by change in population size) population dynamics categories, to test whether the values of Δd are significantly above or below zero. To counteract possible oversampling errors, I drew 1000 subsamples for each population dynamics category, performing the t-test for each subsample. I repeated the subsampling analysis for different subsample sizes N_s = 100, N_s = 200, or N_s = 500. Furthermore, the entire analysis of the change in death rate between timesteps was repeated for different parameter values: Nb = 3, Nb = 10, Nb = 50, and the control simulation with a flat fitness landscape.

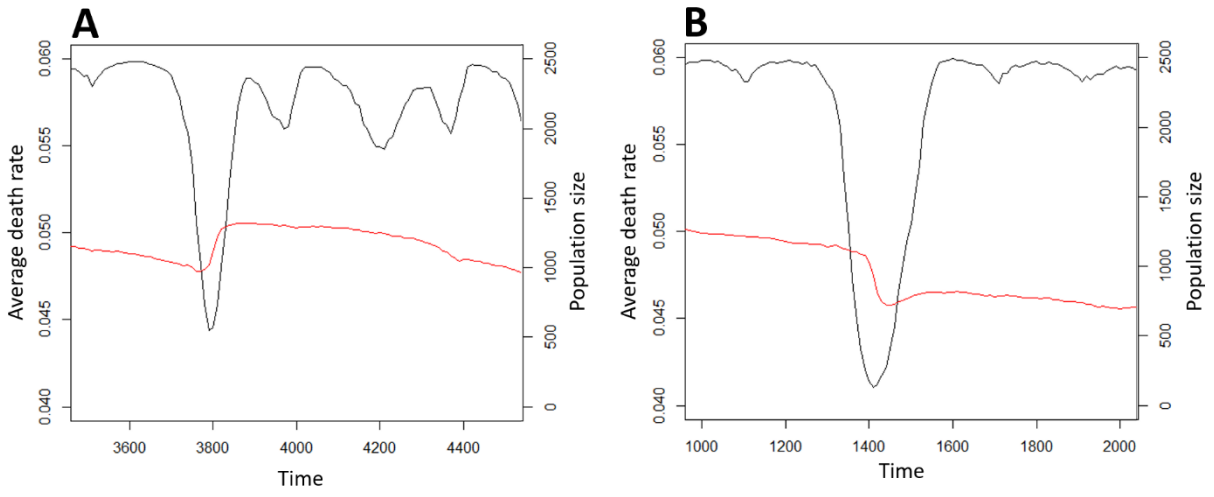

Figure 1: Example plots showing death rate trajectories (red) during population crashes (population size in black). The second approach to data analysis considers how the population average of the death rate changes during a population crash: does it increase (Figure 1 A) or does it decrease (Figure 1 B)? In other words, across all thirty replicates for one parameter combination, do cases such as Figure 1 A or cases such as Figure 1 B predominate?

Simulation results

Evolution of immortality

Across all replicates of all parameter combinations, two main simulation outcomes can be observed: extinction (see Figure 2 A) or immortality (see Figure 2 B). In the majority of parameter combinations, the population average of the death rate declines sharply, allowing the population to reach literal immortality (death rate = 0) or functional immortality (extremely low death rates). Once death rates have dropped below a certain value, they are unlikely to respond to any selection pressure. The remaining analysis of the simulation results therefore focuses not on the death rate obtained at the end of the simulation run time, since these might be subject more to drift than selection, but rather on the slope of the death rate trajectory.

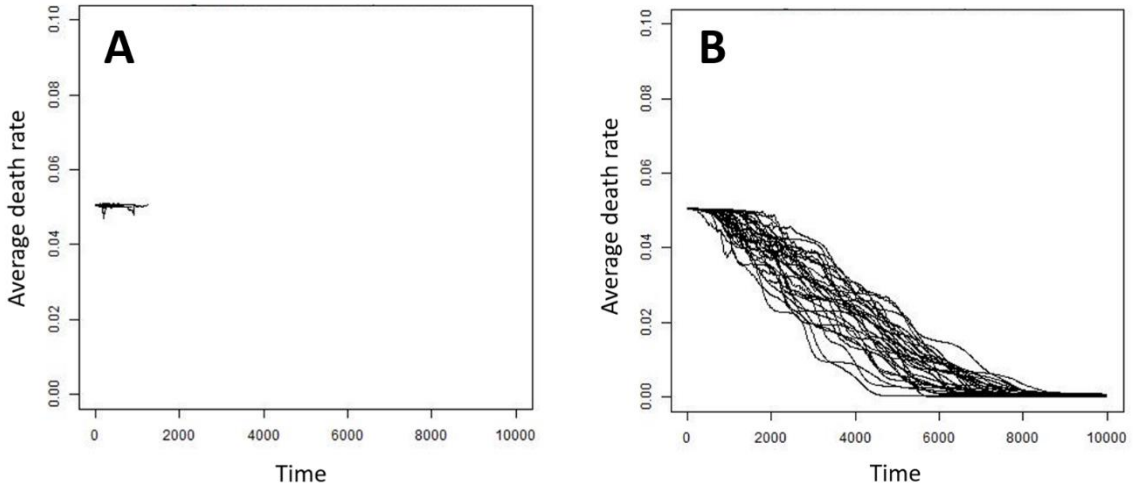

Figure 2: Typical simulation outcomes. Under all tested parameter combinations, the population meets one of two evolutionary fates. Either all populations in all replicates go extinct early in the simulation (Figure 2 A), or the populations evolve immortality by the end of the simulation (Figure 2 B).

Testing selection for evolvability

The hypothesized evolution of adaptive high death rates necessitates two prerequisites. The first one is selection for increased evolvability – otherwise, the death rate would be exposed to selection for elongated lifespans only. The results of model 3 (which differs merely in allowing the mutation rate to evolve, rather than the death rate) show that under environmental oscillations, increased mutation rates evolve, indicating that this model does indeed select for increased evolvability.

Testing kin selection

The second prerequisite is the presence of kin structured populations and kin selection. The death rate shows a shallow decline in small neighbourhood sizes with strong kin structure, and a steep decline in large neighbourhood sizes with weak kin structure (Figure 3). This indicates the presence of kin selection in strongly kin-structured populations, counteracting and weakening selection for immortality: the cost of dying earlier is mitigated if it benefits related individuals by increasing their resource availabilities. This pattern is clearly evident under population-wide low fixed death rates, but less clear under population-wide high fixed death rates and oscillating environments, likely due to death rates and population crashes facilitated by the changing environment increasing population turnover and hence diluting the population kin structure (Figure 4). Overall, this shows that kin selection does indeed act in this system, and is stronger in simulations with smaller neighbourhood sizes.

Figure 3: Comparison of death rate evolution under different levels of kin selection, under low fixed death rates. This plot shows the slope of the evolutionary trajectory of the death rate, for simulations with different neighbourhood sizes. In simulations with small neighbourhood sizes (Nb = 1) the population has a pronounced kin structure, leading to strong kin selection. In simulations with large neighbourhood sizes (Nb = 10) the population is largely homogeneous, leading to weak kin selection. Whilst all trajectory slopes are negative, those from simulations with strong kin selection show a much more shallow decline: strong kin selection thus allows higher death rates to be maintained. Parameters other than the neighbourhood size are kept constant (s = 0.1, d = 0.001, k = 1).

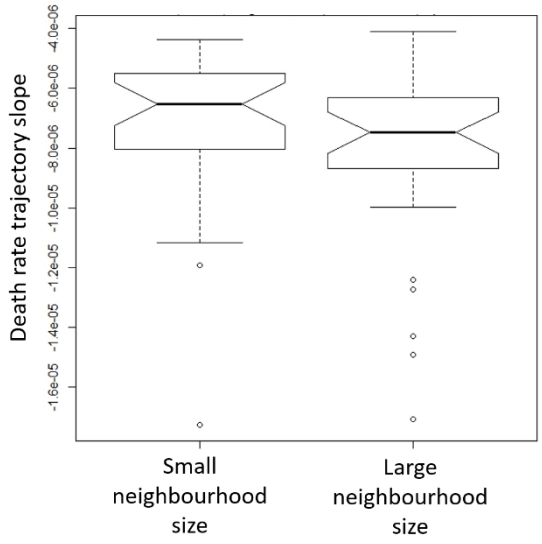

Figure 4: Comparison of death rate evolution under different levels of kin selection, under high fixed death rates. This plot shows the slope of the evolutionary trajectory of the death rate, for simulations with different neighbourhood sizes. Whilst slopes of death rate trajectories evolved in small neighbourhood sizes (Nb = 1) are less steep than those evolved in large neighbourhood sizes (Nb = 10), the pattern is much less pronounced than the one seen in Figure 3 (low fixed death rate). A high fixed death rate thus weakens the impact of kin selection. Parameters other than the neighbourhood size are kept constant (s = 0.1, d = 0.05, k = 1).

The impact of oscillating environments on the slope of death rate trajectories

In kin structured populations (i.e. Nb ≠ 50), the death rate declines significantly slower under environmental oscillations compared with under stable environments. This difference in death rate slope is influenced by several different parameters: slower environmental oscillations (Figure 5), higher population-wide fixed death rates (Figure 7), and stronger selection (Figure 9) all allow the death rate to decline slower. In populations without kin structure (i.e. Nb = 50), the same patterns can be observed: a significantly slower decline of the death rate under environmental oscillations, enhanced by slower environmental oscillations (Figure 6), higher population-wide fixed death rates (Figure 8), and stronger selection (Figure 10).

Under several parameter combinations, the populations of most or all replicates went extinct. Nearly all populations both with and without kin-structure went extinct when evolving under a combination of high population-wide fixed death rates (d = 0.05 or d = 0.075) and strong selection (k = 3). No extinctions happened in the control simulations with a stable environment, regardless of parameter combination.

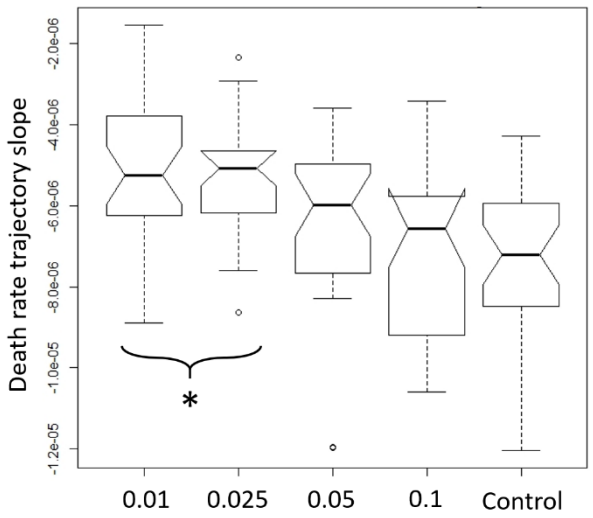

Figure 5: The effect of environmental oscillation speed on death rate evolution, in kin structured populations. The figure shows the slope of the death rate trajectory for 4 different environmental oscillation speeds as well as for the control in which the environment remains stable, with 30 replicates for each. All death rate trajectories show a decline, however, under slow environmental oscillations the decline is less steep. Parameter values for which the slopes of the death rate trajectories are significantly different from those in the control simulations are marked with an asterisk. Parameters other than the speed of environmental oscillations are kept constant (d = 0.075, Nb = 1, k = 1).

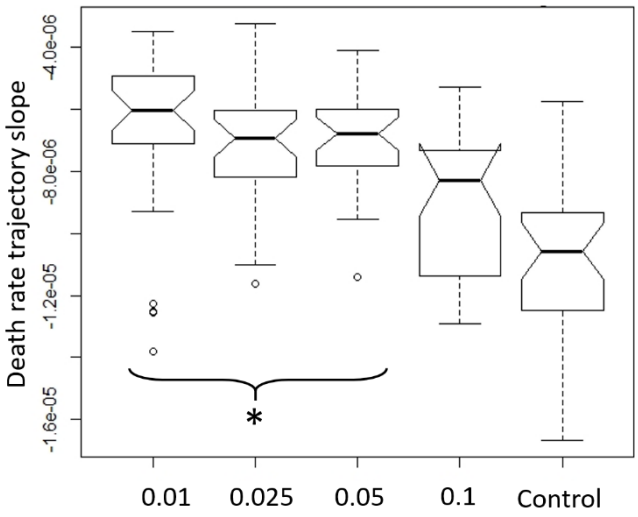

Figure 6: The effect of environmental oscillation speed on death rate evolution, in populations without kin structure. The figure shows the slope of the death rate trajectory for 4 different environmental oscillation speeds as well as for the control in which the environment remains stable, with 30 replicates for each. The results are comparable to those from Figure 5 with kin structure: All death rate trajectories show a decline, under slow environmental oscillations the decline is less steep. Parameter values for which the slopes of the death rate trajectories are significantly different from those in the control simulations are marked with an asterisk. Parameters other than the speed of environmental oscillations are kept constant (d = 0.075, Nb = 50, k = 1).

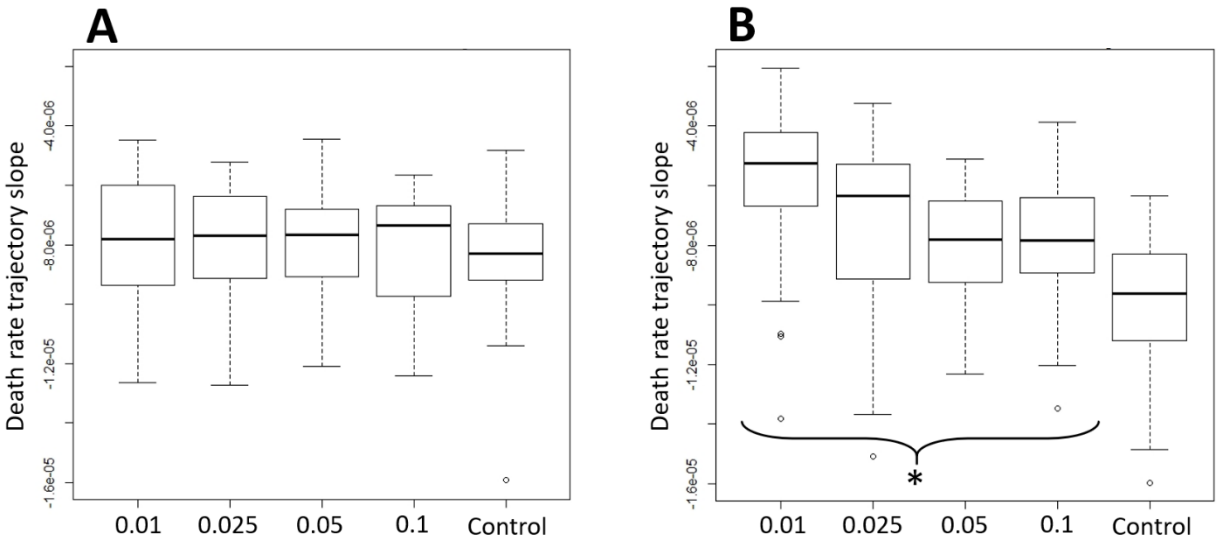

Figure 7: The effect of the fixed population-wide death rate on death rate evolution, in kin structured populations. The figure shows the slope of the trajectory of the individually evolvable death rate for 4 different environmental oscillation speeds as well as for the control in which the environment remains stable, with 30 replicates for each. Figure 7 A: under low fixed death rates (d = 0), the slopes of the death rate trajectory for any environmental oscillation speed are not different from those in the control. Figure 7 B: under high fixed death rates (d = 0.075), the slopes of the death rate trajectory are considerably less steep under environmental oscillations than in the control, especially under slow environmental oscillations. Parameter values for which the slopes of the death rate trajectories are significantly different from those in the control simulations are marked with an asterisk. Parameters other than the fixed death rate and the speed of environmental oscillations are kept constant (Nb = 3, k = 1).

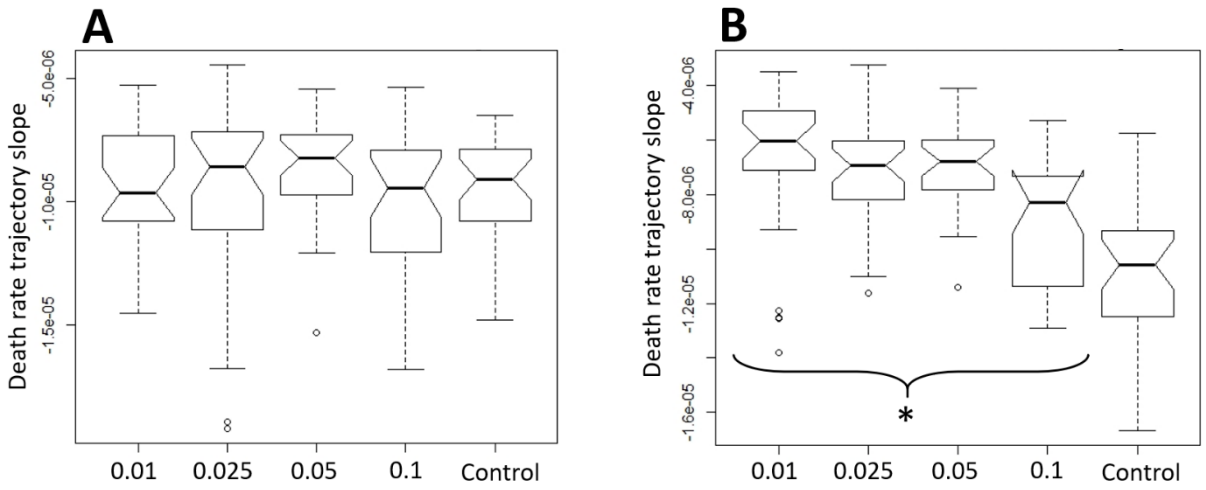

Figure 8: The effect of the fixed population-wide death rate on death rate evolution, in populations without kin structure. The figure shows the slope of the trajectory of the individually evolvable death rate for 4 different environmental oscillation speeds as well as for the control in which the environment remains stable, with 30 replicates for each. The results are comparable to those of Figure 7: under low fixed death rates (d= 0), there is no difference in the slopes of the death rate trajectory between simulations with environmental oscillations and control simulations (Figure 8 A). Under high fixed death rates (d = 0.075), the slopes are less steep under environmental oscillations compared with the control simulations (Figure 8 B). Parameter values for which the slopes of the death rate trajectories are significantly different from those in the control simulations are marked with an asterisk. Parameters other than the fixed death rate and the speed of environmental oscillations are kept constant (Nb = 50, k = 1).

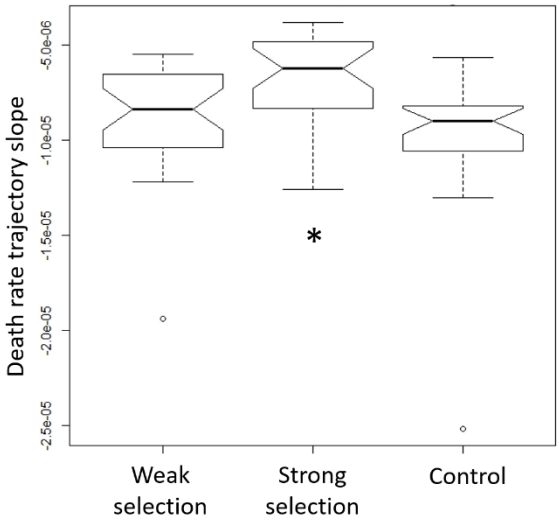

Figure 9: The effect of selection strength on death rate evolution, in kin structured populations. The figure shows the slope of the death rate trajectory for 2 different levels of selection strength, as well as for the control in which the environment remains stable, with 30 replicates for each. All death rate trajectories show a decline, however, under strong selection the decline is less steep. Parameter values for which the slopes of the death rate trajectories are significantly different from those in the control simulations are marked with an asterisk. Parameters other than the selection strength are kept constant (s = 0.01, d = 0.001, Nb = 10).

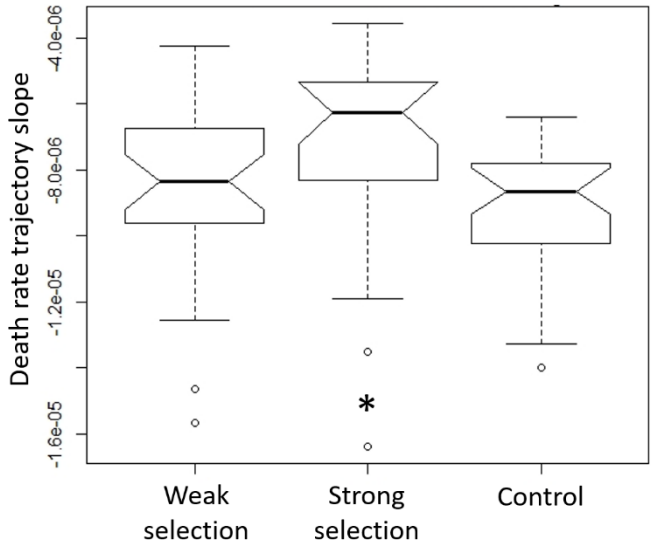

Figure 10: The effect of selection strength on death rate evolution, in populations without kin structure. The figure shows the slope of the death rate trajectory for 2 different levels of selection strength, as well as for the control in which the environment remains stable, with 30 replicates for each. The results are comparable to those from Figure 9 with kin structure: All death rate trajectories show a decline, under stronger selection the decline is less steep. Parameter values for which the slopes of the death rate trajectories are significantly different from those in the control simulations are marked with an asterisk. Parameters other than the selection strength are kept constant (s = 0.01, d = 0.001, Nb = 50).

Linking population dynamics and death rate evolution

Investigating the change in death rate during different phases of population dynamics can reveal more subtle effects of the selection for evolvability. In the first approach I separated stationary phases and crash phases according to the population size at any given point. I will present the results from setting the threshold at 80% of the median, changing the threshold to 90% of the median did not affect the results. For this analysis I chose to concentrate on simulations with fast oscillations (s = 0.1), high fixed death rates (d = 0.05) and strong selection (k = 3). In the control simulations with constant selection pressures and a fixed, flat fitness landscape, the population remains in the stationary phase, i.e. the population is stable and near its maximum during the entire simulation length (Figure 11). The change in death rate is significantly different from zero and negative. Under environmental oscillations, both stationary and crash phases can be observed (Figure 12). During the stationary phase, the change in death rate is significantly different from zero and negative; during the crash phases, the change in death rate is not significantly different from zero. This is observed both in simulations with kin structure (Figure 12 A), and in simulations without kin structure (Figure 12 B). The figures shown in this section are example plots illustrating the influence of the tested parameters, not all explored parameter combinations are shown here. Plots of the other parameter combinations are available in the data archive.

These results are based on a large sample size – during one simulation run time (disregarding the first 500 time steps), 950 data points for the change in death rate are calculated, which for 30 replicates results in a sample size of 28500 for each parameter combination (distributed unevenly across stationary and crash phases). To control for oversampling, I subsampled the dataset 1000 times and repeated the analysis. For histograms of p-values and mean changes in average death rate please see Appendix Figure 5 and Appendix Figure 6, as well as the data archive. Overall, the subsampling analysis confirms the results given above: most subsamples from stationary phases show significantly negative changes in death rate, whereas most subsamples from crash phases do not show changes in death rate that are significantly different from zero. This is true for simulations with and without kin structure.

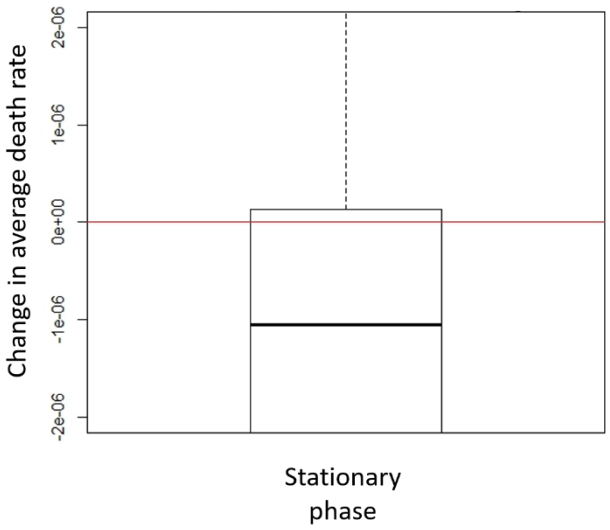

Figure 11: This figure shows the change in average death rate during intervals of 10 timesteps for control simulations with a stable environment. In control simulations, the population is in the stationary, stable phase only. The change in average death rate is significantly different from zero and negative ($p < 0.001$, mean = $-5.19 * 10^{-6}$). In other words, during the control simulations populations are consistently stable and the population average death rate declines, indicating evolution of immortality.

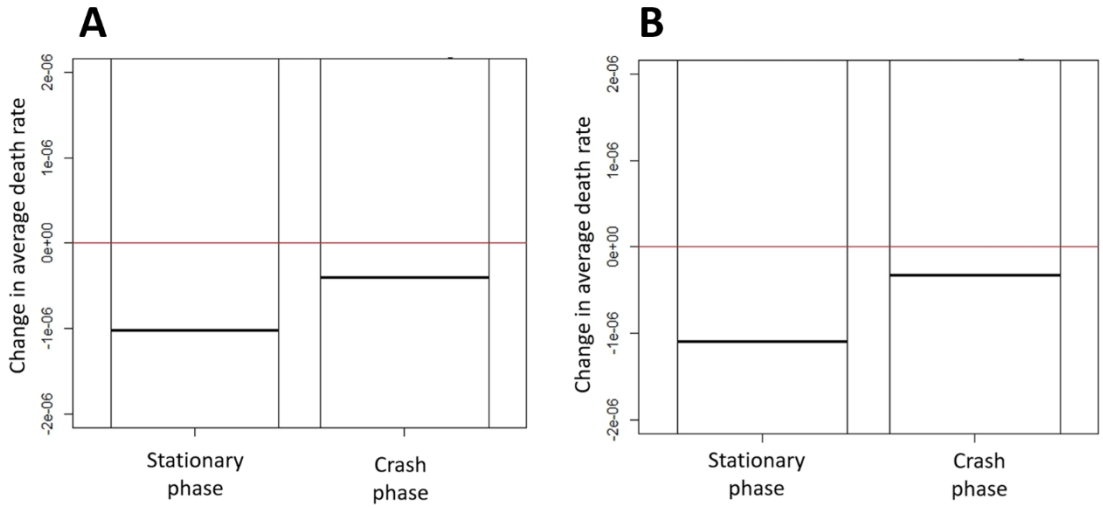

Figure 12: This figure shows the change in average death rate during intervals of 10 timesteps for populations with kin structure (Figure 12 A, Nb = 10) and without kin structure (Figure 12 B, Nb = 50). There is no clear difference between results shown in Figure 12 A and Figure 12 B, indicating that kin structure has no impact on the results. During stationary phases, the mean change in average death rate is negative (Figure 12 A: $p < 0.001$, mean = $-5.67 * 10^{-6}$; Figure 12 B: $p < 0.001$, mean = $-5.53 * 10^{-6}$). During the crash phase, the mean change in average death rate is not significantly different from zero (Figure 12 A: $p = 0.141$, mean = $2.47 * 10^{-6}$; Figure 12 B: $p = 0.662$, mean = $6.04 * 10^{-7}$).

In the second approach I separated stationary phases, crash phases and recovery phases according to the change in population size at any given point. As in the first approach, I chose to concentrate on

simulations with fast oscillations (s = 0.1), high fixed death rates (d = 0.05) and strong selection (k = 3). In the control simulations with constant selection pressures and a fixed, flat fitness landscape, the population remains in the stationary phase, i.e. the population is stable during the entire simulation length (Figure 13). The change in death rate is significantly different from zero and negative. Under environmental oscillations, stationary phases, crash phases and recovery phases can be observed (Figure 14). During the stationary phase, the change in death rate is significantly different from zero and negative. During the crash phases, the change in death rate is not significantly different from zero. This is observed both in simulations with kin structure (Figure 14 A), and in simulations without kin structure (Figure 14 B). Finally, during the recovery phase, the change in death rate is not significantly different from zero in simulations without kin structure (Figure 14 B), but is significantly different from zero and positive in simulations with kin structure (Figure 14 A). The figures shown in this section are example plots illustrating the influence of the tested parameters, not all explored parameter combinations are shown here. Plots of the other parameter combinations are available in the data archive.

As in the first approach, these results too are based on a large sample size, with 28500 data points for each parameter combination (distributed unevenly across stationary, crash and recovery phases). I again subsampled the dataset and repeated the analysis to control for oversampling. Overall, the subsampling analysis confirms the results given above: most subsamples from stationary phases show significantly negative changes in mean death rate, and most subsamples from crash phases do not show changes in mean death rate that are significantly different from zero. This is true for simulations with and without kin structure. Interesting are the results from the recovery phase: in simulations with kin structure, despite not all mean changes in death rate during the recovery phase being significantly positive, there is a clear bias towards low p values (Figure 15) – a pattern not seen without kin structure (Figure 16). Histograms of p-values and mean changes in average death rate for other parameter combinations and other subsample sizes are available in the data archive.

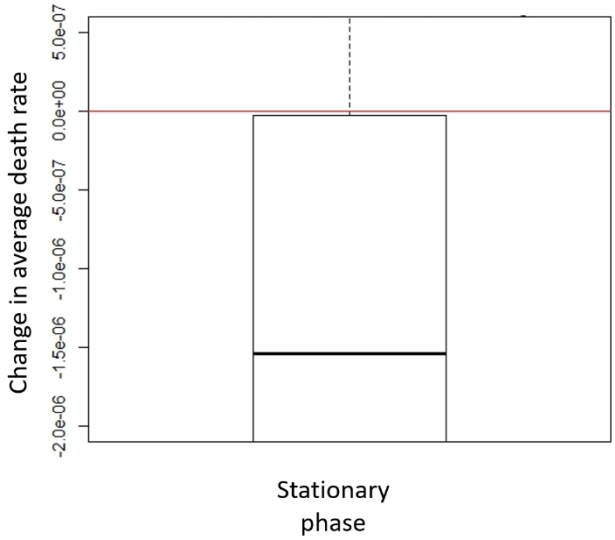

Figure 13: This figure shows the change in average death rate during intervals of 10 timesteps for control simulations with a stable environment. In control simulations, the population is in the stationary, stable phase only. The change in average death rate is significantly different from zero and negative ($p < 0.001$, mean = $-5.24 * 10^{-6}$). In other words, during the control simulations populations

are consistently stable and the population average death rate declines, indicating evolution of immortality.

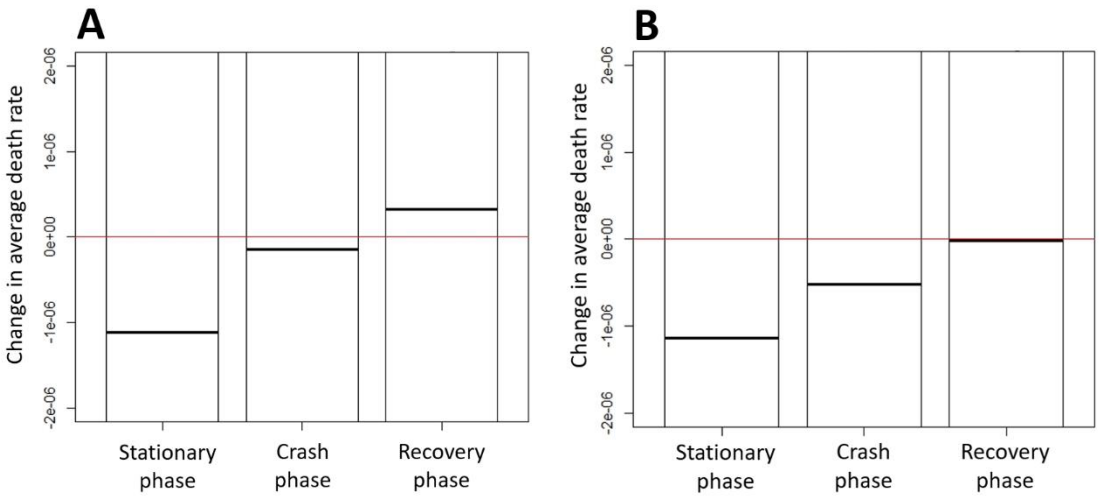

Figure 14: Change in average death rate during intervals of 10 timesteps for populations with kin structure (Figure 14 A, Nb = 10) and without kin structure (Figure 14 B, Nb = 50). During stationary phases, the mean change in average death rate is negative, regardless of kin structure (Figure 14 A: p < 0.001, mean = $-6.43 * 10^{-6}$; Figure 14 B: p < 0.001, mean = $-6.06 * 10^{-6}$). During the crash phase, the mean change in average death rate is not significantly different from zero, regardless of kin structure (Figure 14 A: p = 0.378, mean = $9.65 * 10^{-7}$; Figure 14 B: p = 0.829, mean = $-9.97 * 10^{-7}$). During the recovery phase, the mean change in average death rate is positive in kin structured populations (Figure 14 A: p = 0.001, mean = $8.63 * 10^{-6}$), but not significantly different from zero in populations without kin structure (Figure 14 B: p = 0.086, mean = $3.34 * 10^{-6}$). Kin structure thus impacts the change in average death rate during the recovery from a population crash.

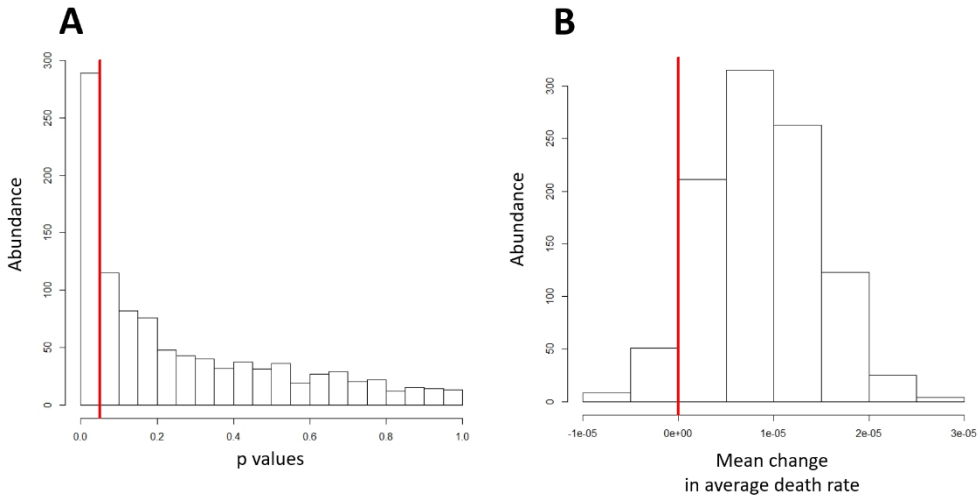

Figure 15: Histograms showing the distribution of p values (Figure 15 A) and of mean changes in average death rate (Figure 15 B) for 1000 subsamples of size 200. These data show how death rate changes in populations with kin structure (Nb = 10). Figure 15 A: despite over half of p values being above α = 0.05 (indicated by the red line), the abundance of p values declines sharply with increasing value, showing a clear bias towards low p values. Figure 15 B: The mean change in average death rate is clearly shifted to positive values (the red line indicates zero).

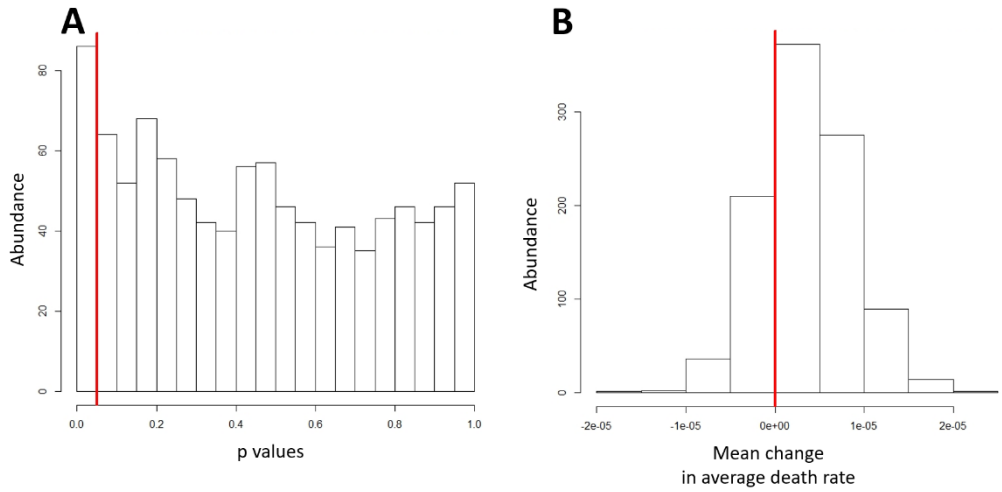

Figure 16: Histograms showing the distribution of p values (Figure 16 A) and of mean changes in average death rate (Figure 16 B) for 1000 subsamples of size 200. These data show how death rate changes in populations without kin structure (Nb = 50). Figure 16 A: only few of the p values are above α = 0.05 (indicated by the red line), and there is no bias towards low p values. The distribution of p values hence contrasts the one seen in Figure 15 A. Figure 16 B: The mean change in average death rate is shifted slightly to positive values (the red line indicates zero), but the shift is less obvious than the one observed in Figure 15 B.

Discussion

In this project I hypothesize that kin and evolvability selection combined lead to the evolution of increased death rates. This hypothesis could not be confirmed: In test and control scenarios alike the death rate declined until functional immortality was reached, with neither the shape of the evolutionary trajectory nor the final value reached showing an influence of selection for evolvability. There may be a transient effect of evolvability observed during population crashes, however this effect if indeed present is very slight. Overall, the question whether selection for evolvability can influence the evolution of death rates cannot be conclusively answered and merits further investigation.

Analysing the evolutionary trajectory of death rates

More specifically, there are indeed differences in death rate evolution between scenarios with oscillating and stable environments, as death rates decline less steeply in changing environments. Whilst this is in accordance with the tested hypothesis, this difference is observed even without kin structure, i.e. even in simulations in which kin selection cannot act. Since kin selection is an essential component of the proposed mechanism, the observed change in death rate evolution is not due to the evolvability effect that I here set out to test. Instead I propose that it may be explained by the selection for immortality being rendered less efficient by population dynamics induced by the oscillating environment: under sufficiently rapid and unpredictable environmental changes large parts of the population are maladapted, causing population crashes. During these population crashes, the efficiency of selection is reduced through increased drift in the small population, slowing down the evolution towards immortality. This is sufficient to explain the steeper decline in death rate in stable environments, observed both with and without kin selection, without the need to invoke evolvability. Furthermore, this explanation is compatible with the patterns seen for different parameter values: the least steep slopes of death rate trajectories are observed for fast environmental oscillations, high fixed

death rates and stronger selection. Fast oscillations increase the incidence of population crashes, as the population struggles more to adaptively track the environmental change, and high fixed death rates and stronger selection increase the severity of the population crashes, by increasing the overall death rate and reducing birth rates respectively. In other words, the least steep declines in death rates are observed under parameter values maximizing the strength of drift, and thus minimizing the strength of selection for immortality. These results also match the conclusions drawn from the mutation rate models, in which the same parameter values also lead to more frequent or severe population crashes, however here with the consequence of maximizing selection for higher mutation rates.

Analysing the change in death rate in relation to population dynamics

Somewhat more promising results can be obtained from a second approach: instead of considering the slope of the death rate trajectory during the entire simulation run time, the dynamics are analysed in small time slices at a time. When splitting the simulation according to whether the population is stable, declining or recovering, the changes in death rate for each of those categories show patterns supporting the hypothesized effect of selection for evolvability on death rates: in stable populations, the death rate declines in disproportionally many time slices. This reflects the observation discussed above: overall, death rates decline and immortality evolves. In other words, negative changes in death rate (declining death rates) are the norm. In contrast to this, the changes in death rate during the recovery phase are predominantly positive or at least not significantly different from zero, depending on the parameter combination. This suggests that there are different selective forces acting on the death rate during phases with different population dynamics: in stable populations, selection for immortality acts alone and drives death rates lower. However, if the population is maladapted to the current environment and hence undergoes a population crash, the situation is complicated by the addition of selection for evolvability: individuals carrying high-evolvability traits are more likely to find the correct phenotypic “solution” to the novel environmental challenge, and the evolvable trait can then hitchhike with this beneficial mutation during the population recovery. The increase in mean death rate during most recovery phases suggest that in most cases, the new phenotypic solution was found by individuals carrying a higher death rate – casting higher death rates as the high-evolvability trait and demonstrating that death rates can and do increase in response to selection for greater evolvability. Thus, the selection for evolvability counteracts, and partly even overrides, selection for immortality during the recovery from a population crash. Similar results are obtained when separating only stable populations and populations undergoing a crash (i.e. combining declining and recovering phases), but the results are less clear.

These patterns attest to the same principle observed in the stepwise increase in mutation rate in model 2 or the burst-like increases in mutation rate in model 3: Otherwise mal-adaptive alleles (increasing the number of deleterious mutations or reducing lifespan) can hitchhike with beneficial mutations, a process facilitated by oscillating environments and the resulting population crashes. It therefore fits the common theme that considering and explicitly modelling population dynamics can be of great importance in the field of evolvability.

However, there are various unresolved issues that may cast doubt on the validity of the results discussed in this section. The first, perhaps most glaring one, is the large sample size: whilst the differences between stable and crash phases are indeed significant, they are very slight, raising the concern that they might be an artefact of oversampling, rather than an indication of a true pattern. To control for this, I subsampled repeatedly with smaller sample sizes. The interpretation of the resulting histograms of p values and mean changes in death rate is not straight-forward, but in general a visual inspection shows that the same results discussed above can also be recovered for smaller sample sizes.

This holds true despite most p-values being non-significant: it is important to realise here a significantly negative change in death rate is the norm due to the selection for immortality, and changes in death rate not significantly different from zero are therefore also a potential result, and can indicate that selection for immortality and selection for evolvability cancel each other out enough to allow a constant death rate to be maintained. The second concern pertains distinguishing these results from the kin-selection control. As with the results discussed in the previous section on analyzing the death rate trajectory, it is necessary for these results as well to only be observed in kin structured populations, i.e. under the action of kin selection: selection for evolvability here acts via kin selection, and any effect observed in the absence of the latter thus cannot be attributed to the former. Interpreting the results from simulations without kin structure is not trivial and subject to some ambiguity: on the one hand, even without kin structure changes in death rate not significantly different from zero are observed, and stationary and recovery phases show clearly different patterns especially when inspecting the histograms of p values and mean changes in death rate – suggesting that the difference between stationary and recovery phases discussed above cannot be credited to evolvability selection in its entirety. Yet on the other hand, the pattern is far clearer in simulations with kin structure than in simulations without, demonstrating that kin selection is at least of some consequence. This is especially evident in the histograms of p values for the subsamples, which show obviously different patterns for simulations with and without kin structure. Based on these deliberations, no definite answer can be reached. The results from this analysis can thus be seen as weak and not entirely conclusive evidence in favor of death rates increasing under selection for evolvability – but evidence nonetheless.

This is, of course, still subject to the caveat that even if death rates do increase during the recovery from a population crash due to selection for evolvability, the effect would be of a very transient nature, as it neither prevents the end result of immortality nor slows down the evolution thereof.

Testing the model assumptions

Overall, no approach to the data yielded a conclusive proof of selection for evolvability causing elevated death rates. To confirm that this lack of effect is not merely due to errors in the model design, I tested that both selection for evolvability and kin selection are indeed acting. The former is confirmed by the results of the third model on mutation rate evolution: using the same model, but allowing mutation rate instead of death rate to evolve, I found a steep increase in mutation rate, demonstrating selection for evolvability. The latter is shown by death rates declining less rapidly in kin-structured populations: kin selection mitigates the cost of cost of dying earlier as it benefits related individuals by increasing their resource availabilities. Additional support comes from a previous paper using a very similar model design, which found both evidence of kin selection and of kin selection increasing death rates (Travis 2004). A comparison with this paper can be illuminating: Travis uses a very similar approach (a spatially explicit patch occupancy model in which limited dispersal distances give rise to kin structure) to test the hypothesis that dying, and thus freeing up resources for (future) relatives, is selected for assuming that future relatives are of high value. In Travis' study, this high value is implemented directly, by making reproductive success inversely proportional to age. Under these conditions, higher death rates evolve, as older individuals die to make space for new, related and more valuable ones (Travis 2004). This is strikingly similar to the hypothesis tested in this project, which essentially only replaces the value of future relatives: instead of a directly implemented reproductive benefit, younger individuals are of greater value in this study due to them representing a chance to be better adapted to the novel environment through mutation. It is puzzling therefore that the increase in death rate observed in Travis 2004 cannot be recovered in this project, despite the similarities in both approach and model design.

Interactions between population dynamics and immortality

The lack of effect of selection for evolvability on death rate evolution might be, at least partially, due to interference by the evolution of immortality. Not only does the selection for immortality overwhelm the opposed selection for evolvability, but more importantly immortality makes the population immune to population crashes: adaptation to the environment affects birth but not death rates, and if the current generation is immortal and not in need of replacing, then the population can persist even in absence of any reproduction, allowing maladapted populations – which would otherwise undergo a population crash – to maintain a stable size. This poses a technical problem for this model: evolution of higher death rates through selection for evolvability is absolutely dependant on population crashes, as it is during the recovery phases of these that higher death rates spread. In other words, as the simulation progresses, death rates decline, the population dynamics stabilize and selection for evolvability, facilitated by those dynamics, is stalled (see Figure 17).

Different parameter combinations amplify or dampen population crashes. I tried to counteract the effect of immortality by choosing parameter combinations optimally suited to promote population crashes. Unfortunately, this was not possible, as parameter combinations minimizing the actions of immortality later in the simulation (when individuals already have low death rates) also maximise the probability of extinctions early in the simulation (when individuals have high death rates and population crashes are therefore already severe). Simulations with the most promising and interesting parameter values hence had very short run times, hampering the understanding gained from the model. Despite exhaustively exploring the parameter space, I could not find any combination allowing something besides the extremes of immortality or extinction.

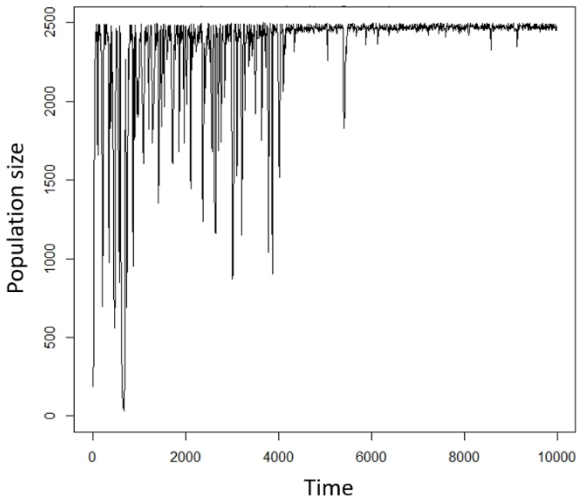

Figure 17: The effect of immortality on population dynamics. This plot shows the change in population size over time, for one example simulation (parameters: s = 0.1, d = 0.05, Nb = 3, k = 3). As the simulation progresses, death rates decline, eventually resulting in (functional) immortality. This has a profound effect on the population dynamics: In the beginning, the population undergoes severe population crashes, even approaching extinction at times. In the second half of the simulation, however, the population is largely stable, with barely any fluctuations in population size. Immortality thus markedly alters the population dynamics.

It is possible that the lack of elevated death rates under selection for evolvability stem from selection for immortality preventing the conditions facilitating the evolution of evolvability in the first place. Whilst this was indeed unfortunate with regards to this project, it opens up avenues for future research: There are already models developed in other contexts in which death rates evolve non-zero values (for an example, see Mitteldorf 2006). Implementing the here explored ideas in such a system can give an opportunity to study the impact of death rate without interference of the adverse effects immortality has on model behaviour.

Conclusion

Overall, this study investigates the effect of kin and evolvability selection on death rates. As death rates universally evolve values low enough to result in functional immortality, no strong effect of selection for evolvability can be observed. However, investigating the population dynamics more closely reveals that during the recovery from population crashes the death rate increases disproportionally often – potentially a subtle and transient signature of the selection for evolvability, thus lending some support to the hypothesis proposed in this study. Unfortunately, the analysis is hampered by the evolution of immortality impairing model function, and I therefore suggest that future studies could explore this question in settings preventing the evolution of immortality.

References

Bartke, A. (2001). Mutations prolong life in flies; implications for aging in mammals. *Trends in Endocrinology & Metabolism*, 12(6), 233-234.

Clark, W.R. (2004). Reflections on an unsolved problem of biology: the evolution of senescence and death. *Advances in Gerontology*, 14, 7–20

DNJ de Grey, A. (2015). Do we have genes that exist to hasten aging? New data, new arguments, but the answer is still no. *Current aging science*, 8(1), 24-33.

Finch, C. E., & Austad, S. N. (2001). History and prospects: symposium on organisms with slow aging. *Experimental gerontology*, 36(4-6), 593-597.

Goldsmith, T. C. (2008). Aging, evolvability, and the individual benefit requirement; medical implications of aging theory controversies. *Journal of theoretical biology*, 252(4), 764-768.

Goldsmith, T. C. (2014). Modern evolutionary mechanics theories and resolving the programmed/non-programmed aging controversy. *Biochemistry (Moscow)*, 79(10), 1049-1055.

Jones, O. R., Scheuerlein, A., Salguero-Gómez, R., Camarda, C. G., Schaible, R., Casper, B. B., ... & Quintana-Ascencio, P. F. (2014). Diversity of ageing across the tree of life. *Nature*, 505(7482), 169.

Kirkwood, T. B. (1977). Evolution of ageing. *Nature*, 270(5635), 301.

Libertini, G. (1983). Evolutionary arguments. *Società Editrice Napoletana*, Naples (Italy).

Libertini, G. (1988). An adaptive theory of the increasing mortality with increasing chronological age in populations in the wild. *Journal of theoretical biology*, 132(2), 145-162.

Libertini, G., Rengo, G., & Ferrara, N. (2017). Aging and aging theories. *Journal of Gerontology and Geriatrics*, 1, 59-77.

Lithgow, G. J., & Gill, M. S. (2003). Physiology: Cost-free longevity in mice?. *Nature*, 421(6919), 125.

López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., & Kroemer, G. (2013). The hallmarks of aging. *Cell*, 153(6), 1194-1217.

Lowe, D., Horvath, S., & Raj, K. (2016). Epigenetic clock analyses of cellular senescence and ageing. *Oncotarget*, *7*(8), 8524.

Masoro, E. J. (2003). Subfield history: caloric restriction, slowing aging, and extending life. *Science of Aging Knowledge Environment*, 2003(8), 2.

Medawar, P. B. (1952). An Unsolved Problem of Biology (London: Published for the college by HK Lewis).

Medawar, P. B. (1946). Old age and natural death. *Modern Quarterly*, *2*, 30-49.

Mitteldorf, J. (2006). Chaotic population dynamics and the evolution of ageing. *Evolutionary Ecology Research*, 8(3), 561-574.

Mitteldorf, J. (2010). Female fertility and longevity. *Age*, 32(1), 79-84.

Olshansky, S. J. (2018). From lifespan to healthspan. *Jama*, 320(13), 1323-1324.

Rando, T. A., & Chang, H. Y. (2012). Aging, rejuvenation, and epigenetic reprogramming: resetting the aging clock. *Cell*, 148(1-2), 46-57.

Rhinn, M., Ritschka, B., & Keyes, W. M. (2019). Cellular senescence in development, regeneration and disease. *Development*, 146(20).

Ricklefs, R. E., & Cadena, C. D. (2007). Lifespan is unrelated to investment in reproduction in populations of mammals and birds in captivity. *Ecology Letters*, 10(10), 867-872.

Tarín, J. J., Gómez-Piquer, V., García-Palomares, S., García-Pérez, M. A., & Cano, A. (2014). Absence of long-term effects of reproduction on longevity in the mouse model. *Reproductive Biology and Endocrinology*, 12(1), 84.

Travis, J. M. (2004). The evolution of programmed death in a spatially structured population. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, 59(4), B301-B305.

Vaupel, J. W., Baudisch, A., Dölling, M., Roach, D. A., & Gampe, J. (2004). The case for negative senescence. *Theoretical population biology*, 65(4), 339-351.

Wallace, A. R. (1889). The action of natural selection in producing old age, decay and death [a note by Wallace written “some time between 1865 and 1870”]. *Essays on hereditary and kindred biological problems*. Clarendon, Oxford.

Weismann, A., Poulton, E. B., Schönland, S., & Shipley, A. E. (1889). *Essays upon heredity and kindred biological problems.* Clarendon press.

Werfel, J., Ingber, D. E., & Bar-Yam, Y. (2015). Programed death is favored by natural selection in spatial systems. *Physical review letters*, 114(23), 238103.

Williams, G. C. (1957). Pleiotropy, natural selection, and the evolution of senescence. *Evolution*, 11(4), 398-411.

Evolution of evolvability

Part III: Dollo's law of irreversibility

Master thesis of Jana M. Riederer

s3507130

February 2019 – December 2019

(50 ECTS)

Supervised by prof. dr. F.J. Weissing & T.J.B van Eldijk

MSc Ecology & Evolution – Track Evolutionary Biology

Theoretical Research in Evolutionary Life Sciences

Rijksuniversiteit Groningen

rijksuniversiteit groningen

MARM GROUP

Evolution of evolvability III: Dollo's Law of Irreversibility

Manuscript submitted to PNAS

Title: Phenotypic plasticity explains violation of Dollo's law

One Sentence Summary: Fat synthesis is a plastic trait; such plasticity explains the maintenance of complex adaptations that are rarely used.

Authors: Bertanne Visser[1], Hans T. Alborn[2], Suzon Rondeaux[1], Manon Haillot[1], Thierry Hance[3], Darren Rebar[4], **Jana M. Riederer**[5], Stefano Tiso[5], Timo J.B. van Eldijk[5], Franz J. Weissing[5] & Caroline M. Nieberding[1]

Affiliations:

[1] Evolutionary Ecology and Genetics group, Biodiversity Research Centre, Earth and Life Institute, Université catholique de Louvain, Croix du Sud 4-5, 1348 Louvain-la-Neuve, Belgium.

[2]Chemistry Research Unit, Center of Medical, Agricultural, and Veterinary Entomology, Agricultural Research Service, United States Department of Agriculture, 1600 SW 23rd Drive, Gainesville, FL 32608, USA.

[3]Ecology of Interactions and Biological Control group, Biodiversity Research Centre, Earth and Life Institute, Université catholique de Louvain, Croix du Sud 4-5, 1348 Louvain-la-Neuve, Belgium.

[4]Department of Biological Sciences, Emporia State University, 1 Kellogg Circle, Campus Box 4050, Emporia, KS 66801, USA.

[5]Groningen Institute of Evolutionary Life Sciences, University of Groningen, Nijenborgh 7, 9747 AG Groningen, The Netherlands.

Abstract

Dollo's law of irreversibility states that once a complex adaptation has been lost in evolution, it will not be regained. Recently, various violations of this principle have been described. Here, we argue that the logic underlying Dollo's law only applies to traits that are constitutively expressed, while it fails in case of 'plastic' traits that are up- or downregulated according to needs. We tested this hypothesis for an archetypal violation of Dollo's law, the loss and regain of fat synthesis in parasitic wasps. Wasps from lineages that supposedly had lost lipogenic ability more than 200 million years ago were grown under various conditions. In line with our hypothesis, it turned out that fat synthesis had not been lost but was only switched on in low-fat environments. Such plasticity cannot only explain supposed violations of Dollo's law, but also the maintenance of adaptations to rarely occurring extreme events.

Main text

Universal laws are rare or absent in ecology and evolution, but Dollo's principle of the irreversibility of evolution seems to come close to such a law (*1*, *2*) (Dollo 1893 Bull Soc Belge Geolog; Gould 1970 J Hist Biol). In the numerous cases where a complex adaptation has been lost in the evolutionary history of a lineage (*3*) (Ellers et al 2012 Ecol Lett), there are only a handful examples where the adaptation has later been regained (*4*, *5*) (Collin and Miglietta 2008 TREE; Esfeld et al 2019 Curr Biol). Snakes did not regain legs, birds did not regain teeth, and ratites did not regain the ability to fly. Complex adaptations seem to be highly vulnerable: if they are of no use for an extended period of time, they can be expected to be selected against and/or to decay by genetic drift and the accumulation of deleterious mutations (*6*)(Lahti et al 2009 TREE); if they are later needed again, they cannot easily re-evolve (*7*)(Tan et al 2011 Phys Rev Lett), since the build-up of a complex trait requires the co-evolution of whole sets of genes that have to work together in harmony. If this logic is indeed correct, Dollo's principle may have important implications. In times of global climate change, previously rare and extreme events (like periods of drought or flooding) are expected to occur much more frequently. How can organisms cope with such events, if their adaptations to such events are irreversibly lost in the intervening period between such events?

The line of argumentation behind Dollo's law sounds plausible for constitutively expressed organs like legs or teeth. It is not surprising that cave fish lose their eyes when living in complete darkness over many generations, and that new eyes do not easily evolve once they change to a habitat where vision is an advantage. However, for three reasons the situation is different for adaptations that are phenotypically plastic (like most physiological and behavioral traits) and expressed according to the 'needs' of the local environment. First, such adaptations are not costly to maintain (and therefore only weakly selected against) in times when they are not needed, since they can just be switched off in such periods. Second, such adaptations will sporadically be expressed (and, hence, tested by selection) when some population members are exposed to non-standard conditions. And third, the regulatory pathways underlying these adaptations tend to combine the properties of robustness and evolvability, meaning that they do not easily decay when subjected to the accumulation of mutations and that only few mutational steps are required for restoring their function when conditions change (*8*)(Wagner 2007 Princeton book). Therefore, we hypothesize that (i) phenotypically plastic traits are much less affected by evolutionary decay and irreversible loss than constitutively expressed traits, and that (ii) many supposed violations of Dollo's law are unrecognized instances of phenotypic plasticity.

In the light of these considerations, we scrutinize an archetypal violation of Dollo's law, the apparent loss and regain of an essential metabolic trait: the synthesis of fat (*9*, *10*)(Visser, Ellers 2008; Visser et al 2010). Fat is synthesized when a surplus of sugars and other carbohydrates is available in the diet (*11*)(Turkish, Sturley 2009 Am J Physiol Gastrointest Liver Physiol), providing a reserve for future use. Fat is critical for survival and reproduction in nearly all living organisms. Fat reserves are built up, for example, in anticipation of periods when food is scarce, e.g., diapause in insects or hibernation in mammals, to fuel intense activities, including flight, as well as being a major component of egg cells in oviparous animals, such as birds, reptiles and insects (*12*, *13*)(Arrese, Soulages 2010; Muller et al 2017 JIP). The importance of fat and the ability to synthesize fat *de novo* explains why underlying metabolic and genetic pathways for fat synthesis are typically highly conserved, from bacteria to humans (*14*–*17*)(Jenke-Kodama et al 2005 Mol Biol Evol; Maier et al 2008 Science; Maier et al 2010 Q Rev Biophys; Bukhari et al 2014 Struct Des). Therefore, it came as a surprise that the ability to synthesize fat for storage appeared to have been repeatedly lost and regained in parasitic insects (*10*)(Visser et al 2010). Parasitic wasps lost the ability to synthesize fat more than 200 million years ago (MA) (*18*) (Peters et

al 2017 Curr Biol) and the trait re-appeared in several lineages, including in the genus *Leptopilina* (*12*, *21*) (Buffington et al 2007 Cladistics; Visser et al 2010) (about 80 MA). Replicated experiments with different populations of two *Leptopilina* species then revealed that some populations synthesized fat, while others did not (*10*, *20*–*23*)(Visser et al 2010; Eijs et al 1998 Ecol Entomol; Moiroux et al 2010 Ecol Entomol; Le Lann et al 2014 Oecologia; Visser et al 2018 Ecol Evol).

We hypothesized that these findings do not reflect the constitutive loss and regain of fat synthesis due to mutational changes in the metabolic pathway, but rather the plastic expression of fat synthesis in response to the local environment. Wasp development occurs in or on a host insect (*24*)(Godfray 1994); hence stored fat of the host can be carried over directly by the wasp. If the host contains plenty of fat, there is no need for *de novo* fat synthesis and the pathway should be shut off. If in contrast a fat-poor host is encountered, the wasp has to synthesize additional fat itself by activating the pathway. It is, therefore, conceivable that fat synthesis became plastic when the wasps started to parasitize fat-rich hosts (more than 200 MA) and that it was switched off except when the wasps encountered fat-poor hosts.

The question arises whether a switching device that is not used for extensive periods of time (more than 100 MY) should not be lost in the course of evolution. To investigate this, we ran individual-based simulations that monitored the sustained functionality of a switching device (a gene regulatory network that could decay by mutation) that is only sporadically used in evolutionary time (for modelling details see Supplementary Material 1: Modelling Study). Figure 1 shows that the switching device rapidly disintegrates (red simulations) if it is never used. However, even very infrequent use (pink: every 100 generations; purple: every 1000 generations) suffices to keep the switching device largely intact. Interestingly, the switching device does not erode gradually, but instead slowly evolves an improved performance over evolutionary time. An inspection of the evolving gene regulatory networks reveals that they become more and more robust (i.e., less and less affected by mutational decay), in line with earlier findings on network evolution (*8*).

From the simulations, we conclude that phenotypic plasticity (i.e., switching on or off of a metabolic pathway) can be maintained over long evolutionary time periods, even if plasticity is only needed sporadically. We, therefore, tested whether wasps that seemingly had lost the ability of fat synthesis were still able to synthesize fat in a low-fat environment. To this end, we let females from four field-caught populations of *L. heterotoma* develop on two naturally co-occurring host species: low-fat ("lean") *Drosophila simulans* and high-fat ("fat") *D. melanogaster* (containing 63 ± 3 µg and 91 ± 4 µg, mean ± 1SE storage fat, respectively; $F_{1,17} = 35.95$; $p < 0.0001$). When developing on the low-fat *D. simulans* host, the fat content of the wasps was in three of the four populations significantly higher after feeding than at emergence, indicating that fat synthesis had indeed occurred (Table 1). In contrast, the wasps showed only a marginal (and non-significant) increase when developing on the fatter *D. melanogaster*. These data suggest that fat synthesis does indeed depend on the host environment, at least in some wasp populations.

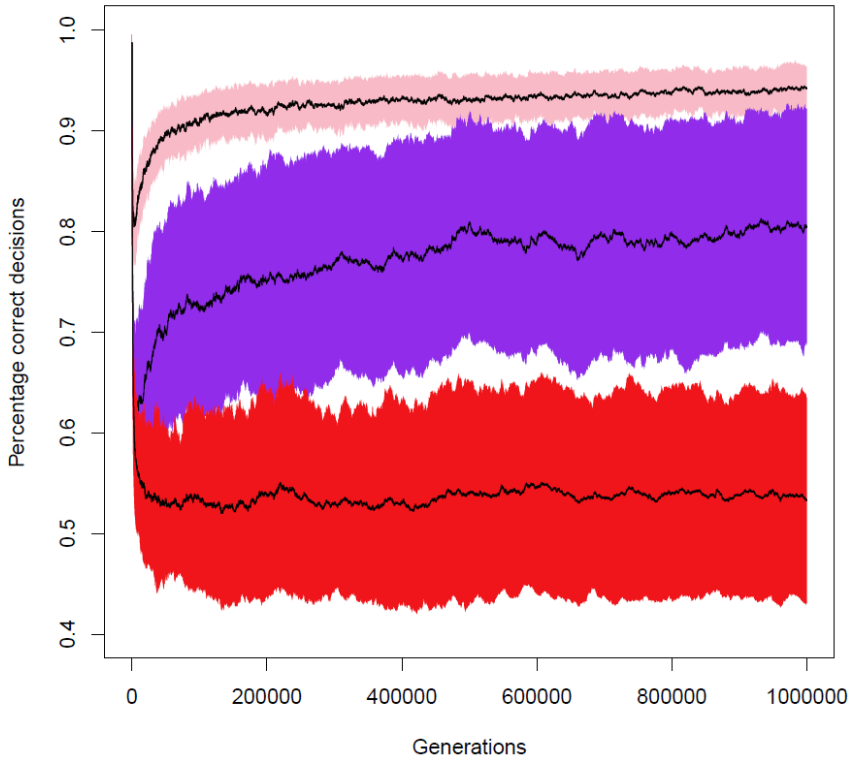

Fig. 1: Sporadic activation is sufficient for the maintenance of adaptive plasticity. *Long-term individual-based simulations showing how the performance of a gene-regulatory network (GRN) underlying adaptive plasticity changes in time when plasticity is only sporadically activated. We first evolved replicate GRNs in a variable environment where it is adaptive to switch on a metabolic pathway (fat synthesis) under low-fat conditions and to switch it off under high-fat conditions. In generation 0, a monomorphic population was established, where all N=10,000 individuals were endowed with the same well-performing GRN (different across replicates). Subsequently, the population evolved subject to selection, mutation (μ=0.001 per gene locus) and genetic drift in a fat-rich environment, where it is adaptive to constitutively switch off the metabolic pathway. Every 100 generations, we monitored the performance of a representative sample of GRNs (percentage correct decisions) in the original (fat-variable) environment: 1.0 means that the GRN is still making 100% adaptive decisions; 0.5 means that the GRN only makes 50% adaptive decision, as would be expected by a random GRN or a GRN that switches the pathway constitutively on or off. The colored graphs show the average performance (± standard deviation) of the GRNs for three scenarios (100 replicates per simulation). Red: the population never again encounters the fat-variable environment, leading to the loss of adaptive plasticity (convergence to 0.5). Pink: the individuals encounter a fat-variable environment on average every 100 generations; after an initial rapid drop in performance, a sustained high performance (>90% correct decisions) of the GRNs is regained after about 100,000 generations. Purple: the individuals encounter a fat-variable environment on average every 1000 generations; after an initial rapid drop in performance, an intermediate performance (>75% correct decisions) is regained gradually.*

Table 1: Wasps supposedly having lost lipogenic ability synthesize fat in a fat-poor environment. *Mean absolute fat amount ± 1se (in µg) was quantified in adult wasps from field-caught L. heterotoma populations raised on two hosts (fat-poor D. simulans, left part of the table; fat-rich D. melanogaster, right part of the table) and at two developmental stages (Emerged: just after emergence; Fed: having fed for 7 days after emergence). P-values reveal whether 7 days of feeding led to a significant increase in fat content, indicating the occurrence of fat synthesis. Three of the four populations tested on D. simulans exhibited fat synthesis on the lean host but no fat synthesis on the fat host, that is, plasticity in fat synthesis. (*) T-tests were performed when data was normally distributed and variances equal with (^) or without log transformation. The non-parametric Mann-Whitney U test was used for non-normal data or data with unequal variances (*).*

	Development on *D. simulans*				Development on *D. melanogaster*			
Population	Sample size	Emerged	Fed	p-value	Sample size	Emerged	Fed	p-value
Belgium 1**	10	17.50 ± 6.84	43.50 ± 11.38	**0.043**	38	36.00 ± 2.54	40.00 ± 3.22	0.336
Belgium 2	38	15.60 ± 1.02	36.83 ± 4.86	**<0.001 (*)**	32	38.50 ± 4.24	43.91 ± 3.45	0.331
UK 1	21	24.20 ± 1.49	30.91 ± 3.22	0.142 (^)	29	40.00 ± 4.41	44.30 ± 2.09	0.375
UK 2	-	-	-	-	17	33.60 ± 3.82	39.50 ± 1.50	0.522
Japan	20	12.20 ± 1.55	24.80 ± 4.86	**0.011 (^)**	13	29.17 ± 6.27	28.67 ± 8.19	0.964

The population-level comparison of wasp fat content at two points in time is only a crude measure that not always detects the occurrence of fat synthesis reliably. Even in case of active fat synthesis, fat content can stay constant or even decrease if, for example, fats are burned at a faster rate than at which they are produced (*25*)(Ament et al 2011 J Exp Biol). To unequivocally demonstrate that fat synthesis can be induced plastically, we turned to stable isotope tracing followed by GC-MS analyses (*26*, *27*) (Visser et al 2012 GBE; Visser et al 2017 Roy Soc Open Sci). Incorporation of stable isotopes after feeding depends on fat synthesis; hence a significant increase in stable isotope levels compared to controls (without access to stable isotopes) demonstrates active fat synthesis, even when lipids are burned. We used a split-brood family design where daughters of a single mother were allowed to develop on either lean *D. simulans* or fat *D. melanogaster*. Seventeen families, belonging to five field-caught populations, showed a (much) higher fat metabolism in the fat-poor environment (*D. simulans*) than in the fat-rich environment (*D. melanogaster*) (Figure 2; see supporting online table 1 for a statistical analysis). These results confirm that fat synthesis is indeed a plastic trait that is induced in response to low host fat content. Notice that the 17 families strongly differ in their environmental response, both in their baseline level of fat synthesis (on fat *D. melanogaster*) and in the slopes of their reaction norms. We conclude that fat synthesis is plastic and that there is (across-family) genetic variation in the degree of plasticity.

To rule out the possibility that the above results reflect a host-species effect (rather than an effect of the fat content of the environment), we repeated the experiment reported in Table 1, but now replacing lean *D. simulans* by lean *D. melanogaster* hosts. By reducing the sugar content in the diet of *D. melanogaster*, we were able to generate leaner flies (i.e. pupae containing 52 ± 3 µg storage lipids, mean ± 1SE, compared to 91 ± 4 µg storage lipids, mean ± 1SE; $F_{1,22} = 71.18$, $p < 0.0001$). In line with our earlier results, three out of four wasp populations showed indeed fat synthesis on these leaner *D. melanogaster* hosts (supporting online table 2). We conclude that plastic fat synthesis is induced by host fat content, rather than other traits differing between *D. melanogaster* and *D. simulans.*

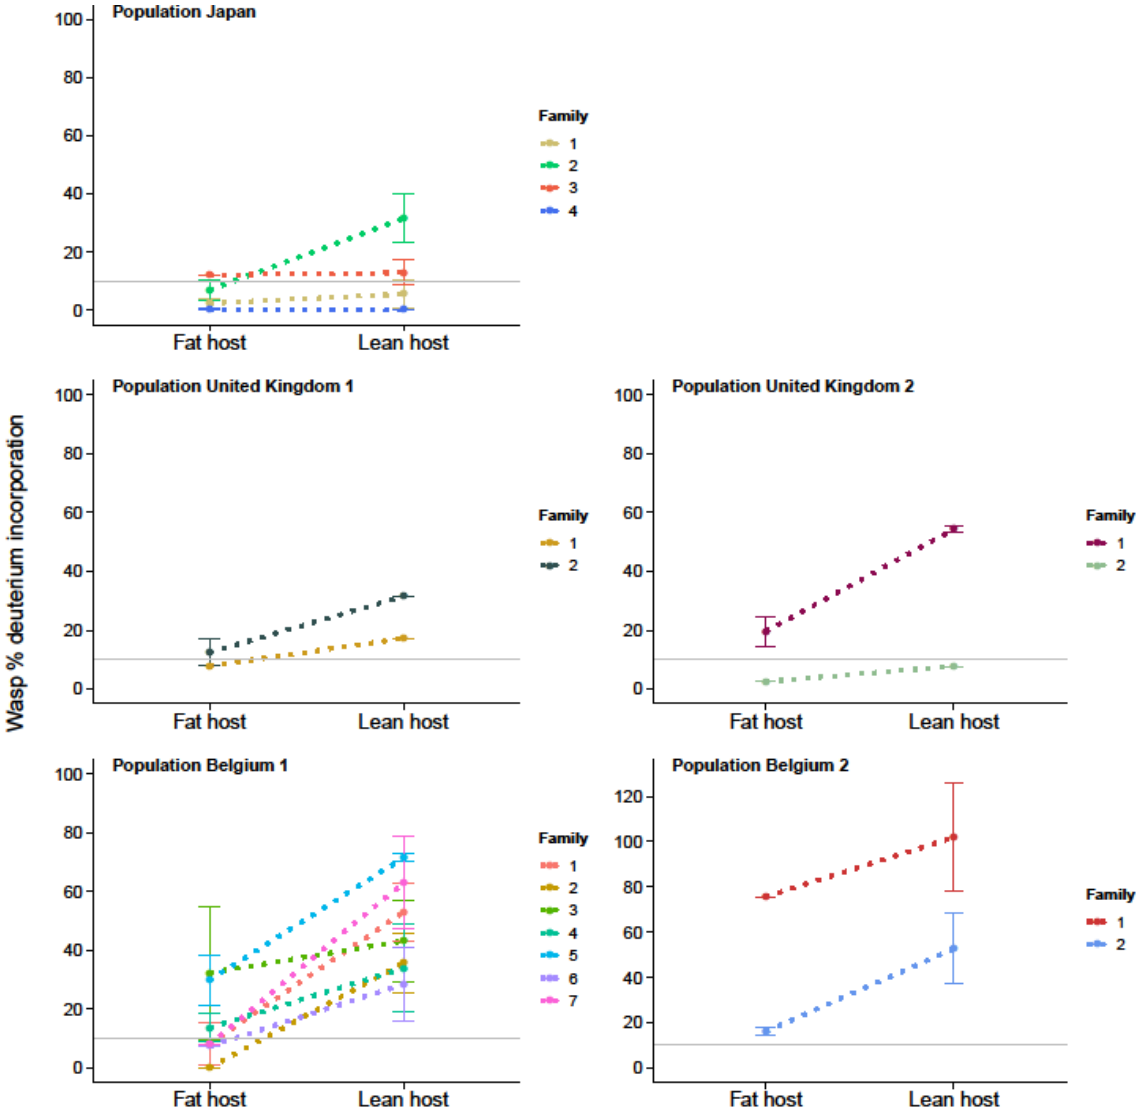

Fig. 2: Phenotypic plasticity in five field-caught wasp populations. *Incorporation of stable isotopes into the fatty acid fraction of offspring from 17 families developing in a fat-rich environment (fat host Drosophila melanogaster, left in each graph) and in a fat-poor environment (lean host D. simulans, right in each graph; n = 138). The horizontal gray line indicates that a stable isotope incorporation below 10% is considered insufficient evidence for the occurrence of fat synthesis.*

a

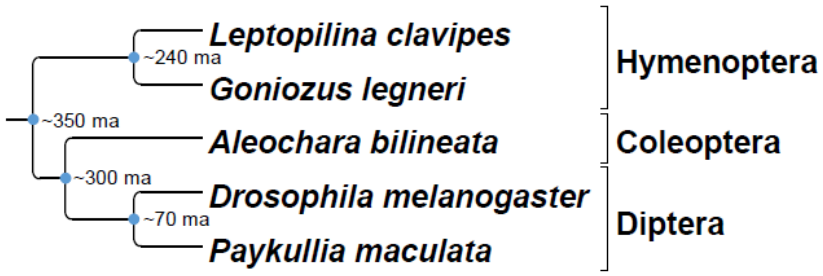

b Acetyl CoA carboxylase

		Leptopilina clavipes scf7180005162486		*Goniozus legneri* scaffold2431		*Aleochara bilineata* scaffold525		*Drosophila melanogaster* NM_136498.3		*Paykullia maculata* scaffold367327	
Domain	Accession	Interval	E-value	Interval	E-value	Interval	E-value	Interval	E-value	Interval	E-value
BC	COG0439	958-2451	5.31E-144	631-2124	2.65E-141	337-1812	3.05E-139	967-2460	1.24E-138	103-1596	1.71E-135
BCCP	cd06850	2851-3084	5.68E-12	2524-2718	4.63E-16	2230-2424	1.28E-13	2860-3045	2.70E-16	1996-2181	1.32E-16
ACCC	pfam08326	3085-5244	0	2719-4874	0	2425-4629	0	3055-5268	0	2191-4401	0
CT	pfam01039	5578-7191	2.44E-159	5221-6852	9.34E-155	4954-6663	3.65E-164	5587-7221	1.53E-160	4720-6138	2.54E-156

c Fatty acid synthase

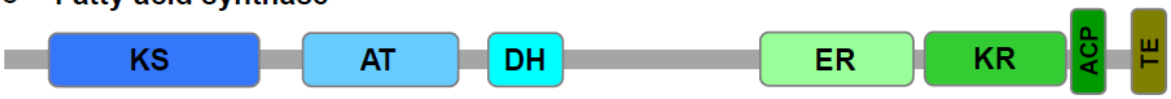

		Leptopilina clavipes scf7180005162060		*Goniozus legneri* scaffold5		*Aleochara bilineata* scaffold43971		*Drosophila melanogaster* FBtr0077659		*Paykullia maculata* scaffold4856	
Domain	Accession	Interval	E-value	Interval	E-value	Interval	E-value	Interval	E-value	Interval	E-value
KS	cd00833	381-1589	6.28E-164	85-1293	1.79E-166	52-1260	7.01E-170	742-1956	9.12E-163	412-1623	9.69E-165
AT	smart00827	1857-2708	1.48E-60	1561-2412	1.86E-64	1525-2376	5.65E-61	2227-3084	2.43E-61	1894- 2775	3.47E-62
DH	pfam14765	2925-3392	1.53E-15	2629-3225	1.23E-15	2626-3327	2.37E-09	3292-4026	4.16E-20	2959-3714	9.21E-15
ER	smart00829	4704-5576	2.79E-127	4408-5280	2.40E-130	4339-5211	2.12E-130	5101-5973	9.57E-126	4768-5640	3.58E-129
KR	cd08954	5637-6368	2.02E-79	5341-6051	9.53E-80	5266-5982	6.33E-77	6025-6747	1.24E-90	5692-6414	1.20E-87
ACP	smart00823	6348-6566	3.89E-07	6091-6249	1.55E-08	6019-6177	4.17E-07	6784-6942	4.39E-08	6451-6609	2.06E-07
TE	pfam00975	6762-7046	9.69E-15	6412-6699	5.41E-17	6346-6783	1.40E-19	7180-7497	4.82E-12	6814-7161	1.76E-17

Fig. 3: Conservation of two genes crucial for fatty acid synthesis in four parasitoid insects that supposedly had lost lipogenic activity. *Long evolutionary divergence times (up to 350 MA) separate the insect Drosophila melanogaster (that synthesizes lipids constitutively) and 4 parasitoid insects that were assumed to have lost the ability to synthesize lipids (10)(A). Acetyl coenzyme A carboxylase (ACC) and fatty acid synthase (FAS) are two essential genes for the production of fatty acids: the presence of all domains of ACC (B) and FAS (C) genes from D. melanogaster in the four parasitoid genomes reveals that the functional activity of the two genes is conserved in these insects. A table containing the detailed length and position of the different functional domains forming the two genes, as well as conservation level of the nucleotide sequence of the domains (e-values; the lower the e-value, the higher the significance of the match) are shown for each species. Abbreviations: BC = Biotin carboxylase; BCCP = Biotin carboxyl carrier protein; ACCC = Acetyl-coA carboxylase central region; CT = Carboxyl transferase domain; KS = Ketoacyl synthase; AT = Acyl transferase; DH = Dehydratase; ER = Enoyl reductase; KR = Ketoacyl reductase; ACP = Acyl carrier protein; TE = Thioesterase. Accession numbers refer to the conserved domain identifier on NCBI's Conserved Domain Database. Parasitoid transcript identifiers are provided underneath each species name.*

The ability to synthesize fat when being placed in a non-standard (low-fat) environment indicates that the key genes for fat synthesis have not lost their functionality in the *Leptopilina* genus. Making use of the fact that the genetic molecular pathway underlying fatty acid synthesis is highly conserved across animal taxa (*14*–*17*)(Jenke-Kodama et al 2005 Mol Biol Evol; Maier et al 2008 Science; Maier et al 2010 Q Rev Biophys; Bukhari et al 2014 Struct Des), we conducted a comparative analysis of coding sequences of acetyl coenzyme A carboxylase (ACC) (*28*)(Abu-Elheiga et al 2005 PNAS) and fatty acid synthase (FAS) (*16*)(Maier et al 2010 Q Rev Biophys), two enzymes that are critical for the production of fatty acids, the raw materials for storage fat. We used the *acc* and *fas* gene coding sequences of *D. melanogaster* as a starting point, because this fly readily synthesizes fat. Similar gene sequences were indeed found in the genome of *L. clavipes,* a sister species of *L. heterotoma,* and all functional domains of ACC and FAS enzymes were recovered, suggesting fully functional coding sequences in the *L. clavipes* genome (Figure 3). We then expanded our search for *acc* and *fas* functional coding sequences and protein domains to more distantly related parasitoids presumed to have lost fat synthesis independently (*10*)(Visser et al 2010): the hymenopteran *Goniozus legneri* (family Bethylidae), the dipteran *Paykullia maculata* (family Rhinophoridae), and the coleopteran *Aleochara bilineata* (family Staphilinidae)(Figure 3). ACC and FAS amino acid sequences of all these species aligned, suggesting that these two critical genes for fat synthesis have been conserved throughout the repeated evolution of parasitism in insects.

Summarizing, we have provided compelling evidence that parasitoid wasps have not lost their ability to synthesize storage fat in the distant past (>200 MA ago), but that fat synthesis is a plastic trait that is switched off when the wasps are developing in a fat-rich environment. It may come as a surprise that a crucial pathway like fat metabolism is not constitutively expressed, but activated or deactivated in response to environmental conditions. However, for parasitoid wasps this makes perfect sense, since they typically develop on fat-rich hosts that provide all the storage fat needed by the wasps. Yet, plasticity is required since there is considerable spatio-temporal variation in host availability and quality. *L. heterotoma* is a generalist wasp that can parasitize more than ten different *Drosophila* species that differ substantially in size and fat availability (*29*)(Fleury 2009 Adv Parasitol). Moreover, there is considerable geographic and seasonal variation in species diversity and community composition (*30*)(Lue et al 2018 Environ Entomol). Hosts are further patchily distributed with overlapping generations, suggesting considerable spatial variation at a local scale (*29*)(Fleury 2009). *Drosophila* are further well known to show large variation in starvation resistance, which is typically correlated with fat content (*31*)(Hoffmann et al 1999 Heredity). Plasticity in wasp fat synthesis is thus likely adaptive and evolved in response to highly variable environmental conditions in host fat content.

Previous documented cases of trait regain over long evolutionary time, in addition to the regaining of fat synthesis in parasitoids (*10*)(Visser et al 2010), include the regaining of wings in stick insects (*32*), the evolution of sexual reproduction from asexuality in mites (*33*)(Domes et al 2007 PNAS), and others (*34*)(Collin et al 2003 Proc R Soc B). These cases were all based on comparative analyses, which was shown to be problematic, because phylogenies do not necessarily provide a reliable representation of trait evolution (*35*–*38*)(Stone, French 2003 Curr Biol; Goldberg, Igic 2008 Evolution; Christin et al 2010 TREE; Galis et al 2010 Evolution;). Our results provide the first experimental evidence that macro-evolutionary patterns of trait reversals may in fact reflect trait plasticity: the trait is not “lost” or “regained” but is rather switched off or on, depending on environmental conditions. Intriguingly, such a regulatory switch can remain largely intact, even if it is only sporadically activated (Fig. 1). We consider it plausible that our findings are not restricted to fat metabolism in parasitoid wasps: the plastic regulation of trait expression could explain more cases of apparent trait loss and reappearance at macro-evolutionary time scales. Wing formation, for example, is often observed as an atavism (the sporadic occurrence of an ancestral phenotype) in otherwise wingless insects (*39*)(Hall 1984 Biol Rev),

and wing polymorphism, i.e. plasticity in wing development, is common in insects in general (*40*)(Zhang et al 2019 Ann Rev Entomol). Similarly, many asexual populations sporadically produce sexually reproducing individuals and plasticity in reproductive mode has evolved in several insect systems (*41–43*). Hence, plasticity may be a common principle explaining apparent violations of Dollo's law. As indicated by our simulation study plasticity can also explain the puzzling fact that adaptations to rare and extreme events are not lost, even if they are only sporadically used.

Acknowledgements

We are grateful to Jerry Coyne, Frietson Galis, Georges Lognay, Philippe Vernon, Claude Remacle, René Rezsohazy, and David Denlinger for helpful suggestions on an earlier draft of this manuscript. We would like to thank Christophe Pels for maintaining hosts and parasitoids and Hanno Hildenbrandt for help with writing the simulation program. Funding This work was supported by the Fonds de la Recherche Scientifique - FNRS under grants CDR and FNRS chargée de recherches n˚ 24905063 and 29109376 and the European Research Council (ERC, grant n˚ 789240). This is publication BRC 427 of the Biodiversity Research Centre. Author contributions Conceptualization: BV, CMN; Formal analysis: BV, CMN; Modelling study: JMR, ST, TJBvE, FJW; Funding acquisition: BV, CMN, FJW; Investigation: BV, HTA, SR, MH; Resources: DR; Supervision: CMN, TH, FJW; Writing - original draft: BV, CMN, FJW; Writing - review & editing: all authors. Competing interests Authors declare no competing interests. Data and materials availability: Data and the simulation program will be made available upon publication.

References

1. L. Dollo, Les lois de l'Evolution. *Bull. Soc. Belge Geol. Paleontol. Hydrol.* **7**, 164–166 (1893).
2. Gould Stephen Jay, Dollo on Dollo's Law, irreversibility and status of evolutionary laws. *J. Hist. Biol.* **3**, 189–212 (1970).
3. J. Ellers, E. T. Kiers, C. R. Currie, B. R. McDonald, B. Visser, Ecological interactions drive evolutionary loss of traits. *Ecol. Lett.* **15**, 1071–1082 (2012).
4. R. Collin, M. P. Miglietta, Reversing opinions on Dollo's law. *Trends Ecol. Evol.* **23**, 602–609 (2008).
5. K. Esfeld *et al.*, Pseudogenization and resurrection of a speciation gene. *Curr. Biol.* **28**, 3776--3786. (2019).
6. D. C. Lahti *et al.*, Relaxed selection in the wild. *Trends Ecol. Evol.* **24**, 487–496 (2009).
7. L. Tan, S. Serene, H. X. Chao, J. Gore, Hidden randomness between fitness landscapes limits reverse evolution. *Phys. Rev. Lett.* **198102**, 1–4 (2011).
8. A. Wagner, *Robustness and Evolvability in Living Systems.* (Princeton Univ. Press, 2007).
9. B. Visser, J. Ellers, Lack of lipogenesis in parasitoids: A review of physiological mechanisms and evolutionary implications. *J. Insect Physiol.* **54** (2008), pp. 1315–1322.
10. B. Visser, C. Le Lann, F. den Blanken, J. Harvey, J. van Alphen, Jacques J.M. Ellers, Loss of lipid synthesis as an evolutionary consequence of a parasitic lifestyle. *Proc. Natl. Acad. Sci.* **107**, 8677–8682 (2010).
11. A. R. Turkish, S. L. Sturley, The genetics of neutral lipid biosynthesis: an evolutionary perspective. *Am. J. Physiol. Metab.* **297**, E19--E27 (2009).
12. E. L. Arrese, J. L. Soulages, Insect fat body: Energy, metabolism, and regulation. *Annu. Rev. Entomol.* **55**, 207–225 (2010).

13. D. Muller *et al.*, Maternal age affects offspring nutrient dynamics. *J. Insect Physiol.* **101**, 123–131 (2017).
14. H. Jenke-Kodama, A. Sandmann, R. Müller, E. Dittmann, Evolutionary Implications of Bacterial Polyketide Synthases. *Mol. Biol. Evol.* **22**, 2027–2039 (2005).
15. T. Maier, M. Leibundgut, N. Ban, The crystal structure of a mammalian fatty acid synthase. *Science (80-.).* **321**, 1315–1323 (2008).
16. T. Maier, M. Leibundgut, D. Boehringer, N. Ban, Structure and function of eukaryotic fatty acid synthases. *Q. Rev. Biophys.* **43**, 373–422 (2010).
17. H. S. T. Bukhari, R. P. Jakob, T. Maier, Evolutionary origins of the multienzyme architecture of giant fungal fatty acid synthase. *Struct. Des.* **22**, 1775–1785 (2014).
18. R. S. Peters *et al.*, Evolutionary history of the Hymenoptera. *Curr. Biol.* **27**, 1013–1018 (2017).
19. M. L. Buffington, J. A. A. Nylander, J. M. Heraty, The phylogeny and evolution of Figitidae (Hymenoptera: Cynipoidea). *Cladistics*. **23**, 403–431 (2007).
20. I. E. M. Eijs, J. Ellers, G.-J. van Duinen, Feeding strategies in drosophilid parasitoids: the impact of natural food resources on energy reserves in females. *Ecol. Entomol.* **23**, 133–138 (1998).
21. J. Moiroux *et al.*, Local adaptations of life-history traits of a Drosophila parasitoid, Leptopilina boulardi: does climate drive evolution? *Ecol. Entomol.* **35**, 727–736 (2010).
22. C. Le Lann *et al.*, Rising temperature reduces divergence in resource use strategies in coexisting parasitoid species. *Oecologia*. **174**, 967–977 (2014).
23. B. Visser *et al.*, Variation in lipid synthesis, but genetic homogeneity, among *Leptopilina*parasitic wasp populations. *Ecol. Evol.* **8**, 7355–7364 (2018).
24. H. C. J. Godfray, *Parasitoids: Behavioural and evolutionary ecology* (Princeton University Press, West Sussex, 1994).
25. S. A. Ament *et al.*, Mechanisms of stable lipid loss in a social insect. *J. Exp. Biol.* **214**, 3808–3821 (2011).
26. B. Visser *et al.*, Transcriptional changes associated with lack of lipid synthesis in parasitoids. *Genome Biol. Evol.* **4**, 752–762 (2012).
27. B. Visser, D. S. Willett, J. A. Harvey, H. T. Alborn, Concurrence in the ability for lipid synthesis between life stages in insects. *R. Soc. Open Sci.* **4**, 160815 (2017).
28. L. Abu-Elheiga *et al.*, Mutant mice lacking acetyl-CoA carboxylase 1 are embryonically lethal. *Proc. Natl. Acad. Sci.* **102**, 12011–12016 (2005).
29. F. Fleury, P. Gibert, N. Ris, R. Allemand, Ecology and life history evolution of frugivorous *Drosophila*parasitoids. *Adv. Parasitol.* **70**, 3–44 (2009).
30. C. Lue, D. Borowy, M. L. Buffington, J. Leips, Geographic and seasonal variation in species diversity and community composition of frugivorous Drosophila (Diptera: Drosophilidae) and their Leptopilina (Hymenoptera: Figitidae) parasitoids. *Environ. Entomol.* **47**, 1096–1106 (2018).
31. A. R. Y. A. Hoffmann, L. G. Harshman, Desiccation and starvation resistance in Drosophila: patterns of variation at the species, population and intrapopulation levels. *Heredity (Edinb).* **83**, 637–643 (1999).
32. M. F. Whiting, S. Bradler, T. Maxwell, Loss and recovery of wings in stick insects. *Nature*. **421**, 264–267 (2003).
33. K. Domes, R. A. Norton, M. Maraun, S. Scheu, Re-evolution of sexuality breaks Dollo's Law. *Proc. Natl. Acad. Sci.* **104**, 7139–7144 (2007).

34. Fr. Collin, R. Cipriani, D. D. E. Ambientales, U. Simo, Dollo's law and the re-evolution of shell coiling. *Proc. R. Soc. B Biol. Sci.* **270**, 2551–2555 (2003).
35. G. Stone, V. French, Evolution: Have wings come, gone and come again? *Curr. Biol.* **13**, 436–438 (2003).
36. E. E. Goldberg, B. Igic, On phylogenetic tests of irreversible evolution. *Evolution (N. Y).* **62**, 2727–2741 (2008).
37. P. Christin, R. P. Freckleton, C. P. Osborne, Can phylogenetics identify C4 origins and reversals ? *Trends Ecol. Evol.* **25**, 403–409 (2010).
38. F. Galis, J. W. Arntzen, R. Lande, Dollo's law and the irreversibility of digit loss in Bachia. *Evolution (N. Y).* **64**, 2466–2476 (2010).
39. B. K. Hall, Developmental mechanisms underlying the formation of atavisms. *Biol. Rev.* **59**, 89–124 (1984).
40. C.-X. Zhang, J. A. Brisson, H.-J. Xu, Molecular mechanisms of wing polymorphism in insects. *Annu. Rev. Entomol.* **64**, 297–314 (2019).
41. D. J. Parker, J. Bast, K. Jalvingh, M. Robinson-rechavi, Repeated evolution of asexuality involves convergent gene expression changes. *Mol. Biol. Evol.* **36**, 350–364 (2018).
42. E. S. Tvedte, J. M. L. Jr, A. A. Forbes, Sex loss in insects: causes of asexuality and consequences for genomes. *Curr. Opin. Insect Sci.* **31**, 77–83 (2019).
43. E. R. Hanschen, M. D. Herron, J. J. Wiens, H. Nozaki, R. E. Michod, Repeated evolution and reversibility of self-fertilization in the volvocine green algae. *Evolution (N. Y).* **72**, 386–398 (2017).

Supplementary Materials

Materials and Methods 1: Modelling study

We consider the general situation where phenotypic plasticity is only sporadically adaptive and ask the question whether and under what circumstances plasticity can remain functional over long evolutionary time periods when the regulatory processes underlying plasticity are gradually broken down by mutations. To fix ideas, we consider a regulatory mechanism that switches on or off a pathway (like fat synthesis) in response to environmental conditions (e.g. fat content).

Fitness considerations

We assume that the local environment of an individual is characterized by two factors: fat content *F* and nutrient content *N*, where nutrients represent sugars and other carbohydrates that can be used to synthesize fat. Nutrients are measured in units corresponding to the amount of fat that can be synthesized from them. We assume that fitness (viability and/or fecundity) is directly proportional to the amount of fat stored by the individual. When fat synthesis is switched off, this amount is equal to *F*, the amount of fat in the environment. When fat synthesis is switched on, the amount of fat stored is assumed to be $N-c+(1-k)F$. This expression reflects the following assumptions: (i) fat is synthesized from the available nutrients, but this comes at a fitness cost *c*; (ii) fat can still be absorbed from the environment, but at a reduced rate (1-*k*). It is adaptive to switch on fat synthesis if $N-c+(1-k)F$ is larger than *F*, or equivalently if $F<\frac{1}{k}(N-c)$.

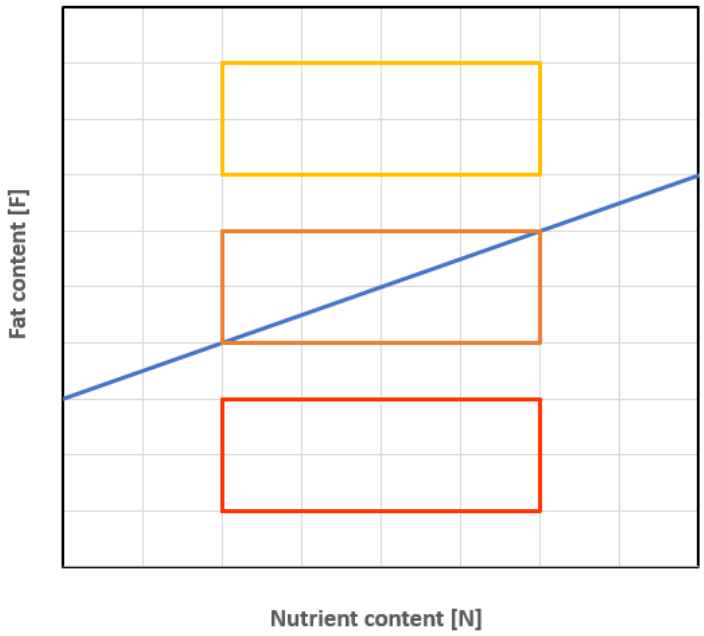

Supp Fig. 1: Environmental conditions encountered by the model organisms. *For a given combination of environmental nutrient content N and environmental fat content F, it is adaptive to switch on fat synthesis if (N, F) is below the blue line (corresponding to* $F<\frac{1}{k}(N-c)$*) and to switch it off otherwise. The three boxes illustrate three types of environment: a low-fat environment (red) where fat synthesis should be switched on constitutively; a high-fat environment (yellow) where fat synthesis should be switched off constitutively; and an intermediate-fat environment (orange) where a plastic switch is selectively favored.*

The right-hand side of this inequality is a straight line, which is illustrated by the blue line in Supplementary Figure 1. The three boxes in this figure illustrate three types of environmental conditions.

- Red box: low-fat environments. Here, $F < \frac{1}{k}(N-c)$ is always satisfied, implying that fat synthesis should be switched on constitutively.
- Yellow box: high-fat environments. Here, $F > \frac{1}{k}(N-c)$, implying that fat synthesis should be switched off constitutively.
- Orange box: intermediate-fat environments. Here, fat synthesis should be plastic and switched on if for the given environment (*N*, *F*) the fat content is below the blue line and switched off otherwise.

The simulations reported here were all run for the parameters $k = \frac{1}{2}$ and $c = \frac{1}{4}$.

Gene regulatory networks (GRN)

In our model, the switching device was implemented by an evolving gene regulatory network (as in REF: Van Gestel & Weissing 2016). The simulations shown in Fig. 1 of the main text are based on the simplest possible network that consists of two receptor nodes (sensing the fat and the nutrient content in the local environment, respectively) and an effector node that switches on fat synthesis if the combined weighted input of the two receptors nodes exceeds a threshold value *T* and switches it off otherwise. Hence, fat synthesis is switched on if $w_F F + w_N N > T$ (and off otherwise), where the weighing factors w_F and w_N and the threshold *T* are genetically determined evolvable parameters. We considered many alternative network structures (all with two receptor nodes and one effector node) and obtained very similar results (see below).

For the simple GRN described above, the switching device is 100% adaptive when the switch is on (i.e. $w_F F + w_N N > T$) if $F < \frac{1}{k}(N-c)$ and off otherwise. A simple calculation yields that this is the case if: $w_N > 0$, $w_F = -k w_N$ and $T = c w_N$.

Evolution of the GRN

For simplicity, we consider an asexual haploid population with discrete, non-overlapping generations and fixed population size $N = 10,000$. Each individual has several gene loci, each locus encoding one parameter of the GRN. (In case of the simple network described above, there are three gene loci, which encode the parameters w_F, w_N and T.) At the start of its life, each individual is placed in a randomly chosen environment (*N*, *T*). Based on its (genetically encoded) GRN, the individual decides on whether to switch on or off fat synthesis. If synthesis is switched on, the individual's fitness is given by $N - c + (1-k)F$; otherwise its fitness is given by *F*. Subsequently, the individuals produce offspring, where the number of offspring per individual is proportional to the individual's fitness. Each offspring inherits the genetic parameters of its parents, subject to mutation. With probability μ (per locus) a mutation occurs. In such a case the parental value is changed by a certain amount; the mutational step size is drawn from a normal distribution with mean zero and standard deviation σ. In the reported simulations, we chose $\mu = 0.001$ and $\sigma = 0.1$.

Preadaptation of the GRNs

Starting with randomly initialized population, we first let the population evolve in the intermediate-fat environment (orange box in Supp. Fig. 1) for 10,000 generations. In all replicate simulations, "perfectly adapted switch" (corresponding to $w_N > 0,\ w_F = -kw_N$ and $T = cw_N$) evolved, typically within 1,000 generations. These evolved networks were used to seed the populations in the subsequent "decay" simulations.

Evolutionary decay of the GRNs

For the decay experiments reported in Fig. 1 of the main text, we initiated a large number of monomorphic replicate populations with one of the perfectly adapted GRNs from the preadaptation phase. These populations were exposed for an extended period of time (1,000,000 generations) to a high-fat environment (yellow box in Supp. Fig. 1), where all GRNs switched off fat synthesis constitutively. However, in some scenarios, the environmental conditions changed back sporadically (with probability *q*) to the intermediate-fat environment, where it is adaptive to switch on fat metabolism in half of the environments. In Fig. 1, we report on the changing rates $q = 0.0$ (no changing back; red), $q = 0.001$ (changing back once every 1,000 generations; purple), and $q = 0.01$ (changing back once every 100 generations; pink). When such a change occurred, the population was exposed to the intermediate-fat environment for *t* generations (Fig. 1 is based on *t* = 3).

Throughout the simulation, the performance of the network was monitored every 100 generations as follows: 100 GRNs were chosen at random from the population, and each of these GRNs was exposed to 100 randomly chosen environmental conditions from the intermediate-fat environment. From this, we could determine the average percentage of "correct" decisions (where the network should be switched on if and only if $F < \frac{1}{k}(N - c)$. 1.0 means that the GRN is still making 100% adaptive decisions; 0.5 means that the GRN only makes 50% adaptive decision, as would be expected by a random GRN or a GRN that switches the pathway constitutively on or off. This measure for performance in the "old" intermediate-fat environment was determined for 100 replicate simulations per scenario and plotted in Fig. 1 (mean ± standard deviation).

Evolving robustness of the GRNs

The simulations in Fig. 1 are representative for all networks and parameters considered. Whenever $q = 0.0$, the performance of the regulatory switch eroded in evolutionary time, but typically at a much lower rate in case of the more complex GRNs. Whenever $q = 0.01$, the performance of the switch went back to levels above 90% and even above 95% for the more complex GRNs. Even for $q = 0.001$, a sustained performance level above 75% was obtained in all cases.

Intriguingly, in the last two scenarios the performance level first drops rapidly (from 1.0 to a much lower level, although this drop is less pronounced in the more complex GRNs) and subsequently recovers to reach high levels again. Apparently, the GRNs have evolved a higher level of robustness, a property that seems to be typical for evolving networks (REF: Book by Andreas Wagner). For the simple GRN studied in Fig. 1, this outcome can be explained as follows. The initial network was characterized by the genetic parameters $w_N > 0,\ w_F = -kw_N$ and $T = cw_N$ (see above), where w_N was typically a small positive number. In the course of evolutionary time, the relation between the three evolving parameters remained approximately the same, but w_N (and with it the other parameters) evolved to much larger values. This automatically resulted in an increasingly robust network, since mutations with a given step size distribution affect the performance of a network much less when the corresponding parameter is large in absolute value.

Materials and Methods 2: Experiments

Insects

Hosts and parasitoids were maintained as previously described *(1)*. Five Leptopilina heterotoma (Hymenoptera: Figitidae) populations were used for experiments: a population from Japan (Sapporo), two populations from the United Kingdom (1: Whittlesworth; 2: Great Shelford) and two populations from Belgium (1: Wilsele; 2: Eupen). Information on collection sites, including GPS coordinates, can be found in (*1*).

Determination of host fat content

D. simulans and *D. melanogaster* hosts were allowed to lay eggs over 24 hours in glass flasks containing ~50mL standard medium (described in (*1*)). After two days, developing larvae were sieved and ~200 were larvae placed in a *Drosophila* tube (ø x h 25x95; Dominique Dutscher) containing ~10mL medium. Seven days after egg laying, newly formed pupae were frozen at -18˚C, after which fat content was determined as described in (*1*), where dry weight before and after neutral fat extraction was used to calculate absolute fat amount (in µg) for each host. The host pupal stage was chosen for estimating fat content, because at this point the host ceases to feed, while the parasitoid starts consuming the entire host (*2*). All data were analyzed using R Project version 3.4.3 (*3*). Fat content of hosts was compared using a one-way ANOVA with host species as fixed factor.

Manipulating host fat content

To generate leaner *D. melanogaster* hosts, we adapted our standard food medium (*1*) to contain 100 times less (0.5g) sugar per liter water. Manipulating sugar content did not alter the structure of the food medium, thus maintaining similar rearing conditions, with the exception of sugar content. Fat content of leaner and fatter *D. melanogaster* hosts was determined and analyzed as described above.

Fat synthesis quantification with wasp populations

Mated female *L. heterotoma* were allowed to lay eggs on host fly larvae collected as described above with *ad libitum* access to honey as a food source until death. Honey consists of sugars and other carbohydrates that readily induce fat synthesis. After three weeks, adult offspring emergence was monitored daily and females were haphazardly placed in experimental treatments: emergence or feeding for 7 days on honey. Wasps were frozen at -18˚C after completion of experiments. Fat content was determined as described above. The ability for fat synthesis was then determined by comparing fat levels of recently emerged and fed individuals, similar to procedures described in (*4*, *5*). An increase in fat levels after feeding is indicative of active fat synthesis; equal or lower fat levels suggest fat synthesis did not take place. Each population was analyzed separately (*1*, *4*) for each host species using one-way ANOVAs with treatment as a fixed factor.

Fat synthesis quantification using a familial design and GC-MS analyses

To tease apart the effect of wasp genotype and host environment, we used a split-brood design where the offspring of each mother developed on lean *D. simulans* or fat *D. melanogaster* hosts. Two experiments were performed, one in which mothers were reared on *D. melanogaster* (experiment 1) and one in which mothers were reared on *D. simulans* (experiment 2). In both experiments, mothers were allowed to lay eggs in ~200 2nd to 3rd instar host larvae of one species for four days, after which ~200 host larvae of the other species were offered during four days. The order in which host larvae were presented was randomized across families. Following offspring emergence, daughters were allocated into two treatment groups: a control where females were fed a mixture of honey and water

(1:2 m/m) or a treatment group fed a mixture of honey and deuterated water (Sigma Aldrich)(1:2 m/m; stable isotope treatment) for 7 days. Samples were prepared for GC-MS as described in (*6*). Incorporation of up to three deuterium atoms can be detected, but percent incorporation is highest when only 1 deuterium atom is incorporated. As incorporation of a single atom unequivocally demonstrates active fat synthesis, we only analyzed percent incorporation (in relation to the parent ion) for the abundance of the m+1 ion. Percent incorporation was determined for five fatty acids, C16:1 (palmitoleic acid), C16:0 (palmitate), C18:2 (linoleic acid), C18:1 (oleic acid), and C18:0 (stearic acid), and the internal standard C17:0 (margaric acid). Average percent incorporation for C17:0 was 19.4 (i.e. baseline incorporation of naturally occurring deuterium) and all values of the internal standard remained within 3 standard deviations of the mean (i.e. 1.6). Percent incorporation of control samples was subtracted from treatment sample values to correct for background levels of deuterium (i.e. only when more deuterium is incorporated in treatment compared to controls fat is actively being synthesized).

For statistical analyses, percent incorporation was first summed for C16:1, C16:0, C18:2, C18:1 and C18:0 to obtain overall incorporation levels, as saturated C16 and C18 fatty acids are direct products of the fatty acid synthesis pathway (that can subsequently be desaturated). We then used a linear mixed effects model (% incorporation ~ host environment (fixed) + family (random) + population (random) + experiment (random) and found that neither population nor experiment contributed significantly to the model. The random term family contributed significantly to the model ($X = 8.697$; $df = 1$; $p = 0.003$), as did the fixed factor host environment (t-value = 4.346). To be able to determine the role of potential genotype by environment interactions, for subsequent analysis of interactions wasps developing on either lean or fat hosts were grouped according to percent deuterium incorporation: <10% incorporation is indicative of lacking fat synthesis, >10% incorporation is indicative of active fat synthesis. A threshold of 10% incorporation was chosen, because the range of natural variation in deuterium incorporation is ~2% (i.e. differences between the minimum and maximum percent incorporation of the internal standard C17:0). As we sum percent incorporation for five fatty acids the threshold value for active fat synthesis is thus 10%. More variation was observed for wasps developing on lean hosts, typically leading to higher standard errors compared to wasps developing on fat hosts. We then used a linear mixed effects model (GLMM, lme4 package) to test for the fixed effects of host (lean *D. simulans* and fat *D. melanogaster*) and family (--, -+, or ++), as well as the random effects of population (Japan, United Kingdom 1 and 2, Belgium 1 and 2) and experiment (experiment 1 and experiment 2)($n = 138$). Non-significant terms (i.e. population and experiment) were sequentially removed from the model to obtain the minimal adequate model (a GLM for which values are reported in Table 2).

Identification of functional *acc* and *fas* genes in distinct parasitoid species

To obtain *acc* and *fas* nucleotide sequences for L. *clavipes, G. legneri, P. maculata* and *A. bilineata*, we used *D. melanogaster* mRNA ACC transcript variant A (NM_136498.3 in Genbank) and FASN1-RA (FBtr0077659 in FlyBase) and blasted both sequences against transcripts of each parasitoid (using the blast function available at parasitoids.labs.vu.nl (*7*, *8*)). Each nucleotide sequence was then entered in the NCBI Conserved Domain database (*9*) to determine the presence of all functional protein domains. All sequences were then translated using the Expasy translate tool (web.expasy.org/translate/), where the largest open reading frame was selected for further use. Protein sequences were then aligned using MAFFT v. 7 (*10*) to compare functional amino acid sequences between all species.

Supplementary Table 1

Supplementary table 1: Fat synthesis is induced plastically based on host fat content. *We analyzed the data presented in Figure 2 in the main text statistically by means of a linear mixed effects model (GLMM, lme4 package) with host (lean D. simulans and fat D. melanogaster) as fixed effect, population (Japan, United Kingdom 1 and 2, Belgium 1 and 2), family within population, and experiment (each experiment was conducted twice) as random factors, and incorporation of stable isotopes as dependent variable (n = 138). Non-significant terms (i.e. population and experiment) were sequentially removed from the model to obtain the minimal adequate model (the GLM reported in the table). When referring to "families," we are referring to the comparison of daughters of singly inseminated females, which (in these haplodiploid insects) share 75% of their genome.*

	Estimate	Std. error	t-value	p
Intercept	1.602	0.327	4.905	**< 0.001**
Host simulans	1.802	0.391	4.607	**< 0.001**
Group "- -" for lipid synthesis	-0.785	0.566	-1.389	0.169
Group "+ +" for lipid synthesis	1.208	0.400	3.020	**0.003**
Host simulans x Group "- -"	-1.499	0.738	-2.033	**0.045**
Host simulans x Group "+ +"	-1.040	0.505	-2.061	**0.043**

Supplementary Table 2

Supplementary table 2: Fat synthesis is plastically induced in lean D. melanogaster hosts. *Mean absolute fat amount ± 1se (in µg) was quantified in adult wasps from field-caught L. heterotoma populations raised on lean D. melanogaster hosts at two developmental stages (Emerged: just after emergence; Fed: having fed for 7 days after emergence). Lean D. melanogaster hosts were created by rearing the larvae on a medium containing 100 times less sugar than usual P-values based on t-tests (with log-transformed data in case of the Belgium 2 population) reveal whether 7 days of feeding led to a significant increase in fat content, indicating the occurrence of fat synthesis.*

Population	n	Emerged	Fed	p-value
Belgium 2	31	27.40 ± 1.82	38.08 ± 3.50	**0.018 (^)**
UK 1	33	25.09 ± 2.51	40.50 ± 2.66	**<0.001**
UK 2	35	27.25 ± 2.60	38.62 ± 2.70	**0.006**
Japan	31	34.70 ± 2.72	34.36 ± 3.69	0.943

Supplementary references

1. B. Visser *et al.*, Variation in lipid synthesis, but genetic homogeneity, among Leptopilina parasitic wasp populations. *Ecol. Evol.* **8**, 7355–7364 (2018).

2. F. Fleury, P. Gibert, N. Ris, R. Allemand, Ecology and life history evolution of frugivorous *Drosophila* parasitoids. *Adv. Parasitol.* **70**, 3–44 (2009).

3. R Development Core Team, R: A language and environment for statistical computing, (available at http://www.r-project.org/).

4. B. Visser *et al.*, Loss of lipid synthesis as an evolutionary consequence of a parasitic lifestyle. *Proc. Natl. Acad. Sci.* **107**, 8677–8682 (2010).

5. B. Visser *et al.*, Transcriptional changes associated with lack of lipid synthesis in parasitoids. *Genome Biol. Evol.* **4**, 752–62 (2012).

6. B. Visser, D. S. Willett, J. A. Harvey, H. T. Alborn, Concurrence in the ability for lipid synthesis between life stages in insects. *R. Soc. Open Sci.* **4**, 160815 (2017).

7. K. Kraaijeveld *et al.*, Decay of sexual trait genes in an asexual parasitoid wasp. *Genome Biol. Evol.* **8**, 3685–3695 (2016).

8. K. Kraaijeveld, P. Neleman, J. Marien, E. de Meijer, J. Ellers, Genomic resources for *Goniozus legneri*, *Aleochara bilineata* and *Paykullia maculata*, representing three independent origins of the parasitoid lifestyle in insects. *G3 Genes Genomes Genet.* (2019).

9. A. Marchler-Bauer *et al.*, CDD/SPARCLE: Functional classification of proteins via subfamily domain architectures. *Nucleic Acids Res.* **45**, D200–D203 (2017).

10. K. Katoh, D. M. Standley, MAFFT multiple sequence alignment software version 7: Improvements in performance and usability. *Mol. Biol. Evol.* **30**, 772–780 (2013).

Evolution of evolvability : Appendix

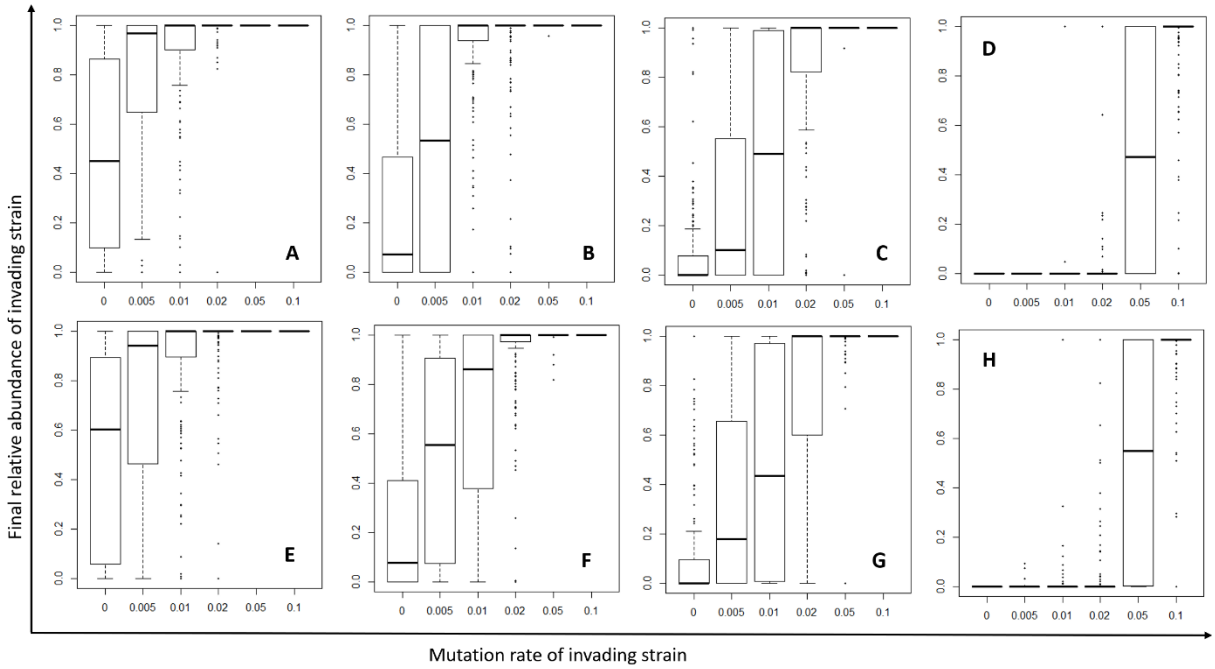

Appendix Figure 1: Frequencies of the invading mutation rate at the simulation end points, under oscillating resource levels. This figure shows the frequencies of the invading mutation rate at the end of the simulation, for six different invading mutation rates, four different resident mutation rates and two different oscillations speeds, with 200 replicates for each parameter combination. In plots A – D, the period length of the environmental oscillations is 125.66, in plots E – H the period length of the environmental oscillations is 62.83. The resident mutation rates are the following: in plots A and E: 0, in plots B and F: 0.005, in plots C and G: 0.01, and in plots D and H: 0.05. In all plots, the higher mutation rate, be it resident or invading, achieves a higher frequency at the end of the simulation. In cases where resident and invading mutation rate are of equal value, the frequencies of the invading mutation rate is distributed evenly around 0.5.

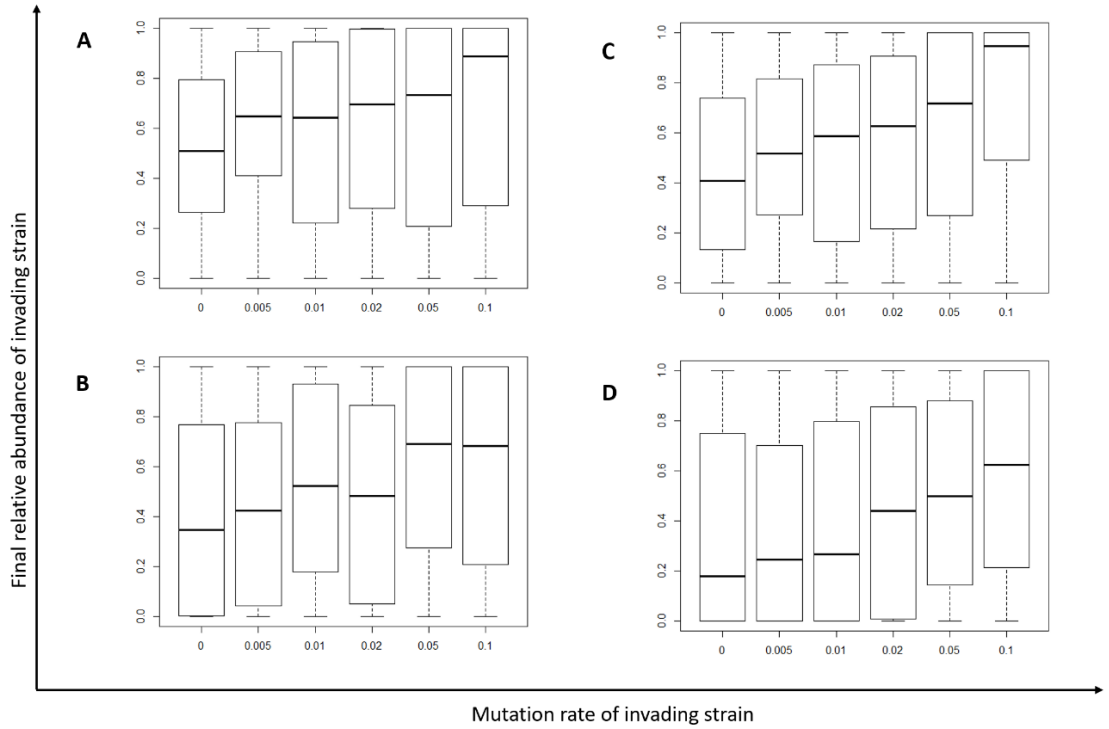

Appendix Figure 2: Frequencies of the invading mutation rate at the simulation end points under negative frequency dependent selection. This figure shows the frequencies of the invading mutation rate at the end of the simulation, for six different invading mutation rates and four different resident mutation rates, with 200 replicates for each parameter combination. The resident mutation rates are the following: in plot A: 0, in plot B: 0.005, in plot C: 0.01, and in plot D: 0.05. Generally, the higher mutation rate, be it resident or invading, achieves a higher frequency at the end of the simulation. However, this pattern is not as clear as it was for mutation rate evolution under environmental oscillations. In cases where resident and invading mutation rate are of equal value, the frequencies of the invading mutation rate are distributed evenly around 0.5.

Appendix Table 1: The correlation between the change in mutation rate (averaged across all individuals with a specific genotype) and the number of individuals with this genotype in Model 2. This table shows the p-values and correlation coefficients for 30 replicates of a correlation analysis testing the correlation between the change in mutation rate, averaged across all carriers of genotype A, and the number of individuals carrying genotype A (second and third column), and the correlation in mutation rate, averaged across all carriers of genotype B, and the number of individuals carrying genotype B (fourth and fifth column). In the correlation analysis I used Pearson's correlation coefficient.

Replicate Number	Gen. A p-value	Gen. A Correlation Coefficient	Gen. B p-value	Gen. B Correlation Coefficient
1	< 0.0001	-0.2795	< 0.0001	-0.2756
2	< 0.0001	-0.2219	0.0001	-0.1784
3	< 0.0001	-0.2026	< 0.0001	-0.2288
4	< 0.0001	-0.2753	< 0.0001	-0.2794
5	< 0.0001	-0.2604	< 0.0001	-0.2608
6	< 0.0001	-0.2313	< 0.0001	-0.2221
7	< 0.0001	-0.2506	< 0.0001	-0.2821
8	< 0.0001	-0.2813	< 0.0001	-0.2579
9	< 0.0001	-0.3156	< 0.0001	-0.3010
10	< 0.0001	-0.2612	< 0.0001	-0.2554
11	< 0.0001	-0.2244	< 0.0001	-0.2178
12	< 0.0001	-0.2265	< 0.0001	-0.2431
13	< 0.0001	-0.2681	< 0.0001	-0.3274
14	< 0.0001	-0.2410	< 0.0001	-0.2400
15	< 0.0001	-0.2842	< 0.0001	-0.2721
16	< 0.0001	-0.2134	< 0.0001	-0.2501
17	< 0.0001	-0.2698	< 0.0001	-0.2491
18	< 0.0001	-0.2459	< 0.0001	-0.2681
19	< 0.0001	-0.2324	< 0.0001	-0.2592
20	< 0.0001	-0.2582	< 0.0001	-0.2430
22	< 0.0001	-0.2716	< 0.0001	-0.2732
23	< 0.0001	-0.2237	< 0.0001	-0.2748
24	< 0.0001	-0.2623	< 0.0001	-0.2513
25	< 0.0001	-0.2427	< 0.0001	-0.2522
26	< 0.0001	-0.2638	< 0.0001	-0.2282
27	< 0.0001	-0.2479	< 0.0001	-0.2780
28	< 0.0001	-0.2897	< 0.0001	-0.2852
29	< 0.0001	-0.2434	< 0.0001	-0.2682
30	< 0.0001	-0.2258	< 0.0001	-0.2682

Appendix Figure 3: Evolution of the mutation rate in the control scenario of Model 2, with a constant fitness pay-off without environmental fluctuations. In the control scenario, no increase in mutation rate is observed beyond what is consistent with drift and a bounded random walk. Parameter values for these two plots are: k=1, Nb= 50, d=0.075. Appendix Figure 3A: In the control scenario, some replicates seem to evolve a higher mutation rate, however this increase is very weak, especially compared to the strong increase seen under environmental fluctuations. This weak increase is consistent with a bounded random walk: through drift alone some replicates will achieve a higher mutation rate, and some replicates will achieve a lower mutation rate. However, since the mutation rate was initialized close to zero, replicates which through drift achieve a lower mutation rate stop at zero (since negative mutation rates are biologically not sensible) – the random walk is bounded at zero. The lower half of the random walk is therefore effectively "chopped-off", resulting in the semblance of a weak increase in mutation rate. Appendix Figure 3B: To check that the weak increase in the mutation rate in the control simulations is indeed due to it being a bounded random walk, I repeated the control scenario but initialized the mutation rate at a higher value (0.5). In this scenario, the replicates of mutation rate show a "cone shape" centered around the value of initialization (shown by the red line at average mutation rate = 0.5), demonstrating that the mutation rate does indeed not evolve a higher average value, and changes through drift in a random manner. Essentially, the difference between 3A and 3B is that in 3A the random walk is bounded at zero, removing the lower half of the cone shape seen in 3B.

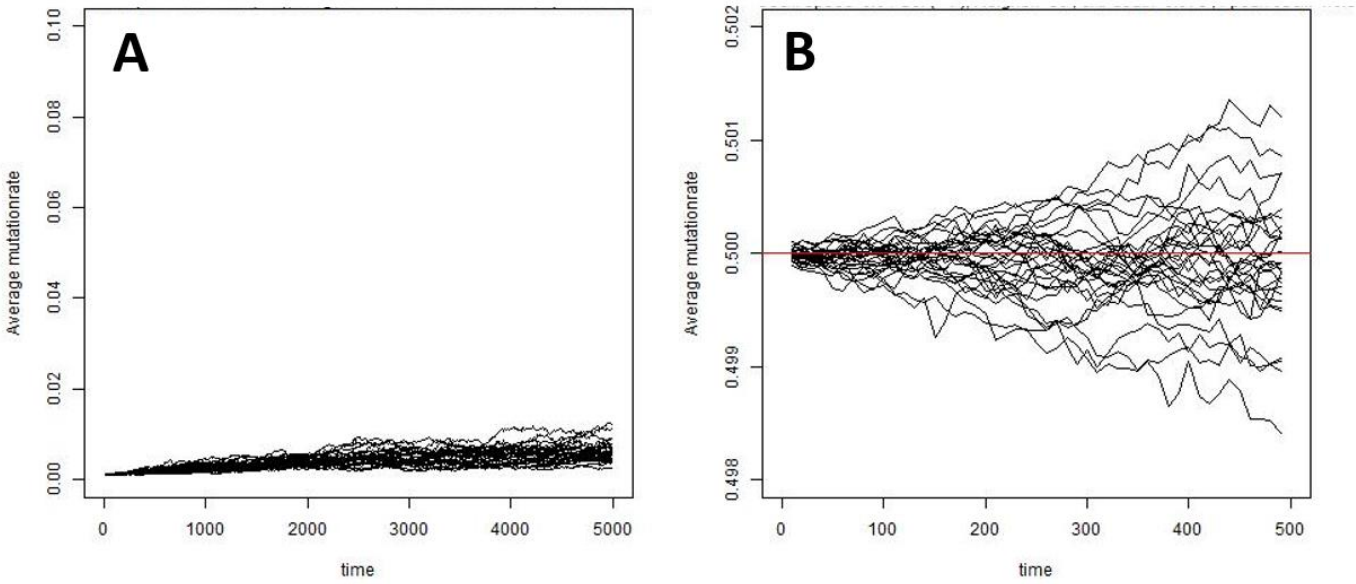

Appendix Figure 4: Evolution of the mutation rate in the control scenario of Model 3, with a constant fitness pay-off without environmental fluctuations. In the control scenario, no increase in mutation rate is observed beyond what is consistent with drift and a bounded random walk. Parameter values for these two plots are: k=1, Nb= 1, d=0.001. Appendix Figure 4A: In the control scenario, some replicates seem to evolve a higher mutation rate, however this increase is very weak, especially compared to the strong increase seen under environmental fluctuations. This weak increase is consistent with a bounded random walk: through drift alone some replicates will achieve a higher mutation rate, and some replicates will achieve a lower mutation rate. However, since the mutation rate was initialized close to zero, replicates which through drift achieve a lower mutation rate stop at zero (since negative mutation rates are biologically not sensible) – the random walk is bounded at zero. The lower half of the random walk is therefore effectively "chopped-off", resulting in the semblance of a weak increase in mutation rate. Appendix Figure 4B: To check that the weak increase in the mutation rate in the control simulations is indeed due to it being a bounded random walk, I repeated the control scenario but initialized the mutation rate at a higher value (0.5). In this scenario, the replicates of mutation rate show a "cone shape" centered around the value of initialization (shown by the red line at average mutation rate = 0.5), demonstrating that the mutation rate does indeed not evolve a higher average value, and changes through drift in a random manner. Essentially, the difference between 4A and 4B is that in 4A the random walk is bounded at zero, removing the lower half of the cone shape seen in 4B.

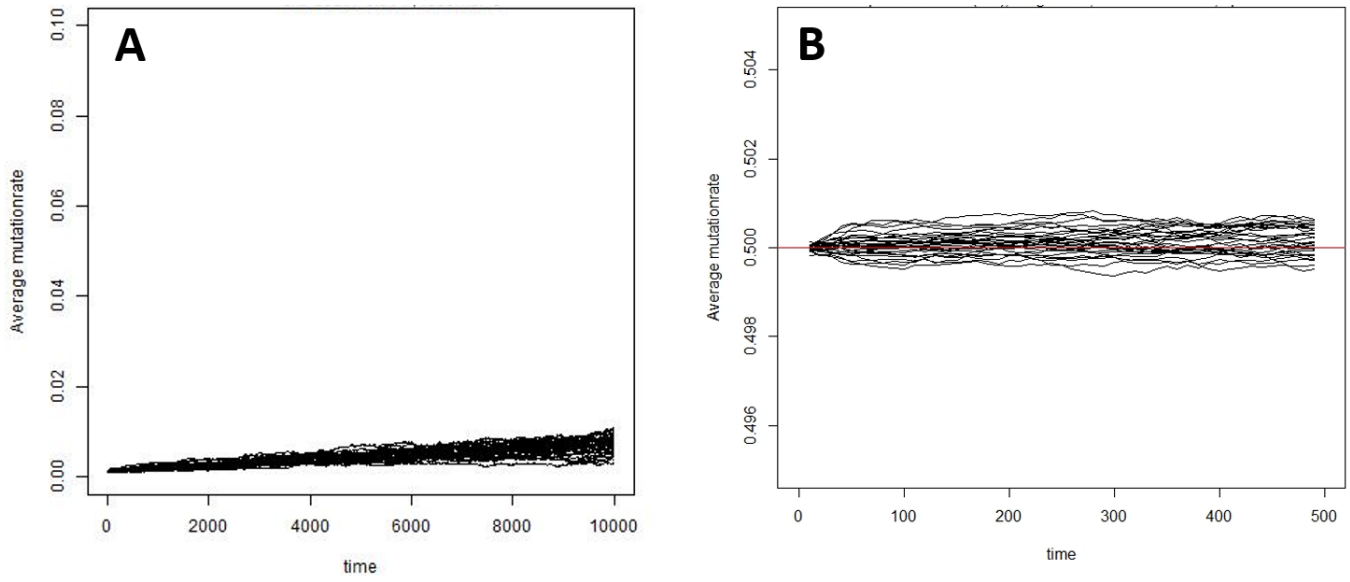

Appendix Table 2: The correlation between the change in average mutation rate and population size in Model 3. This table shows the p-values and correlation coefficients for 30 replicates of a correlation analysis testing the correlation between the change in average mutation rate and population size. In the correlation analysis I used Pearson's correlation coefficient. Statistically significant replicates are shown in green.

Replicate Number	p-value	Correlation Coefficient
1	< 0.0001	-0.1534
2	0.0671	-0.0595
3	< 0.0001	-0.2711
4	< 0.0001	-0.2216
5	0.0246	-0.0730
6	0.5689	0.0185
7	< 0.0001	-0.2270
8	< 0.0001	-0.3569
9	0.0237	-0.0734
10	0.0104	-0.0832
11	< 0.0001	-0.2024
12	< 0.0001	-0.4149
13	< 0.0001	-0.1686
14	< 0.0001	-0.3385
15	< 0.0001	-0.2180
16	< 0.0001	-0.1485
17	< 0.0001	-0.3089
18	< 0.0001	-0.3766
19	< 0.0001	-0.2335
20	< 0.0001	-0.2371
21	< 0.0001	-0.2916
22	< 0.0001	-0.4310
23	< 0.0001	-0.1589
24	0.4837	0.0228
25	< 0.0001	-0.1913
26	0.0064	-0.0884
27	< 0.0001	-0.1757
28	< 0.0001	-0.1631
29	< 0.0001	-0.3157
30	0.1060	-0.0525

Appendix Figure 5: Subsampling the analysis of the change in average death rate during stationary population dynamic phases in Model 3. As the analysis of the change in death rate during different population dynamic phases (see Part 2 of the thesis, subheading "Linking population dynamics and death rate evolution"), I subsampled the dataset 1000 times and repeated the analysis on each subsample, resulting in 1000 p-values and 1000 mean changes in death rate. This figure thus shows histograms of the p-values (Appendix Figure 5A & 5C) and of the mean changes in death rate (Appendix Figure 5B & 5D) for 1000 subsamples of subsample size N_{sub}=100 (Appendix Figure 5A & 5B) and for 1000 subsamples of subsample size N_{sub}=500. Histograms for other subsample sizes are available in the data archive. The red line indicates 0.05 in Appendix Figure 5A and 5C, and 0 in Appendix Figure 5B and 5D. The parameter values used in this analysis are s = 0.1, k = 3, Nb = 10 and d = 0.05, histograms for other parameter combinations are available in the data archive. The histograms shown here pertain to the subsampling analysis carried out in Approach 1 (categorizing population dynamic phases by population size), histograms for the second approach are available in the data archive. Overall, most subsamples in the here shown stationary phase show significantly negative changes in death rate.

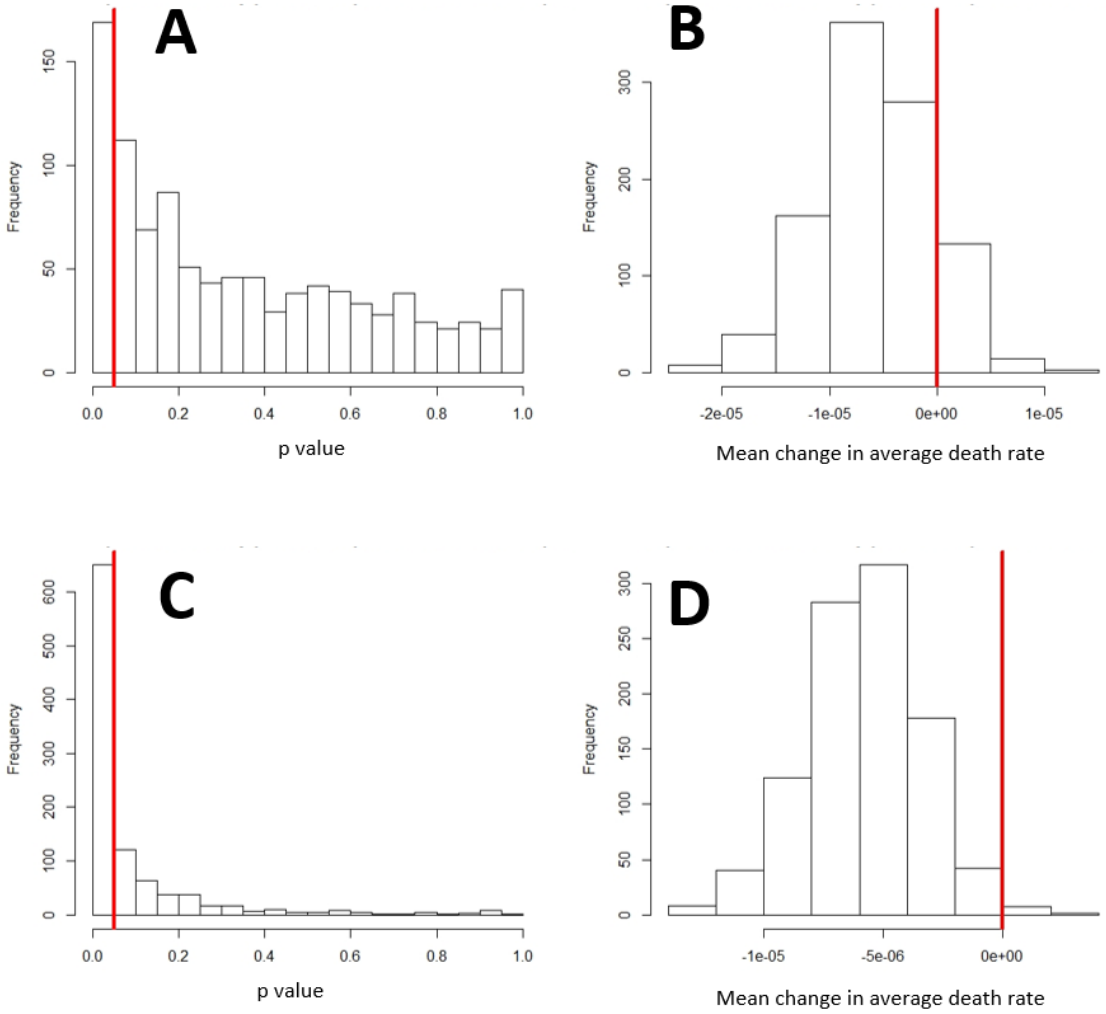

Appendix Figure 6: Subsampling the analysis of the change in average death rate during population crash phases in Model 3. As the analysis of the change in death rate during different population dynamic phases (see Part 2 of the thesis, subheading "Linking population dynamics and death rate evolution"), I subsampled the dataset 1000 times and repeated the analysis on each subsample, resulting in 1000 p-values and 1000 mean changes in death rate. This figure thus shows histograms of the p-values (Appendix Figure 6A & 6C) and of the mean changes in death rate (Appendix Figure 6B & 6D) for 1000 subsamples of subsample size N_{sub}=100 (Appendix Figure 6A & 6B) and for 1000 subsamples of subsample size N_{sub}=500. Histograms for other subsample sizes are available in the data archive. The red line indicates 0.05 in Appendix Figure 5A and 5C, and 0 in Appendix Figure 5B and 5D. The parameter values used in this analysis are s = 0.1, k = 3, Nb = 10 and d = 0.05, histograms for other parameter combinations are available in the data archive. The histograms shown here pertain to the subsampling analysis carried out in Approach 1 (categorizing population dynamic phases by population size), histograms for the second approach are available in the data archive. Overall, most subsamples in the here shown population crash phase do not show changes in death rate that are significantly different from zero – in contrast to most subsample from stationary phases, as shown in Appendix Figure 5.