The distribution of fitness effects during adaptive walks using a simple genetic network

Nicholas L. V. O'Brien[1,2], **Barbara Holland**[3,4], **Jan Engelstädter**[1,2,‡], **and Daniel Ortiz-Barrientos**[1,2,‡]

[1]School of the Environment, The University of Queensland, St Lucia 4072, QLD, Australia
[2]ARC Centre of Excellence for Plant Success in Nature and Agriculture, The University of Queensland, St Lucia 4072, QLD, Australia
[3]School of Natural Sciences, University of Tasmania, Hobart 7001, Tasmania, Australia
[4]ARC Centre of Excellence for Plant Success in Nature and Agriculture, University of Tasmania, Hobart 7001, Tasmania, Australia
[‡]These authors contributed equally to this work.
Correspondence: n.obrien@uq.edu.au and d.ortizbarrientos@uq.edu.au

November 1, 2023

Abstract

The tempo and mode of adaptation is dependent on the availability of beneficial alleles. Genetic interactions arising from gene networks can restrict this availability. However, the extent to which networks affect adaptation remains inconclusive. Current models of evolution consider additive genotype-phenotype relationships while often ignoring the contribution of gene interactions to phenotypic variance. In this study, we present a novel approach to modeling quantitative traits as products of genetic networks via systems of ordinary differential equations. This allows us to mechanistically explore the effects of network structures on adaptation. We present the first step of this approach by studying a simple gene regulatory network, the negative autoregulation motif. Using forward-time genetic simulations, we measure adaptive walks towards a phenotypic optimum in both additive and network models. A key expectation from adaptive walk theory is that the distribution of fitness effects of new beneficial mutations is exponential. We found that although both models strayed from expectation, harboring fewer beneficial mutations than expected, they took a similar number of adaptive steps to reach the optimum. Network model populations were 10% less likely to reach the optimum and took longer to adapt. This was partly due to a complex and largely deleterious distribution of fitness effects from new mutations. This behavior is reminiscent of the cost of complexity, where correlations among traits constrain adaptation. Our results suggest that the interactions emerging from genetic networks can hinder adaptation by generating complex and multimodal distributions of fitness effects.

***Keywords**: population genetics; quantitative genetics; systems biology; adaptive walk; genotype-phenotype map*

Introduction

A lingering question in the study of adaptation concerns the distribution of effect sizes of adaptive alleles that affect trait and fitness variation in natural populations. Research across various species and traits has yielded mixed results. In some cases, natural selection favors alleles with small phenotypic effects (e.g., Visscher *et al.* 2012; Boyle *et al.* 2017; Kosheleva and Desai 2018), such as the alleles contributing to the evolution of body size in mammals (Kemper *et al.* 2012). Other times, however, alleles with larger effects are preferred, such as those found in flower color changes that affect pollinator preferences (e.g., Dorweiler *et al.* 1993; Bradshaw *et al.* 1995; Bomblies and Peichel 2022). Understanding when allelic effects of different sizes are favored is crucial to comprehend the "adaptive walk" – the metaphorical path a population takes through genetic space as it adapts to its environment (Walsh and Lynch 2018b). Each "step" of the walk represents a beneficial allele becoming fixed in the population (i.e. reaching 100% frequency), moving it closer to phenotypic states that maximize fitness (Maynard Smith 1970; Kauffman and Levin 1987). The distribution of fitness effects (DFE) among beneficial alleles is a key feature of adaptation, as it describes the tempo and mode of the adaptive walk in phenotypic space: how many steps there are, how large each step is, and in which direction each step drives the population.

Models of adaptive walks

The theory of adaptive walks is crucial for understanding how organisms evolve to have optimal phenotypes. Much of this theory is built on Fisher's geometric model (Fisher 1930). Accord-

ing to this model, an organism's phenotype is represented as a singular point in a multi-dimensional space. The optimal phenotype is depicted by another point. The space between these two points can be traversed as the population accumulates mutations that gradually shift the population from its original phenotype to the new optimum. Fisher suggested that adaptation under this model is primarily driven by mutations that have minor effects. The rationale behind this idea is quite simple: mutations with small phenotypic effects are more likely to move the population towards the optimum (Fisher 1930). This is because these minor mutations are more likely to be beneficial, thereby increasing their chances of being preserved and spreading in the population through natural selection. Thus, over time, these small steps can lead to significant evolutionary changes.

In Fisher's model, the probability that a random mutation of a given phenotypic size, r, is favorable is given by $1 - \Phi(x)$, where Φ is the cumulative distribution function of a standard normal distribution and x is a mutational effect size, given by $x = r\sqrt{n} \cdot 2z$, where n is the number of traits and z is the distance to the optimum (Fisher 1930). For small values of x, $1 - \Phi(x)$ will be closer to 1, indicating a higher probability of the mutation being favorable. Kimura (1983) pointed out that Fisher's argument was not entirely correct. He argued that mutations with small effects are especially susceptible to being lost by genetic drift. Instead, he proposed that mutations with a moderate effect on an organism's phenotype are less likely to be lost by drift and therefore more likely to contribute to adaptation (Kimura 1983; Charlesworth and Charlesworth 2010).

In this chronicle of ideas, Orr (1998) pointed out that both Fisher's and Kimura's theories do not capture the full breadth of an adaptive walk. Using a previous model of adaptive walks developed by Gillespie (1983, 1984), Orr studied adaptive mutations via the DFE over the course of a complete adaptive journey. One of his critical observations was that as populations drew nearer to their phenotypic optimum, mutations causing significant phenotypic shifts became progressively disadvantageous, primarily due to the increased risk of "overshooting" the optimum. Hence, the Gillespie-Orr model is marked by a pattern of diminishing returns in fitness improvement as organisms tread along their evolutionary paths towards a phenotypic optimum (Gillespie 1983, 1984; Orr 1998, 2003a). This model, often referred to as the Gillespie-Orr model, provides a more nuanced understanding of how organisms adapt and evolve over time. It takes into account not only the size and direction of phenotypic changes caused by mutations but also the changing landscape of fitness effects as populations approach their optimum (Orr 1998, 2003a).

The development of the Gillespie-Orr model marked a shift in theories of adaptation; however, rumblings of adaptive walk models could be heard decades before Gillespie's (1983; 1984; 1991) or Orr's (1998; 2003a) publications. The Gillespie-Orr model builds on the conceptual sequence-space model proposed by Maynard Smith in the early 1970s. In this model, each unique genetic sequence is envisioned as a point in a hypothetical "genotype space", with distances between points representing their genetic differences. Adaptive walks, according to Maynard Smith, could be visualized as sequences of amino acid changes, each constituting an adaptive step in the organism's evolutionary journey (Maynard Smith 1970). Through a series of fixations whereby each increases fitness relative to the last one, natural selection could modify a protein incrementally until the population reaches the optimal protein configuration (Maynard Smith 1970).

Building on Maynard Smith's foundational work, Gillespie integrated principles from extreme-value theory – a statistical branch focused on sampling from extreme tails of distributions – into the genetic theory of adaptation. Gillespie posited that adaptive steps, such as successive amino acid replacements, should be seen as samples from an extreme value distribution (EVD). In essence, any chosen genotype during an adaptive walk could come from the tail of a given distribution, as long as it confers greater fitness than the wild-type configuration (Gillespie 1984). Orr later extended this work using simulations to explore the model's overtones. He showed that adaptive walks generally comprise a few large-effect fixations followed by numerous smaller-effect ones (Orr 1998, 2003a). Orr's model underscores that as a population approaches its optimum, the adaptive walk comprises the largest available effect size, yet still features mutations of increasingly smaller phenotypic effect size, demonstrating the diminishing-returns nature of adaptive walks.

The role of the EVD in understanding the statistical behavior of beneficial alleles during adaptive walks is now foundational in evolutionary genetics. Gillespie's work was instrumental in eliciting this. By assuming that a wild-type genotype already has high fitness, he found that the EVD aptly describes the distribution of beneficial mutations (DBFE) in a population (Gillespie 1984, 1991). When the wild-type has high fitness, it is situated on the right tail of the DFE among all genotypes. Any genotype with superior fitness is then an outlier in the tail of the DFE of all new mutations. This tail is represented by the EVD. The EVD can take various forms, generally categorized into one of three statistical domains: the Gumbel, Weibull, or Fréchet family of distributions. Each domain has its implications for our understanding of the DBFE, reflecting the genetic architectures, selection pressures, and environmental variability that drive the fixation of alleles during adaptation (Fig. 1; Rokyta *et al.* 2008; Joyce *et al.* 2008; Martin and Lenormand 2008).

The different EVD domains represent alternative forms for the DBFE. Gillespie favored the Gumbel domain, arguing that it is most likely to represent empirical DFEs as it captures common distributions such as the normal, gamma, and log-normal (Gillespie 1983, 1984). Under the Gumbel, the DBFE takes a negative exponential form (hereafter referred to as exponential). This implies that during an adaptive walk, the DFE consists of a few large-effect alleles and numerous smaller, incremental improvements (Gillespie 1983). Empirical findings, including mutagenesis studies and spontaneous mutation experiments in viruses, largely support this exponential representation (e.g. Orr 1998; Imhof and Schlötterer 2001; MacLean and Buckling 2009; Kassen and Bataillon 2006; Sanjuán *et al.* 2004; Vale *et al.* 2012; Eyre-Walker and Keightley 2007). However, evidence in favor of alternative domains does exist, suggesting

the field is still open to exploration. The Weibull domain, for instance, is characterized by its lighter tail, suggesting fewer large-effect beneficial mutations compared to the Gumbel domain (Fig. 1, Charras-Garrido and Lezaud 2013). Rokyta *et al.* (2008) observed a Weibull distribution in the ID11 ssDNA phage, hinting at an upper limit on beneficial fitness effects. On the other hand, the Fréchet domain's heavy tail implies more frequent strongly beneficial mutations (Fig. 1, Charras-Garrido and Lezaud 2013). While these mutations are typically rare (Eyre-Walker and Keightley 2007), Schenk *et al.* (2012) observed a Fréchet-like DFE in *Escherichia coli* adapting to antibiotics. This diversity in observations suggests that while Gillespie's Gumbel assumption is broadly accurate, exceptions exist.

Differences in the shape of the EVD between populations can be driven by differing selection pressures, environmental variability, and genetic architectures. For instance, stabilizing selection can generate a Weibull-distributed EVD (Martin and Lenormand 2008). Fluctuating environments, on the other hand, can change the optimal fitness landscape, thereby altering the form of the DFE over time or across different habitats. This violates one of the Gillespie-Orr model's assumptions: that the DFE does not change over time. Empirical studies indicate that a population's DFE tends to shift towards a more "leptokurtic" shape, with heavier tails and a sharper peak, as populations approach their optimal states (Cowperthwaite *et al.* 2005; Eyre-Walker and Keightley 2007). Hence, in a variable environment, the EVD might take a form that allows for more rapid adaptation, with a wider range of beneficial mutations becoming plausible candidates for fixation (i.e. a Fréchet-distributed EVD). The effects of genetic architecture on the shape of the EVD are less well known.

Gene interactions, pleiotropic effects, and linkage disequilibrium are known to create complex distributions of deleterious fitness effects (Eyre-Walker and Keightley 2007). Schenk *et al.* (2012) proposed that pleiotropy could describe the variability in DBFEs between traits after observing a variety of DBFEs that depended on environmental conditions. For instance, consider the fixation of antibiotic resistance alleles in a bacterial population. If these alleles have pleiotropic effects on other traits which reduce fitness, the fixation probability will depend on the presence of the antibiotic in the environment (Schenk *et al.* 2012). A similar effect would be observed if the resistance allele was linked to an allele that lowered fitness in other traits. In summary, there are a large number of factors which can influence the DBFE and these can produce non-Gumbel EVDs.

Deviations from the Gumbel domain in the Gillespie-Orr model have many implications. As observed by Rokyta *et al.* (2008), minor shifts might not drastically alter the predicted exponential relationship of beneficial mutations. Second, since the Gillespie-Orr model predicts a sequence of beneficial fixations leading to adaptation (which requires strong per-locus selection relative to mutation), it can be difficult to reconcile it with alternative models of adaptation, such as polygenic adaptation, where adaptation occurs through small frequency shifts at many loci under weak selection (Visscher and Goddard 2019; Wellenreuther and Hansson 2016).

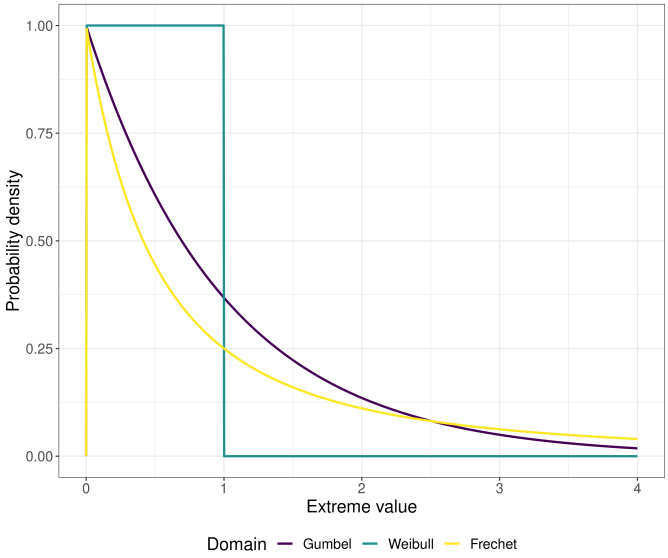

Figure 1 The different domains of attraction for the generalized Pareto distribution (GPD). The GPD, an extreme value distribution, characterizes the behavior of extreme values: random samples from the right tail of a continuous distribution. Extreme values are used to represent beneficial mutations during an adaptive walk. Depending on the shape parameter, κ, of the GPD, it can belong to one of three domains: Gumbel, Weibull, or Fréchet. For $\kappa = 0$, the extreme values align with the Gumbel distribution, which, based on Gillespie's assumptions, is the anticipated distribution for beneficial mutations during an adaptive walk (shown in purple). If $\kappa < 0$, the extreme values take on a Weibull distribution, exhibiting a truncated right tail (shown in green). Decreasing κ shifts the maximum extreme value (at extreme value $= 1$ in the figure) towards zero. In the figure, $\kappa = -1$. Conversely, for $\kappa > 0$, the extreme values follow the Fréchet distribution, characterized by an extended right tail (as depicted in the plot where $\kappa = 1$, shown in yellow).

The foregoing reflection illuminates the complexities inherent in adaptive walk models, revealing that a single-model approach may be insufficient to capture the full scope of evolutionary dynamics. Indeed, the difficulty in reconciling adaptive walk models like the Gillespie-Orr with alternative paradigms such as polygenic adaptation highlight the need for a more nuanced understanding. These complexities are deeply embedded in the genotype-phenotype-fitness (GPW) landscape, a concept that integrates genetic, developmental, and ecological factors to offer a comprehensive view of evolutionary processes. A thorough exploration of the GPW map, therefore, is not just an intellectual exercise but a necessary step in progressing our understanding of how organisms adapt and evolve. The following section introduces the intricacies of the GPW map and how it helps to build the model presented here.

The genotype-phenotype-fitness map

The genotype-phenotype-fitness (GPW) map is composed of two distinct components: the genotype-phenotype (GP) and

the phenotype-fitness (PW) relationships. The GP map describes how developmental and physiological processes translate genotypes into phenotypes, while the PW map describes how ecological and selection regimes create fitness differences between phenotypes. For example, stabilizing selection favors intermediate phenotypes, while directional selection favors extreme phenotypes (Simpson 1944). The PW map, along with the related genotype-fitness (GW) map, has been extensively studied in the fields of quantitative and population genetics (e.g., Wright 1931, 1932; Thornton 2019; Maynard Smith *et al.* 1985; Lande and Shannon 1996). Because the GP map is notoriously challenging to estimate, many quantitative genetics models assume an additive relationship between genotype and phenotype (Walsh and Lynch 2018a; Lande 1976).

Additive GP models have a storied history, beginning with early biometricians' efforts to measure the correlations between the trait values of parents and offspring. These correlations were measured using linear models, with the slope of the model describing the amount of heritable variability in traits of interest (Falconer 1996). However, these models were largely removed from the emerging Mendelian genetics theory, which viewed trait inheritance as a discrete process. In 1918, Fisher published a paper reconciling the continuous distributions of trait values that biometricians measured with the discrete Mendelian understanding of gene action (Fisher 1918). Fisher's "infinitesimal" model supposes that continuous trait distributions can be produced by loci under Mendelian segregation as long as those loci are a) many, and b) have small, additive effects on the phenotype (Fisher 1918). This model has become the basis of modern quantitative genetics, providing a simple GP map that requires no information about the underlying genetic systems that describe the developmental and physiological underpinnings of traits (Falconer 1996; Walsh and Lynch 2018a).

Overwhelming evidence suggests that non-additive gene interactions (epistasis) are ubiquitous in nature (Álvarez-Castro and Carlborg 2007; Hansen 2013; Ang *et al.* 2023) and play a crucial role in adaptation (Draghi *et al.* 2011; Breen *et al.* 2012; Bendixsen *et al.* 2017; Weinreich *et al.* 2006). Additive models, including the infinitesimal, often fail to capture these gene interactions, and when identified, it is unclear how they represent biological gene action (Huang and Mackay 2016). Integrating information about the underlying systems describing trait development enables us to capture genetic interactions affecting fitness at the cost of increased model complexity. Fortunately, with the power of contemporary cluster computing technologies and the wealth of knowledge that has emerged from systems biology (Alon 2019), this is not as difficult a cost to ameliorate as it would have been in Fisher's day. modeling the molecular networks underpinning trait development and expression can help us understand how gene interactions and regulatory processes shape the survival and reproductive success of organisms, and provide a mechanistic view of how variation arises in natural populations. As we delve deeper into the nature of the genotype-phenotype-fitness landscape, it becomes important to re-evaluate our foundational quantitative genetic models in light of our current understanding of the molecular basis of traits. This paper aims to bridge this gap, offering insights into the nuances of the GPW map and its implications for evolutionary genetics.

Modeling molecular networks during trait development

At the foundation of any trait lies a complex network of interacting genes and regulatory elements that respond to environmental cues by producing specific proteins (Doane and Elemento 2017). Systems biologists employ mathematical networks known as gene regulatory networks (GRNs) to model such systems (Alon 2019). As an example, consider a small GRN motif called negative autoregulation (NAR; Fig. 2). This simple motif consists of two genes: gene X activates gene Z (indicated by the pointed arrow in Fig. 2A), while the presence of Z's product limits its further expression (indicated by the flat arrow in Fig. 2A).

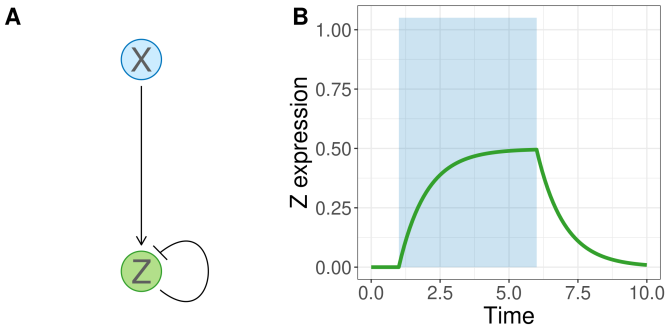

Figure 2 A negative autoregulation (NAR) motif and its expression curve. The NAR motif consists of two genes, X and Z. The expression of X activates Z, and the expression of Z inhibits further Z production (A). This results in a characteristic expression curve (B). Z production begins when X is activated (blue shaded area) and stops when it is inactivated. Gene Z product approaches a steady state concentration and then quickly falls off in the absence of X.

NAR motifs are highly prevalent in biological networks owing to their ability to reduce gene expression variability between cells and to accelerate transcriptional responses to environmental cues (Rosenfeld *et al.* 2002; Nevozhay *et al.* 2009). They can be thought of as self-balancing systems: when a gene is overexpressed, the NAR mechanism triggers to reduce that gene's production, leading to a steady state of gene expression (Alon 2019). To model such networks, systems biologists often use systems of ordinary differential equations (ODEs), with the solutions used to predict gene expression over some time period, such as during cell development. For the NAR motif, the ODE is given by:

$$\frac{dZ}{dt} = \beta_Z \frac{X^{n_{XZ}}}{K_{XZ}^{n_{XZ}} + X^{n_{XZ}}} \frac{K_Z^{n_Z}}{K_Z^{n_Z} + Z^{n_Z}} - \alpha_Z Z \qquad \textbf{(1)}$$

Here, X and Z represent the cellular concentrations of gene products X and Z, n_{XZ} and n_Z are steepness coefficients (Hill coefficients), α_Z is the rate of Z degradation, β_Z is the production rate of Z, and K_{XZ} and K_Z are activation and repression

coefficients, respectively (Alon 2019). X follows a step function and is only expressed within the interval $t \in [t_{start}, t_{stop}]$:

$$X(t) = \begin{cases} 1, & \text{if } t_{start} \leq t \leq t_{stop}, \\ 0, & \text{otherwise} \end{cases} \quad \textbf{(2)}$$

The solution to this ODE generates a nonlinear approach to maximum Z expression, followed by a decline after X expression ceases (Figure 2B). The expression curve can be measured to yield a quantitative trait value. The behavior of the NAR motif introduces such nonlinearities to the GPW map because the effect of a mutation in the production rate parameter β_Z depends on the removal rate α_Z. With a low α_Z, increasing β_Z results in a large phenotypic effect as protein production is not hampered. But with a high α_Z, modulation of β_Z has little impact as proteins are rapidly degraded anyway. This nonlinearity imposes a more rugged, complex fitness landscape, where the effect of an allele depends on other loci. For instance, in an additive model, an α+1 effect allele always increases the phenotype by +1 regardless of context. But with a NAR network, the same allele could produce a variety of phenotypes depending on the genetic background. Consequently, the DFE of new mutations in the network model may be multimodal and more skewed towards deleterious effects than in the additive model.

Traits that measure cellular and biochemical processes (such as our NAR example) are referred to as molecular traits (Claringbould *et al.* 2017). Molecular traits are intermediate traits lying between genotype and phenotype. Instead of relying on an additive model where numerous quantitative trait loci (QTLs) contribute small, additive effects to the trait value, we can utilize a biological network to approximate the underlying mechanisms of trait expression. In the context of quantitative traits, a network-based approach could offer insights into how adaptation is influenced by the genetic architecture of a trait. We propose a novel method which integrates GRNs into existing population and quantitative genetics models to examine how gene networks constrain or facilitate responses to selection on quantitative traits. While there have been previous attempts to model network evolution (e.g., Slatkin 1987; Rice 1998, 2008; François 2014), these have primarily focused on canalization and developmental biology and lack the population and quantitative genetics perspective that our approach includes.

Our model treats the evolution of the network as the result of mutations in "molecular components", the coefficients of the network's ODE/s. This is a novel solution that allows epistatic interactions between loci to emerge as properties of the network. The hierarchical structure of the model drives these emergent interactions: mutations alter the molecular components, which then creates variation in the molecular traits, which then inform quantitative trait expression. By focusing on the underlying molecular components and their interactions, our model offers a more granular and accurate picture of trait evolution, capturing nuances that might be overlooked in more simplified models. Our approach offers a more holistic understanding of the adaptive constraints that non-additive GP maps might impose upon natural populations. By helping to create a theoretical bridge between population genetics, quantitative genetics, and systems biology, our method has the potential to inform the conditions that limit adaptation under complex genetic architectures.

Below we introduce our approach to model GP maps using a simple implementation of the NAR motif (Fig. 2, Eq. 1) to investigate the effects of network-mediated interactions on the distribution of fitness effects during an adaptive walk to a phenotypic optimum. We begin by presenting a mathematical model that can be applied to any network alongside an illustrative implementation using the NAR motif. We use Wright-Fisher simulations to describe the adaptive walks of populations with either an additive or NAR GP map. Subsequently, we examine which extreme value distribution domain (Gumbel, Weibull, or Fréchet) best captures the behavior of the proposed systems model of adaptation. We then identify if the network imposes constraints on adaptation compared to an additive model. Finally, we discuss the contribution of genetic network architectures to the process of adaptation and propose avenues for further exploration using similar systems models.

A hierarchical molecular network model for quantitative traits

To capture the genetic architecture of a trait mediated by a genetic network, our model incorporates a hierarchical structure of components (Fig. 3). The highest level is the phenotypic trait, such as height or body mass, which emerges from the interactions of underlying molecular traits. These molecular traits encompass various cellular mechanisms, including gene expression, methylation, and protein production (Claringbould *et al.* 2017). By combining these molecular traits through a function $P(\mathbf{M})$, where $\mathbf{M}$ is a vector of the relevant molecular traits, we obtain the phenotypic trait value. In the negative autoregulation (NAR) example, we simplify the model by considering total Z expression, a molecular trait, as the phenotypic trait value. However, in more complex scenarios, multiple molecular traits contribute to the phenotypic trait. For instance, the pathogenicity of *Sclerotinia sclerotiorum*, a fungal plant pathogen, is influenced by several molecular traits, including the molecular weights, isoelectric points, and transcriptional regulation of cell-wall-degrading enzymes (Keon and Byrde 1987; Bolton *et al.* 2006). These molecular traits combine via $P(\mathbf{M})$ to determine the phenotypic trait (Fig. 3). Mathematically, the values of these molecular traits are solutions to a system of ordinary differential equations (ODEs). Modeling molecular trait variation requires an additional level of organization to account for the dynamics of the ODEs over developmental/physiological time. These underlying "molecular components" govern the behavior of the molecular traits.

Molecular components capture the variation in transcription and translation that drive the expression of molecular traits. For example, variation in the binding affinity of RNA polymerase or the production rate of transcription factors can impact gene expression dynamics over time. In our NAR example (Eq. 1), the molecular components include n_Z, n_{XZ}, α_Z, β_Z, K_Z, and K_{XZ}. In this study, we focus on the dynamics of α_Z and β_Z for simplicity, fixing the other molecular components at constant values (found in Table 1 – NAR molecular components). α_Z represents the rate at which Z is removed from

the cell. This might reflect, for example, the activity of ubiquitinase in tagging Z for removal and the 26S proteasome in breaking down the protein (Morimoto *et al.* 2016; Schmidtke *et al.* 2006). β_Z represents the production rate of Z, which is influenced by factors such as transcription factor binding affinity, enhancers, silencers, and *trans*-acting regulatory elements (Halfon 2020). To account for variation in molecular components, we adopt a model where loci contribute multiplicatively to the values of molecular components. We refer to these loci as molecular quantitative trait loci (mQTLs). The evolution of these mQTLs can be studied using the Wright-Fisher model, a foundational model in population genetics that describes the stochastic process of allele frequency change in a finite population (Wright 1931; Fisher 1930). We will now present a detailed mathematical framework of the hierarchical molecular network model for quantitative traits, dissecting each tier in the hierarchical structure, as illustrated in Fig. 3). For a comprehensive list of parameters and symbols pivotal to the model, refer to Supplementary Table 1.

Genotype to molecular component

Let $\mathbf{C}$ be a vector of n_C molecular components:

$$\mathbf{C} = (C_1, C_2, ..., C_{n_C}) \in \mathbb{R}^{n_C}$$

In the NAR example, we have:

$$n_C = 2$$
$$\mathbf{C} = (\alpha_Z, \beta_Z)$$

An allele at locus i and chromosome j is associated with a vector A^{ij} of effects on molecular components:

$$A^{ij} = \left(a_1^{ij}, a_2^{ij}, ..., a_{n_C}^{ij}\right)$$

In our NAR example, each locus affects a single molecular component for simplicity, i.e., A^{ij} takes the form $(a_1^{ij}, 0)$ or $(0, a_2^{ij})$.

Given an individual's alleles, the value for each molecular component is calculated by exponentiation of the sum of allelic effects across all mQTLs and chromosomes. For molecular component k:

$$C_k = C_k^{\text{base}} \cdot \exp\left(\sum_{i=1}^{L_Q}\sum_{j=1}^{2} a_k^{ij}\right), \quad \textbf{(3)}$$

where C_k^{base} is the baseline value of C_k and $C_k \in [0, \infty)$. L_Q represents the number of causal loci along the genome. The exponential transformation ensures that the molecular component values are always positive, which is essential as these values represent rates or measurements of concentration that cannot be negative.

Molecular component to molecular trait

After calculating $\mathbf{C}$, we can determine the values of the molecular traits $\mathbf{M}$. We achieve this by constructing a system of ODEs and using a functional, which takes a function as its input and produces a scalar (a single number) as its output, which in this case is a summary statistic of the ODE's solution. ODEs are commonly used to model gene networks due to their balance between efficiency and realism (Schlitt and Brazma 2007; Karlebach and Shamir 2008).

Let the system of ODEs describing the expression of molecular traits $\mathbf{M}$ be:

$$\mathbf{0} = F\left(\mathbf{E_M}(t), \mathbf{E_M}'(t), t\right) \quad \mathbf{E_M}(t) \in \mathbb{R}^{n_M} \quad \textbf{(4)}$$

where n_M is the number of molecular traits, $\mathbf{E_M}(t)$ is a vector of n_M functions describing the number of molecules of gene product/s $\mathbf{E_M}$ produced at time t, and $\mathbf{E_M}'(t)$ is a vector of the first derivatives of those functions.

We can express the ODEs as an initial value problem where given the "initial" or starting value of the function we can determine the function's behavior for subsequent time using the ODE:

$$\mathbf{0} = F\left(\mathbf{E_M}(t), \mathbf{E_M}'(t), t\right)$$
$$\mathbf{E_M}(0) = \mathbf{E_{M0}}$$

Let $\hat{\mathbf{E}}_\mathbf{M}$ be the solution to this initial value problem. Then, the molecular traits can be calculated as:

$$\mathbf{M} = h(\hat{\mathbf{E}}_\mathbf{M}) \quad \textbf{(5)}$$

where $h(\cdot)$ is a functional.

Note that $\mathbf{E_M} \neq \mathbf{M}$. $\mathbf{E_M}$ represents the expression curves of the molecular traits, while $\mathbf{M}$ represents the summary statistics of those curves.

In the NAR example, the functional is represented by a definite integral, which calculates the total amount of Z produced during trait development. The gene product of interest is Z, so $n_M = 1$ and $M \in \mathbb{R}$. The ODEs describing the production of Z over time are given by Eq. 1 and 2. In Eq. 2, $t_{start} = 1$ and $t_{stop} = 6$. Z production is evaluated for $0 \leq t \leq 10$, with $t_{start} \leq t \leq t_{stop}$ representing 50% of the total measured time. To calculate M, we integrate the expression curve:

$$M = h(\hat{E}_M) = \int_0^{t_{\max}} Z(t)dt \quad \textbf{(6)}$$

where $Z(t) = E_M(t)$ and $t_{\max} = 10$.

Molecular trait to complex trait

The complex trait value $P \in \mathbb{R}$ is computed by the function:

$$P = P(\mathbf{M}) \quad \textbf{(7)}$$

In the NAR example, where only one molecular trait is considered, we have:

$$P = M \quad \textbf{(8)}$$

Complex trait to fitness

To map a complex trait to fitness, we utilize a Gaussian fitness function, which allows us to model directional and stabilizing

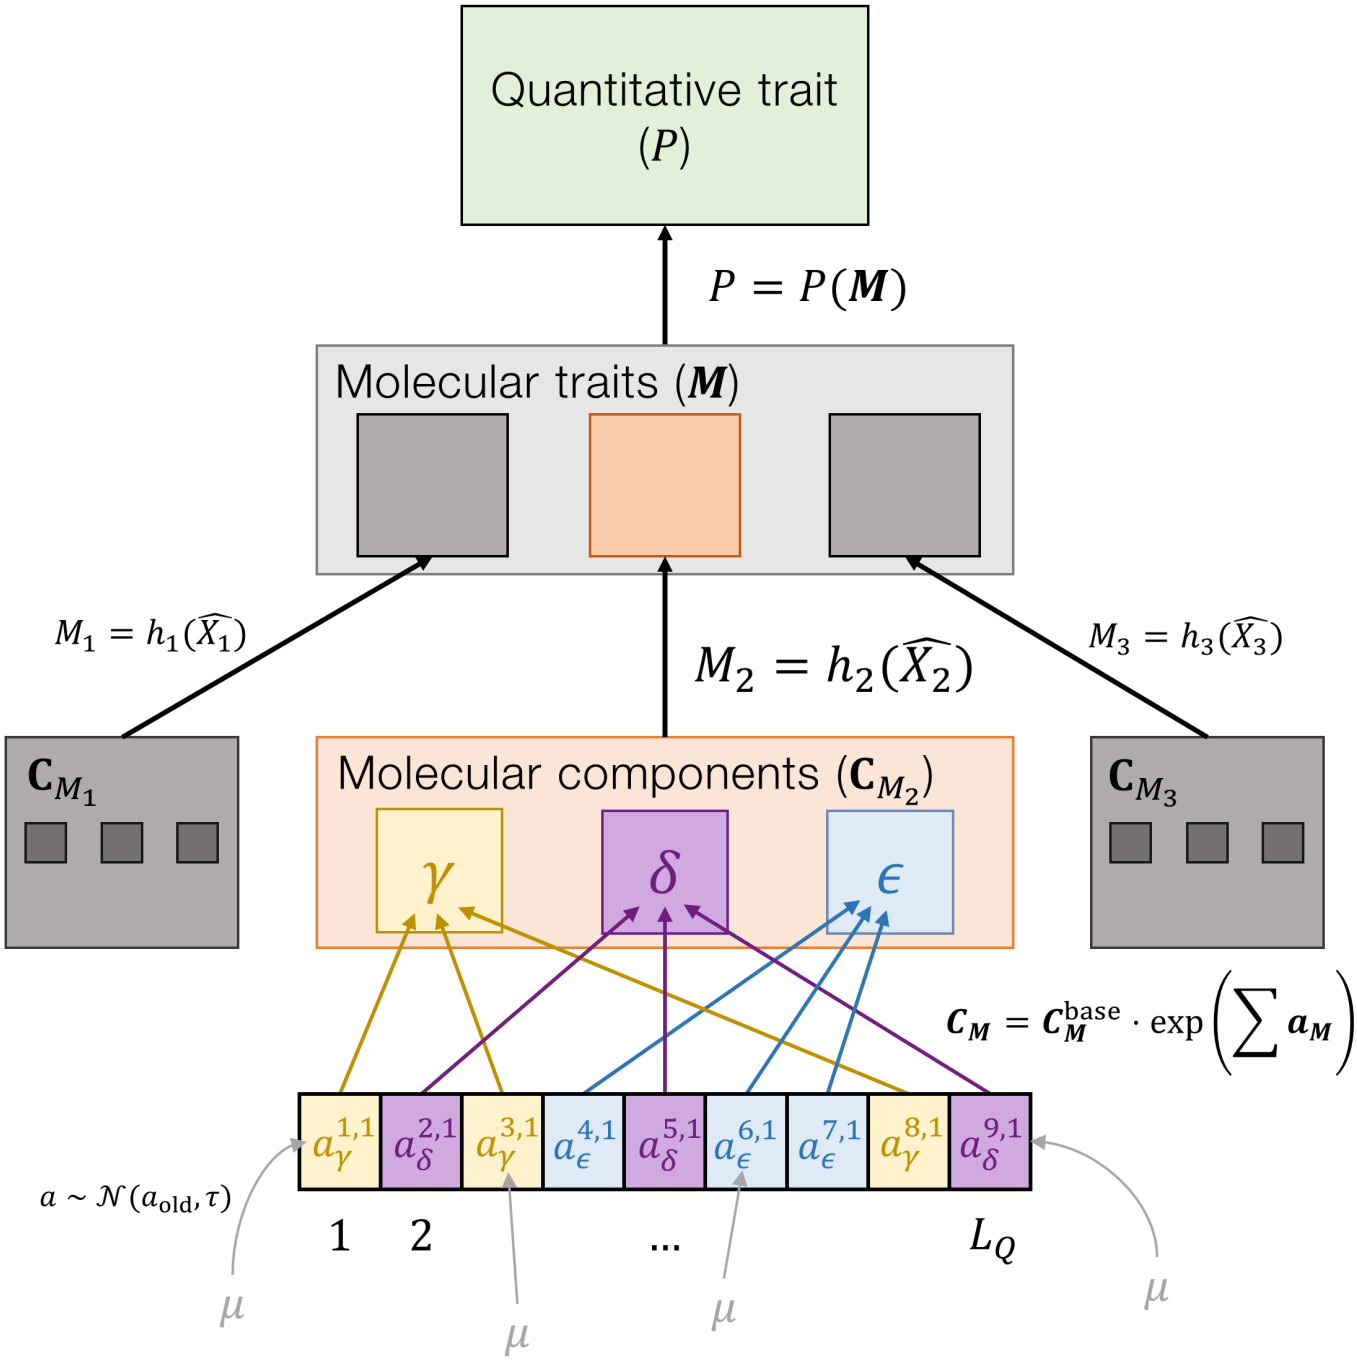

Figure 3 A hierarchical model of a quantitative trait influenced by a genetic network. The network is mathematically modeled using a system of ordinary differential equations (ODEs). Solving the ODEs yield molecular trait values via some measurement function $M = h(\hat{\mathbf{X}})$. The quantitative trait is derived from a combination of molecular traits ($P = P(\mathbf{M})$). To contribute variation to the molecular traits, we consider the coefficients of the ODE, the molecular components, as targets of evolution. L_Q molecular quantitative trait loci (mQTLs) are scattered across the genome, contributing small amounts towards the molecular components in a multiplicative fashion ($\mathbf{C_M} = \mathbf{C_M}^{base} \cdot \exp(\Sigma \mathbf{a_M})$). Mutations (represented by μ) in mQTLs are sampled such that they alter the previous value of the allele ($a_k^{i,j} \sim \mathcal{N}(a_{\text{old}}, \tau)$), where $a_k^{i,j}$ is the allelic effect on molecular component k at locus i on chromosome j. τ is the variance in the distribution of effects on molecular components among new mutations. L_Q represents the number of causal loci along the genome. The colored boxes show an example of how one set of loci affecting one molecular trait (M_2) through its molecular components ($\mathbf{C}_{M_2}$) can percolate to trait variation. Dark grey boxes show the same process for different molecular components and traits, with the genomic loci and molecular components hidden for clarity.

selection. The Gaussian fitness function, originally described by Lande (1976), is defined as follows:

$$w(\mathbf{z}) = \exp\left(\frac{-\mathbf{z}^2}{2\sigma_w^2}\right) \quad \textbf{(9)}$$

In this function, $\mathbf{z}$ represents an individual's phenotypic deviation from the optimal phenotype ($\mathbf{P_O}$), and σ_w is the width of the fitness function. Wider functions represent weaker selection. For our simulations, we used $\sigma_w = \sqrt{10}$, which represents a $\sim 5\%$ drop in fitness when an individual is one phenotypic unit from the optimum. In the NAR example, where only a single trait is considered ($n_P = 1$), we use this Gaussian fitness function.

For a multi-trait space with $n_P > 1$, a multivariate Gaussian fitness function can be employed, as proposed by Zeng (1988):

$$w(\mathbf{z}) = \exp\left(-\frac{1}{2}\mathbf{z}^\top \mathbf{\Sigma}_w^{-1}\mathbf{z}\right) \quad \textbf{(10)}$$

In this equation, $\mathbf{\Sigma}_w$ is a positive definite symmetric matrix that represents the intensities of stabilizing selection on the n_P traits. In the context of the Gaussian fitness function, the diagonal elements of this matrix (the variances) dictate how quickly fitness decreases as a trait deviates from its optimum. The off-diagonal elements (the covariances) describe how the traits interact with each other concerning fitness. If two traits have a positive covariance, it means that increasing one trait will also increase the other, and this combined increase might have a compounded effect on fitness. $\mathbf{z}$ is a vector of size n_P that contains the deviations of an individual's P values, $(P_1, P_2, ..., P_{n_P})$, from the optimal P values, $(P_1^{opt}, P_2^{opt}, ..., P_{n_P}^{opt})$. $\mathbf{z}^\top$ denotes the transpose of the vector $\mathbf{z}$. Transposing a vector turns its rows into columns (or vice versa). This is done so that the resulting product $\mathbf{z}$ is a scalar fitness value $w \in \mathbb{R}$. The equation calculates the fitness $\mathbf{w}$ of an individual based on how far its traits deviate from their optimal values. If an individual's traits are exactly at their optimum, $\mathbf{z}$ would be a vector of zeros, making $\mathbf{w}$ equal to 1 (because $e^0 = 1$). As the traits deviate from their optima, $\mathbf{w}$ decreases, indicating reduced fitness.

It is worth noting that alternative fitness functions, such as those describing directional or disruptive selection (see Felsenstein 1979; Walsh and Lynch 2018d) can also be used. However, they are beyond the scope of this study.

Fitness to the next generation

To model the evolution of the network, we employ a Wright-Fisher (WF) model (Fisher 1923; Wright 1931). We consider a diploid, hermaphroditic population with random mating and non-overlapping generations. Offspring are generated by sampling with replacement from parents, with the sampling probability weighted by their fitness. Individuals possess a pair of homologous chromosomes with n_l loci. Random mutations occur at L_Q molecular quantitative trait loci (mQTLs), which are randomly distributed across the genome as $l_1^Q, l_2^Q, ..., l_{L_Q}^Q$. These mutations alter the values of molecular components multiplicatively, contributing an allelic effect $a \sim \mathcal{N}(a_{\text{old}}, \tau)$, where a_{old} represents the previous allelic effect at the mQTL. The allelic effects across the L_Q loci are combined to produce the molecular component values using Eq. 3. In this study, we consider the case where there is free recombination between loci, although this assumption can be relaxed.

The model we have presented outlines a genetic architecture that connects the genotype, molecular components, molecular traits, complex traits, and fitness. mQTLs contribute variation to molecular components which in turn contribute variation to the molecular traits, which finally lead to variation in the quantitative trait. We have also employed the Gaussian fitness function to map the phenotype to fitness. For scenarios involving multiple traits, the model can be extended by using a multivariate Gaussian fitness function (Eq. 10).

Materials and methods

SLiM implementation

To investigate the effects of our model on the distribution of fitness effects during adaptation, we utilized a custom version of the forward-time, individual-based simulation software SLiM 3.7.1 (Haller and Messer 2019). Our modified SLiM implementation can be accessed openly at `https://github.com/nobrien97/SLiM/releases/tag/AdaptiveWalks2023`. SLiM scripts are available at `https://github.com/nobrien97/NARAdaptiveWalk2023`. A flowchart of the SLiM implementation is given in Supplementary Fig. 1.

SLiM simulates a Wright-Fisher process in a population of N diploid individuals over time. Individuals consist of a pair of genomes, which are split into n_l loci. Mutations were randomly and uniformly sampled across the n_l loci at a rate of $\mu = 9.1528 \times 10^{-6}$ per locus per generation. This mutation rate is based on the average mutation rates observed in *Arabidopsis thaliana* and adjusted to per-locus rates by multiplying it with the average length of a eukaryotic gene (1346bp) (Aston *et al.* 2017; Xu *et al.* 2006). *A. thaliana* was chosen as it is a model plant species with good estimates of mutation rates. The population consisted of $N = 5000$ diploid individuals. This satisfies the condition $2N\mu \ll 1$ (where $2N$ is the total number of genomes in the population; $2N\mu = 0.0915$), which is necessary for the Gillespie-Orr model (Gillespie 1983; Orr 2010). Each molecular component was influenced by a single molecular quantitative trait locus (mQTL), which was randomly assigned to one of the n_l loci at the beginning of each simulation. The mutational effects of mQTLs on molecular components were drawn from a standard normal distribution and the effect added to the previous allele at that locus such that $a_{\text{new}} \sim \mathcal{N}(a_{\text{old}}, 1)$. In each generation, the mutational effects of each individual were combined using Eq. 3 and then used as input for Eq. 1. The ODE was solved by integrating the Ascent numerical solution library (Berry *et al.* 2021) into SLiM and using a custom C++ function to call it. The ODE solver produced molecular trait values for the individuals in the population. Since some individuals shared genotypes (and hence had the same ODE inputs), ODE solutions were cached in memory

to reduce redundant solutions and improve performance. Fitness was calculated by inputting the molecular trait value into Eq. 9. The Lande fitness function was used as it models adaptive walks well: fitness improvements face diminishing returns as the population approaches the phenotypic optimum.

We also considered an additive model. In this scenario, causal loci were sampled and distributed randomly across the genome, similar to the network model. However, in the additive model, the trait value was given by summing allelic effects across QTLs instead of solving an ODE:

$$P_{\text{Additive}} = \sum_{i=1}^{L_Q} \sum_{j=1}^{2} a_k^{ij}, \tag{11}$$

where a_k^{ij} is the allelic effect of locus i on chromosome j.

The additive models consisted of two QTLs to match the mutational target size of the network model (which had 1 locus each for α_Z and β_Z). Both models underwent 50,000 generations of burn-in. The burn-in ensured that populations had high fitness just prior to the optimum shift. During burn-in, populations adapted to a phenotypic optimum at $P_O = 1$. The phenotypic optimum was then instantaneously shifted to $P_O = 2$, and the population was monitored for 10,000 generations of adaptation. We measured the phenotypic means and allelic effects of segregating and fixed mutations every 50 generations. We replicated each model 2,880 times with 32-bit integer seeds sampled from a uniform distribution in R 4.3.1 (R Core Team 2023). Detailed simulation parameters can be found in Table 1. Simulations were executed on the National Computational Infrastructure's Gadi HPC system.

Of the simulation parameters, the most important were the fitness function width, σ_w, and the optimum shift amount, O_{shift}. σ_w and O_{shift} are important because the Gillespie-Orr model assumes that the wild-type has relatively high fitness, so that in the total space of genotypes, it exists on the right-tail of the distribution (Orr 1998). Hence, the fitness function and optimum shift needs to be chosen so that at the optimum shift (generation 50,000) fitness is still relatively high. We chose σ_w and O_{shift} so that this assumption is met: at the optimum shift, individuals perfectly adapted to the burn-in optimum suffered a $\sim 5\%$ drop in fitness.

Fitness effect calculation

To assess the phenotypic and fitness effects of mutations in the network models, we performed single gene knockouts and measured the difference between the "wild-type" and knockout phenotypes and fitness. A flowchart is given in Supplementary Fig. 2A. For the knockout simulation, we recalculated the molecular component values (α_Z and β_Z) by subtracting the effect of a given allele. Subsequently, we recalculated the trait value for the knocked-out genotype, and the effect size was determined as the difference between the phenotype with the allele knocked out (P_k) and the original trait value:

$$\Delta P = P - P_k \tag{12}$$

A similar approach was applied to calculate the selection coefficient, where we subtracted the fitness of the individual with the mutation knocked out from the fitness of a "wild-type" individual:

$$s = w - w_k \tag{13}$$

where w_k represents the fitness of an individual with the allele knocked out. Both the phenotype and fitness recalculations were performed using a custom tool that implemented the Ascent C++ library (Berry *et al.* 2021). This library was also used in the SLiM simulations, keeping the phenotype calculation consistent. The source code for our tool can be found at `https://github.com/nobrien97/odeLandscape`.

For the additive models, the phenotypic and fitness effect sizes were calculated in a similar manner, using Eq. 12 and Eq. 13. However, the phenotype was computed as the sum of additive effects instead of using the network ODE solution. In both the additive and network models, we employed the Lande fitness function (Eq. 9) to recalculate fitness for the phenotype with the allele knocked out.

In silico mutant screen

To assess the distribution of mutational effects, we conducted an *in silico* mutant screen experiment. A flowchart is given in Supplementary Fig. 2B. We sampled 1,000 mutations from a standard normal distribution and added them to the molecular component values/quantitative trait values of individuals at each step of the adaptive walk. This was done for both network and additive models. We chose a standard normal distribution as it matched the distribution of new mutations that populations faced in the simulations. For the network models, mutations were applied separately to α_Z and β_Z, while in the additive models, the mutation effect was added to the quantitative trait value. We then recalculated the phenotypic and fitness effects of the sampled mutations using the methods described above. The source code for this experiment can be found at `https://github.com/nobrien97/NARAdaptiveWalk2023`.

Expected waiting times to beneficial mutation

To estimate the expected waiting time to a beneficial mutation between models, we first calculated waiting times, t_{wait} for a particular replicate at a given adaptive step using the equation

$$t_{\text{wait}} = \frac{1}{2N\mu L_Q p_{s>0}} \tag{14}$$

where N is the population size, μ is the per-locus, per-generation mutation rate, L_Q is the number of causal loci, and $p_{s>0}$ is the probability that a new mutation is beneficial, which is given by

$$p_{s>0} \approx n_{s>0}^{ISMS} / n^{ISMS}$$

where n^{ISMS} is the number of mutations generated in the *in silico* mutation screen experiment and $n_{s>0}^{ISMS}$ is the number of mutation screen alleles with beneficial effects on fitness. $p_{s>0}$ was compared between additive and NAR simulations using a linear model:

$$
\begin{aligned}
p_{s>0} = &\alpha + \beta_1(\text{Model}_{\text{NAR}}) + \beta_2(\text{Adaptive step}_1) + \\
&\beta_3(\text{Adaptive step}_2) + \beta_4(\text{Adaptive step}_{\geq 3}) + \\
&\beta_5(\text{Model}_{\text{NAR}} \times \text{Adaptive step}_1) + \\
&\beta_6(\text{Model}_{\text{NAR}} \times \text{Adaptive step}_2) + \\
&\beta_7(\text{Model}_{\text{NAR}} \times \text{Adaptive step}_{\geq 3}) + \epsilon
\end{aligned}
$$

Note that the α and β coefficients here are regression coefficients and do not refer to α_Z and β_Z. To estimate marginal means and contrasts, we used the R package `emmeans` 1.8.5 in R 4.3.1 (Lenth 2023; R Core Team 2023). Contrast p-values were adjusted for multiple comparisons with the Šidák correction. Waiting times were calculated for each adaptive step and model. Comparisons between models were done using a bootstrap analysis. We sampled 100,000 different NAR-Additive model pairs across all adaptive steps, calculating the difference between their waiting times. We then calculated the means and 95% CIs for each model's waiting times and the difference between models (Δt_{wait}). We tested the statistical significance of the mean difference with a paired t-test with the null hypothesis $H_0 : \Delta t_{\text{wait}} = 0$. The source code for this analysis can be found at `https://github.com/nobrien97/NARAdaptiveWalk2023`. A flowchart describing this methodology is shown in Supplementary Fig. 3.

Strong selection weak mutation (SSWM) checks

To ensure we were in the SSWM domain assumed by adaptive walk theory, we measured population heterozygosity and the effects of segregating alleles on trait variation. To measure heterozygosity, H_O, we used the equation

$$H_O^i = N_{het}^i / N \tag{15}$$

where N_{het} is the number of individuals heterozygous at locus i, and N is the total population size. To measure the effects of segregating alleles on the trait, we measured the ratio between the phenotype created by only fixations and the population mean phenotype,

$$r_{fs} = P_{Fixed} / \bar{P} \tag{16}$$

where P_{Fixed} is the phenotype due to only fixed alleles and $\bar{P}$ is the mean population phenotype.

Fitness landscapes

To estimate the fitness landscape of the network system, we generated 160,000 combinations of α_Z and β_Z, sampling α_Z and β_Z values between 0 and 3. We then calculated the phenotype and fitness for each combination using the ODE landscaper tool we used to calculate fitness effects. We plotted the resulting fitnesses against α_Z and β_Z using `ggplot2` 3.4.2 (Wickham 2016) in R 4.3.1 (R Core Team 2023). A similar method was used to estimate the fitness landscape of the ratio β_Z/α_Z as α_Z increased. For this analysis, we generated another 160,000 combinations of α_Z and β_Z/α_Z ratios. Both α_Z and β_Z/α_Z values were sampled between 0 and 3.

DFE analysis

To compare the distributions of beneficial mutational effects in the additive and network models, we used a likelihood ratio test to fit a generalized Pareto distribution (GPD) to the observed distribution of beneficial fitness effects (DBFE) obtained from the mutant screen (Beisel *et al.* 2007). A flowchart describing this approach is provided (Supplementary Fig. 4). We chose to use the GPD rather than alternative extreme value distributions because it is well suited for analysing the distribution of extreme events past a threshold. In our case, the threshold is the high-fitness wild-type and the extreme events are beneficial mutations. The GPD is characterized by two parameters: κ (shape) and τ (scale). The parameter κ determines the tail behavior of the distribution, indicating which domain of attraction the underlying distribution belongs to. Under the Gillespie-Orr model, it is assumed that the DFE follows a Gumbel distribution, and the tail (i.e. the DBFE) is exponentially distributed (Orr 2003a; Gillespie 1984). The method proposed by Beisel *et al.* (2007) uses a likelihood ratio test to evaluate if $\kappa = 0$, corresponding to the Gumbel case. If $\kappa > 0$, the DFE belongs to the Fréchet domain, and if $\kappa < 0$, it belongs to the Weibull domain (Fisher and Tippett 1928; Beisel *et al.* 2007).
We used two different sampling approaches to fit our data to GPDs (Supplementary Fig. 4). In the first method, we pooled mutant screen alleles across all replicates to create a pooled DFE of beneficial mutations (pDBFE) for additive and network models. We sampled 1,000 mutations from the pDBFE and fit the GPD to the resulting sampling distribution. We repeated this process 10,000 times, yielding a distribution of estimates of κ. In the second method, we instead sampled up to 100 mutations from each simulation's mutant screen DFE and fit the GPD to that sample (i.e. there was no pooling between replicates in this method). This was done per adaptive step. Since not all simulation-step pairs had 100 beneficial alleles, some analyses were conducted with fewer than 100 sampled mutations. Simulation-step pairs with fewer than 10 beneficial alleles were discarded as they could not be reliably used to estimate the GPD shape (Beisel *et al.* 2007). From these GPD fits, we obtained a distribution of κ estimates, one for each simulation-step pair. In summary, both analyses fit GPDs to the DBFE and estimate the shape and scale parameters of the distribution. However, the pooled approach shows the average shape of the DBFE across many adaptive walks, whilst the per-simulation approach shows the variability in the shape of the DFE during different adaptive walks. In both the pooled and non-pooled approach, p-values from the likelihood ratio tests were combined across replicates using Fisher's method (Fisher 1936). Both approaches fit the GPD using the R package `GenSA` 1.1.8 (Xiang *et al.* 2013) in R 4.3.1 (R Core Team 2023), with code modified from Lebeuf-Taylor *et al.* (2019).

Results

We first examined the phenotypic response to selection during adaptive walks in populations with either an additive or a network (NAR motif) GP map. We defined an adapted population as one with a mean phenotype within 10% of the optimum ($1.9 < \bar{P} < 2.1$). We chose 10% as the threshold as it only

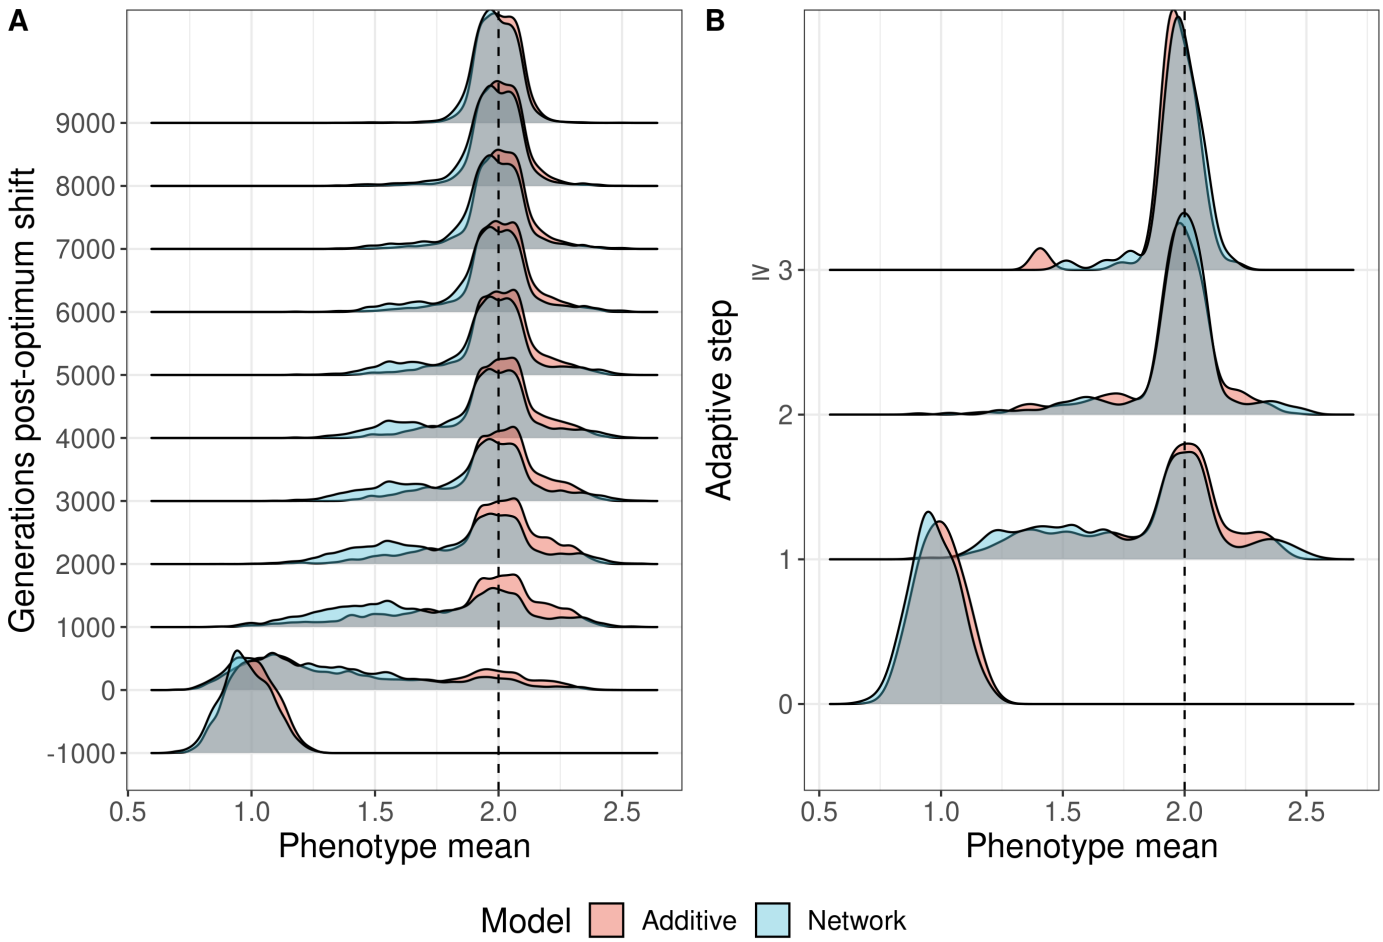

Figure 4 Phenotypic evolution over 10,000 generations of adaptation across 2,880 replicates of additive (red) and network (blue) models. (A) The distribution of mean phenotypes among replicate populations during the adaptive period. The -1000 timepoint represents 1,000 generations before the optimum shifted. (B) The mean phenotype at each adaptive step (each fixation along the adaptive walk). The last step is pooled across all steps greater than or equal to step three in the walk. The dotted line in A and B represents the post-shift optimum. The grey area shows the overlap between additive and network distributions.

decreased fitness by 0.04% relative to a 5% threshold, but increased the sample size of adapted populations by 42% (2636 adapted populations vs. 1529). By the end of the simulation, 50.5% of the additive simulations and 41% of the network simulations had adapted. A 2×2 Fisher's exact test revealed a highly significant difference in the probability of adaptation between the two models ($N = 5760, p < 0.0001$). For the remainder of the results, we consider only the adapted populations.

Among the adapted populations, there were differences in the rate of trait evolution between the models (Fig. 4). Network model populations took longer to reach the optimum on average (Fig. 4A), although the changes in phenotype across adaptive steps were similar between the models (Fig. 4B). There was no significant difference between the models in the number of adaptive steps taken (2×2 Fisher's exact test: $N = 3022, p = 0.590$). Of the populations that adapted, 66.3% did so in a single step, 28% in two steps, and the remaining 5.7% adapted in three or more steps. On average, the number of steps taken was 1.4 ± 0.026 (95% CI). This is in-line with results from Orr (2003b), who found a mean walk of ≈ 1.72 steps, assuming that the genotype was comprised of many loci and that the starting genotype of the walk was randomly sampled.

When considering the timing of adaptive steps among all populations, there was no significant difference between the models. However, among populations that had not yet adapted after one step (the left tail in Fig. 4), the median step in a network model occurred 900 (95% CI [150, 1600]) generations later than in the additive models (Supplementary Fig. 5; Wilcoxon rank sum test: $W = 4749, p = 0.018$). To investigate the underlying causes of the differences in adaptation success and rate between models, we examined the distribution of fitness effects (DFE) among fixations during the adaptive walk.

Fixations had similar effect sizes in both models and became smaller over the course of the adaptive walk (Fig. 5). We found that the mode of the distribution of beneficial fixations was offset from 0 (Supplementary Fig. 6), supporting a gamma distribution of fixations rather than an exponential distribution. We fit gamma distributions to the DFE among fixations in both the

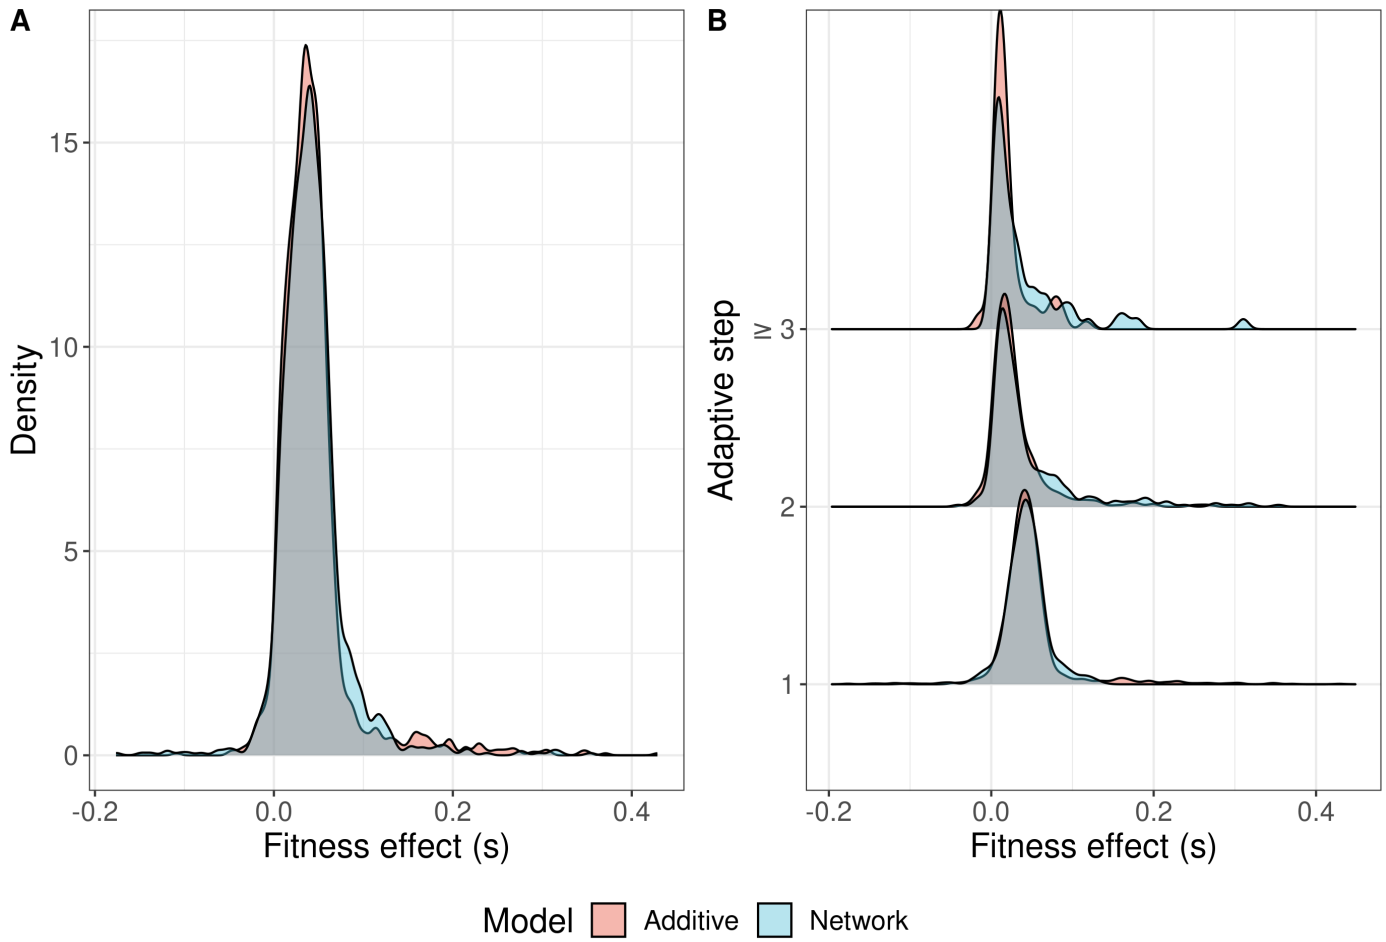

Figure 5 Fitness effects among fixations during adaptive walks to a phenotypic optimum across 2,880 replicates of additive (red) and network (blue) models. The grey area shows overlap between the additive and network distributions of effects. (A) The overall distribution of all fixations across the entire walk. (B) The distribution of fixations at each adaptive step.

additive and network models using `fitdistrplus` 1.1.8 in R 4.3.1 (R Core Team 2023; Delignette-Muller and Dutang 2015). The deleterious fixations shown in Fig. 5 were excluded from the dataset when fitting the gamma distribution. These deleterious fixations were neutral or adaptive prior to the shift in the optimum and typically reached high frequencies before the adaptive walk commenced (Supplementary Fig. 7). Therefore, these fixations can be attributed to the action of genetic drift.

We compared the shape and rate parameters of the gamma fits between the models. The network model populations exhibited slightly fewer small effect fixations and had a shorter tail compared to the additive models (NAR: Shape $= 2.032 \pm 0.147$ 95% CI, Rate $= 44.392 \pm 3.651$ 95% CI; Additive: Shape $= 1.722 \pm 0.108$ 95% CI, Rate $= 37.588 \pm 2.730$ 95% CI). However, the DFE among beneficial fixations did not differ strongly between the models. Hence, differences between models in the DFE of beneficial fixations did not explain differences in the rate of trait evolution. To investigate this further, we needed to determine the availability of beneficial mutations to the populations. We did this by conducting an *in silico* mutant screen experiment, examining the DFE among new mutations at each adaptive step.

The *in silico* mutation screen experiment revealed that the DFE among new mutations differed dramatically between models (Fig. 6A). Both were strongly negatively skewed; however, mutations in NAR populations were overall more deleterious than those in the additive map, with a longer and more pronounced tail of strongly deleterious mutations. In addition, the DFE in NAR populations was bimodal (Fig. 6A). The first mode was at $s \approx 0$, similar to the additive DFE mode (NAR: $s = -0.0003$; Additive: $s = -0.0004$). The second mode was at $s = -0.113$. Among beneficial mutations, the distributions were more similar.

Before we analyzed the DBFE, we first checked if we were in the strong selection weak mutation (SSWM) regime required for the Gillespie-Orr predictions to apply. Under SSWM, each beneficial mutation should fix before another arises. Hence, heterozygosity at QTLs/mQTLs should be close to zero most of the time. We found observed heterozygosity (H_O) at the QTLs/mQTLs was low on average, peaking just after the optimum shift at $11.39 \pm 0.423\%$ and $9.637 \pm 0.400\%$ in additive and NAR models, respectively (Supplementary Fig. 8). The

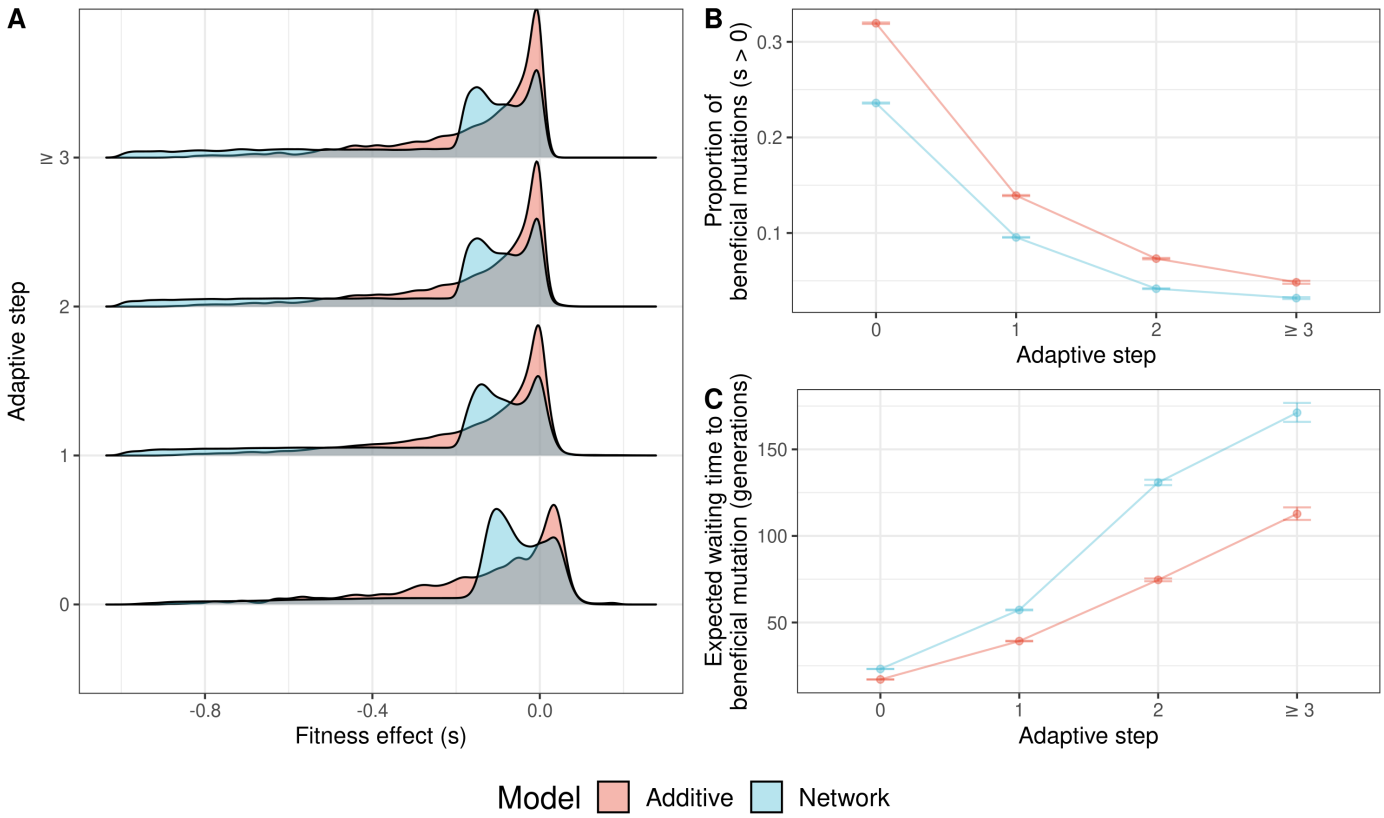

Figure 6 Fitness effects among possible mutations and the consequences for adaptation in 2,880 replicates of additive (red) and network (blue) models. (A) The distribution of fitness effects among 1,000 randomly sampled mutations in additive and network models. The grey area represents overlap between additive and network models. (B) The proportion of mutations that are beneficial over the adaptive walk. (C) The change in beneficial mutation probability corresponds to the expected waiting time for a beneficial mutation.

grand mean H_O across all time points was $8.368 \pm 0.028\%$ and $7.841 \pm 0.027\%$ for additive and NAR models. Although H_O was not zero, the majority of variation arising was not beneficial/likely to fix (Fig. 6B), meaning that the chance that two beneficial mutations could co-segregate was lower than H_O suggests.

To further test the SSWM assumption, we also measured the ratio between phenotypes generated by only fixations (i.e. the phenotype in absence of all segregating variants) and the mean population phenotype (Eq. 16. This ratio describes how much trait variation was contributed by segregating variants compared to fixations. When $r_{FS} = 1$, segregating variants contribute no trait variation. When $r_{FS} > 1$, segregating variants decrease the trait value on average compared to the phenotype due to fixations, and vice versa when $r_{FS} < 1$. We found that most populations had no variation contributed by segregating variants, with relatively few exceptions (Supplementary Fig. 9). The average r_{FS} was 1.019 ± 0.007 (95% CI) and 1.024 ± 0.011 (95% CI) for additive and NAR populations respectively. From these figures, it seems that the populations are in the SSWM domain, enabling us to explore the extreme value theory expectations set out by Gillespie (1983, 1984) and Orr (1998).

We fit a generalized Pareto distribution (GPD) to the DBFEs generated by the mutant screen experiment using methods developed by Beisel *et al.* (2007). This was done separately for additive and network models. The shape parameter of the GPD (κ) describes what domain the DFE belongs to: $\kappa = 0$ specifies a Gumbel domain, $\kappa < 0$ specifies a Weibull domain, and $\kappa > 0$ specifies a Fréchet domain. We fit this in two ways: first, by pooling all replicates and sampling from this pooled distribution to estimate the GPD parameters; and second, by sampling each replicate at each of its adaptive steps separately and estimating the GPD for each estimate (for more information, see the Methods section). Under the pooled approach, we found evidence for a minor deviation from an exponential distribution towards the Weibull domain in both the additive and network models (Network: LR = 0.994 ± 0.00005 (95% CI), $p < 0.0001$, $\hat{\kappa} = -0.130 \pm 0.0007$; Additive: LR = 0.996 ± 0.00004 (95% CI), $p < 0.0001$, $\hat{\kappa} = -0.118 \pm 0.0005$). However, this deviation is small enough that the adaptive walk should behave as if its DFE belonged to the Gumbel domain (Joyce *et al.* 2008).

Under the per-simulation method, we saw a different result. The DBFE in both NAR and additive populations were strongly Weibull distributed. In the NAR model, κ tended to decrease over the adaptive walk. However, the additive model κ estimate was relatively stable at around $\kappa \approx -2$, slightly increasing over the adaptive walk (Supplementary Table 2; Supplementary Fig. 11). In addition, these non-pooled fits were considerably

more variable than under the pooled approach (Supplementary Fig. 10), particularly at early adaptive steps. Under the per-simulation/adaptive-step approach, across all adaptive steps, at least 99.8% of additive replicates and at least 99.7% of NAR replicates had DFEs of beneficial mutations that fit the Weibull domain. Under the pooling method, all replicates in both models were approximately Gumbel-distributed (i.e. $\kappa \approx 0$).

Despite sharing similar DFEs, beneficial mutations were on average less common in network models than in additive models ($F_{7,5171} = 1239, p < 0.001, R^2 = 0.626$; Fig. 6B). The difference between models diminished with adaptive steps (as populations neared the optimum). Immediately after the optimum shift, mutations in NAR populations were $8.35 \pm 0.006\%$ (95% CI; $t_{5171} = 25.766, p < 0.0001$) less likely to be beneficial. At the first adaptive step, this difference decreased to $4.37 \pm 0.006\%$ (95% CI; $t_{5171} = 13.494, p < 0.0001$). At the second adaptive step, the NAR models were $3.15 \pm 0.011\%$ (95% CI; $t_{5171} = 5.613, p < 0.0001$) less likely to produce a beneficial mutation. At further adaptive steps, there was no significant difference between models ($t_{5171} = 1.088, p = 0.726$). We measured the effect of the lower beneficial mutation rate on adaptation by estimating the waiting time for a new beneficial mutation to arise. We found that the mean waiting time for a beneficial mutation in NAR populations was 98.555 ± 5.8155 (95% CI) generations longer than in additive populations ($t_{99999} = 33.315, p < 0.0001$, Additive mean waiting time: 136.382 ± 2.131 (95% CI) generations; NAR mean waiting time: 234.937 ± 5.414 generations; Fig. 6C). Given that the DFE differed between models, we wanted to explore the causes underpinning this. To do so, we investigated the fitness landscape of α_Z and β_Z.

We found a pronounced fitness ridge on the NAR fitness landscape, surrounded by valleys of low fitness (Fig. 7A). This ridge lay diagonally across α_Z and β_Z space, suggesting that fitness depended on the ratio of the molecular components. This relates to the systems biology of the ODE: β_Z/α_Z is the level of steady state Z expression under simple gene regulation (i.e. without any negative feedback or induction) (Alon 2019). Plotting this ratio over α_Z confirmed this: β_Z/α_Z was the arbiter of fitness as long as $\alpha_Z \gtrapprox 0.5$ (Fig. 7B). We investigated how the β_Z/α_Z ratio contributed to the shape of the GP map, finding a nonlinear relationship between β_Z/α_Z and phenotype (Fig. 7C). Around the optimum, increasing β_Z/α_Z quickly increases the trait value, however with larger β_Z/α_Z increases, the trait value changes less (Fig. 7C). We approximated the optimum β_Z/α_Z ratio ($\theta_{\beta_Z/\alpha_Z}$), by substituting our fixed molecular components ($K_{XZ} = 1, K_Z = 1, n_Z = n_{XZ} = 8$) into Eq. 1. We considered the case where $X = 1$, simplifying to

$$\frac{dZ}{dt} = -Z\alpha_Z + \frac{\beta_Z}{2+2Z^8}$$

The equilibrium fulfills the equation

$$\frac{\beta_Z}{2\alpha_Z} = Z(1+Z^8)$$

If we assume a model of simple gene regulation,

$$\theta = \frac{\beta_Z}{\alpha_Z}$$

$$\frac{\theta}{2} = Z(1+Z^8)$$

where θ is the steady state of gene expression under simple gene regulation (Alon 2019). For small values of Z (where $Z < 1$), this is well approximated by

$$Z = \frac{\theta}{2}$$

Using this, we approximate the solution of the ODE by the equilibrium whenever $X = 1$ and by 0 when $X = 0$. Hence, the solution can be given by

$$M = \int_{t_{start}}^{t_{stop}} = \int_1^6 \frac{\theta}{2} dt = \frac{5\theta}{2}$$

Hence, when the optimum is at $P = M = 2$, the optimum steady state is

$$\theta_{\beta_Z/\alpha_Z} = \frac{2M}{5} = 0.8$$

This assumes that the increase in Z after X is activated is rapid, and likewise for the decrease in Z when X is deactivated (i.e. in Fig. 2B, the green curve should be close to rectangular around X activation). For our n_Z and n_{XZ}, this assumption is well met, and the approximation predicts the observed $\theta_{\beta_Z/\alpha_Z}$ in Fig. 7B.

To explore how the nonlinearity of the GP map might have influenced how populations navigated molecular component space, we examined the ratio of absolute additive effect sizes on molecular components among populations that adapted by two steps or more, denoted by $\phi_{\alpha_Z\beta_Z}$ (Fig. 7D). Mathematically,

$$\phi_{\alpha_Z\beta_Z} = \sum_{i=1}^{L_\alpha} |2a^i_{\alpha_Z}| - \sum_{j=1}^{L_\beta} |2a^j_{\beta_Z}|,$$

where a is the allelic effect of locus i or j on α_Z or β_Z, L_α is the number of loci contributing to α_Z, and L_β is the number of loci contributing to β_Z. Allelic effects were multiplied by two since we simulated a diploid population. When $\phi_{\alpha_Z\beta_Z} > 0$, fixations have larger effects on the α_Z axis than the β_Z axis, and vice versa for when $\phi_{\alpha_Z\beta_Z} < 0$. When $\phi_{\alpha_Z\beta_Z} = 0$, there are equal contributions from mutations affecting α_Z and β_Z. On average, $\phi_{\alpha_Z\beta_Z} = 0.164 \pm 0.150$, suggesting that α_Z mutations contributed most to adaptation, albeit by a small amount. Of the populations considered for this analysis, 29.5% adapted with only α_Z mutations, 25% with only β_Z mutations, and the remaining 45.5% used a combination of both to reach the optimum.

Together, these results indicate that the network model harbors a complex, multimodal distribution of fitness effects, partly reflecting the nonlinear genotype-phenotype map imposed by the network. For our choice of phenotypic optimum, this distribution led to network models adapting more slowly on average than additive populations.

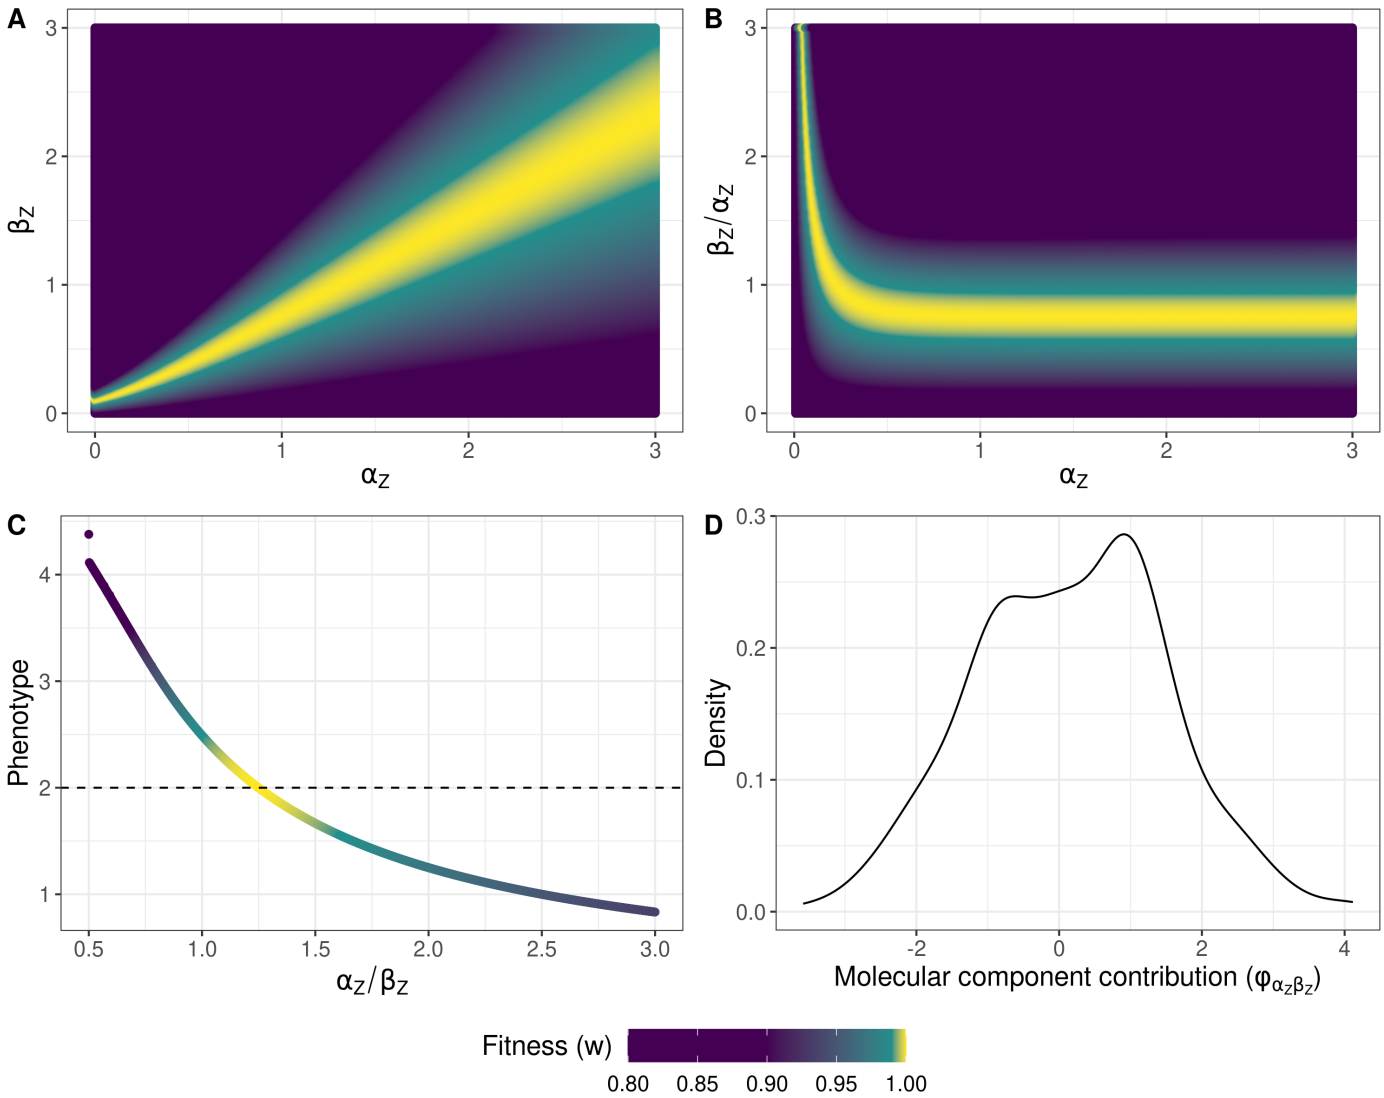

Figure 7 The fitness and phenotype landscapes of the network model. (A) The fitness landscape with regards to the molecular components of the network, α_Z and β_Z. The landscape consists of a high-fitness ridge (yellow) surrounded by low-fitness valleys (purple). (B) The ratio of β_Z/α_Z, the steady state of the network, is what matters for adaptation. As α_Z increases relative to β_Z/α_Z, fitness becomes constant over α_Z. (C) Fitness is a nonlinear function of the network steady state, β_Z/α_Z. An optimum ratio at $\theta_{\beta_Z/\alpha_Z} = 0.8$ maximizes fitness. The dashed black line shows the phenotypic optimum. (D) The contributions of mutations in each molecular component to the adaptive walk. This is taken for populations with at least 2 steps in the walk. $\phi_{\alpha_Z\beta_Z}$ shows the difference in sums of absolute allelic effects among fixations in α_Z and β_Z across an adaptive walk. Values greater than 0 represent larger changes in α_Z than β_Z, and vice versa for values smaller than 0. At 0, both α_Z and β_Z changed by the same amount during the adaptive walk.

Discussion

Overview

In this study, we explored how a simple gene network might influence trait evolution by natural selection. Models of adaptation have traditionally relied on direct genotype-phenotype relationships (Falconer 1996; Kopp and Hermisson 2009; Matuszewski *et al.* 2015), overlooking the intricate gene interaction networks that drive trait expression at different stages of development (but see Rice 1998; Slatkin 1987; Wagner 1996). Understanding how genetic networks influence adaptation is an under-explored area of research. Hence, we developed a novel approach to modeling the evolution of quantitative traits through the lens of genetic networks. As a first step in our research program, we illustrated our model by studying the evolution of the negative autoregulation (NAR) motif within a Wright-Fisher model of evolution. We modified the software SLiM to solve differential equations, allowing us to capture the dynamics of adaptive walks in network model populations and compare them to those of additive models.

We observed that both network and additive models exhibit similarities in key aspects of their adaptive walks, specifically the

number of adaptive steps and the distribution of fitness effects (DFE) among both fixations and beneficial mutations (DBFE) (Fig. 4, 5). However, the characterization of the DBFE was sensitive to the estimation method employed. Utilizing a pooling approach under the assumptions of the Central Limit Theorem (CLT) leads to a homogenization of the shape parameter κ, nudging the distribution towards a Gumbel distribution, and thus suggesting exponential spacing between adjacently-ranked beneficial alleles (Joyce *et al.* 2008). The averaging effect of pooled samples comes with its own problems as it can mask underlying Fréchet or Weibull behaviors in κ, effectively reducing our ability to discriminate between different extreme value distributions (EVDs).

In contrast, using a non-pooled approach preserves the contributions of each replicate's unique combination of mutations to the shape parameter, improving the detection of non-Gumbel behavior. However, this is not always possible in experimental studies where the number of beneficial mutations obtained per round of adaptation is small. Under these conditions, pooling between replicates is required to reduce the risk of incorrectly rejecting a Gumbel distribution (Joyce *et al.* 2008; Beisel *et al.* 2007). In general, both pooling and non-pooling approaches come with their own sets of limitations and advantages, making it crucial to consider the trade-offs involved, particularly when the objective is to reliably identify Fréchet or Weibull regimes within the data. Given these considerations, we next discuss the variation in estimates of the shape parameter κ when using pooled and non-pooled approaches.

When we fit GPDs to each replicate with their own set of mutations, almost all the fitted GPDs belonged to the Weibull domain (Supplementary Fig. 10). In addition, κ was quite variable between simulations, models, and adaptive steps. κ affects the GPD by modulating an upper bound on the effect size of beneficial mutations. With decreasing κ, this upper bound approaches zero. We observed that κ tended to decrease over longer walks under the network model. In the additive model, κ was relatively stable, however it slightly increased (particularly when comparing κ before the optimum shift and at adaptive steps ≥ 3 (Supplementary Fig. 11; Supplementary Table 2). Regardless of this increase, both models were largely Weibull (over 99% of simulations/adaptive steps belonged to the Weibull domain), which rejects Gillespie's Gumbel expectation.

Weibull-distributed DFEs of beneficial effects have been observed in a number of empirical studies in viral and bacterial populations (e.g. Rokyta *et al.* 2008; MacLean and Buckling 2009; Foll *et al.* 2014). Furthermore, a Weibull distribution of effects is predicted when populations are close to an optimum (as is the case in later adaptive steps) (Martin and Lenormand 2008). Evidence for limits on adaptation seems common and both models seem to be affected by such limitations. While both models were starved of large-effect beneficial mutations relative to the exponential expectation, this effect was more limiting for network models than additive as the walk went on. This was further highlighted by a major difference between network and additive responses to selection: the rate of trait evolution.

Trait adaptation and network evolution

In our simulations, we found that network simulations approached the optimum more slowly than additive models on average (Fig. 4A). To decipher the underlying cause of this disparity, we explored the distribution of new mutations and discovered a complex DFE in the network models, contrasting with the simple distributions observed in additive models (Fig. 6). The additive model distribution largely met Gillespie's expectations: the DFE was comprised of mostly neutral mutations with a small, exponential tail of beneficial mutations, and a longer tail of deleterious mutations (Fig. 6, Gillespie 1991 pg. 267, Fig. 6.5).

In the network model, we found a bimodal distribution of fitness effects. This observation aligns with previous studies that have reported complex and multimodal distributions of deleterious fitness effects (e.g. Kousathanas and Keightley 2013; Zeyl and DeVisser 2001), potentially attributed to different classes of mutations with distinct DFEs (Eyre-Walker and Keightley 2007). Mutations in gene regulatory regions can have widely varying effects on fitness (e.g. Halligan *et al.* 2013; Williamson *et al.* 2014; Ohta 2002), suggesting that network structure strongly affects the fitness distribution of mutants on a given molecular component (Erwin and Davidson 2009). This structure presents a considerable challenge to adaptation in the NAR model. The ridge-like shape of the α_Z/β_Z fitness landscape is evidence of this (Fig. 7).

Unlike the additive model, the fitness landscape between the molecular components of the NAR revealed non-interchangeable per-locus effects (Fig. 7A,C). Consequently, NAR populations faced lower beneficial mutation rates and a slight bias favoring α_Z mutations over β_Z (Fig. 6, 7D). This suggests α_Z mutations should be less deleterious on average than β_Z mutations. Supporting this, α_Z mutations were slightly less likely to be strongly deleterious, instead falling between the distribution's modes (Supplementary Fig. 13). However, both mutation types remained largely deleterious, and had similar DBFEs. A potential explanation for the bias towards α_Z mutations involves the optimal β_Z/α_Z ratio ($\theta_{\beta_Z/\alpha_Z}$) of 0.8 post-shift, versus 0.4 at burn-in. To reach this optimum from burn-in, populations could increase β_Z or decrease α_Z. However, the required α_Z decrease exceeded the β_Z increase. In essence, larger α_Z mutations were needed to hit the new optimum. Crucially, α_Z and β_Z mutations did not affect the phenotype equally.

Complexity and the cost of adaptation

In our study, the NAR model demonstrated greater complexity compared to traditional additive models. We discovered a nonlinear relationship between β_Z/α_Z and phenotype, embodied by a nonlinear genotype-phenotype (GP) map (Fig. 7C). Such maps are known to impose adaptive constraints (Wade *et al.* 2001; Wagner 2011), reflected in the intricate shape of the deleterious DFE for NAR populations (Fig. 6A). The NAR model's increased complexity through multiple molecular components and dynamic biochemical interactions renders it more susceptible to deleterious mutations. This vulnerability could amplify Hill-Robertson interference, with beneficial alleles overshad-

owed by deleterious backgrounds (Hill and Robertson 1966). Thus, network complexity may trade off with adaptive potential, partially mitigated by fixing certain molecular components to restrict deleterious mutation targets. In summary, the NAR model balances meaningful biological complexity with constraints on uncontrolled complexity that could overwhelm evolution.

Trade-offs between complexity and adaptability echo the cost of complexity described by Orr (2000). The cost of complexity postulates that an increasing number of axes (traits, or molecular components) leads to an increasingly deleterious mutation space, impeding adaptation. This notion finds support in studies that have explored network topology in different species. For instance, Costanzo *et al.* (2010) found that increasing the degree of genetic interaction in *Saccharomyces cerevisiae* genes was negatively correlated with single mutant fitness, suggesting an adaptive cost to maintaining a highly connected gene network. This cost might lead to the long-term stability of gene networks or complex traits: comparisons between *S. cerevisiae* and *Schizosaccharomyces pombe* genetic interaction networks revealed remarkable similarity in network structure (Koch *et al.* 2012). Another study by Barua and Mikheyev (2021) found features of housekeeping gene networks involved in reptile venom production were conserved between amniote clades, with those genes contributing to saliva production in mammals (Barua and Mikheyev 2021). The cost of complexity suggests some rather strong limitations on the evolution of complex genetic networks: in Orr's (2000) model, the rate of adaptation declines according to the inverse square root of the number of traits, $1/\sqrt{n}$. Empirical evidence suggests that biological networks can be extremely complicated, connecting hundreds of nodes. Adaptation involving such a highly connected network would be extremely slow if Orr's (2000) theory is correct. However, aspects of network structure, including local connectivity and modularity, might alleviate such a cost of molecular complexity.

One empirical example of network structure modulating the rate of evolution is the biofilm production network in *Candida* species (Mancera *et al.* 2021). This network consists of seven highly-connected master regulator genes connected to hundreds of target genes. These master regulators control the expression of about one-sixth of the *C. albicans* genome and are required to drive biofilm production (Mancera *et al.* 2021). Each target gene is connected to one or more master regulators, meaning they are much less connected than the master regulators and more "free" to evolve without affecting the expression of hundreds of other genes. Mancera *et al.* (2021) found that the master regulators evolve slowly compared to the target genes, possibly due to the structure of the network. Since the master regulators affect the expression of hundreds of genes, the space of beneficial mutations at master regulator loci might scale according to the cost of complexity. On the other hand, the authors found that target genes are considerably more divergent between *Candida* species and populations, perhaps due to target genes harboring fewer interactions within the network than master regulators (Mancera *et al.* 2021).

The *Candida* biofilm network is reminiscent of the omnigenic model of polygenic traits, where the degree of pleiotropy depends on gene regulatory network structure (Boyle *et al.* 2017). The omnigenic model considers complex traits as the result of their underlying gene regulatory networks. Under this model, loci of large effect represent "core" genes that describe variation in coding regions and *cis*-regulatory features (Liu *et al.* 2019; Boyle *et al.* 2017). Core gene expression is regulated by the far more numerous "peripheral" genes, which are comprised of *trans*-regulatory elements (Boyle *et al.* 2017). These networks are "small-world": that is, each gene in the network is usually connected to a core gene by only a few intermediates. Hence, the effects of a mutation in any given peripheral gene are likely to percolate to a connected core gene due to their closeness. The *Candida* example is a special case where peripheral genes contribute large effects to the phenotype (Liu *et al.* 2019). Here, the master regulators are peripheral genes, altering the expression of core (target) genes. This is one of the few examples where peripheral genes might be more highly constrained than core genes, owing to their high connectivity and hub-like structure (Liu *et al.* 2019). Some evidence suggests that hubs might be under stronger selective constraints than peripheral genes (Jordan *et al.* 2004; Hahn *et al.* 2004). Hence, network structure could be a driving factor in modulating the cost of complexity on the tempo and mode of adaptation, as well as the modularity of the network and its robustness to mutation.

The NAR motif serves as a simplified regulatory network, effectively sidestepping the full cost of complexity outlined by Orr (2000). Despite this, it imposes greater adaptive constraints than the additive model, a point underlined by the strongly deleterious DFE of new mutations (Fig. 6). Fixing multiple NAR molecular components (e.g. K_Z, K_{XZ}, n_Z, n_{XZ}) renders the model's behavior more predictable. For sufficiently large α_Z, individual fitness correlates solely with the ratio β_Z/α_Z (Fig. 7B), reflecting the steady-state cellular concentration under simple regulation (Alon 2019). This outcome may be attributed to our parameter choices; specifically, the rapid onset of Z activation/suppression due to $n_Z = n_{XZ} = 8$. Varying these Hill coefficients or treating them as evolvable components in future work could offer deeper insights into how more gradual responses to X activation/deactivation and Z production might modulate total Z expression (Alon 2019). Moreover, evolving parameters like K_Z and K_{XZ} could elucidate the full NAR fitness landscape. In essence, our simplified NAR model balances complexity and tractability, avoiding certain complexities while retaining meaningful constraints.

Should K_Z be allowed to evolve, thereby becoming a molecular component, the fitness landscape would likely exhibit increased ruggedness. In this context, the optimal β_Z/α_Z ratio would be contingent upon the K_Z value. A similar effect is anticipated for K_{XZ}, the regulator of Z's production rate in response to X. Consequently, both steady-state and fitness would depend on K_Z and K_{XZ}. Introducing more evolvable molecular components is expected to amplify this landscape ruggedness due to interdependencies that influence fitness, thereby exacerbating the previously noted cost of complexity. To substantiate these projections, future research should explore more intricate networks, as exemplified by empirical studies such as Bertheloot

et al. (2020). Advancing our understanding of how genetic networks inform fitness landscapes will yield invaluable insights into the adaptive strategies of organisms. In summary, making additional NAR parameters evolvable could unmask pivotal complexities in the genotype-phenotype map and their evolutionary ramifications.

Ramifications of strong selection weak mutation (SSWM) and alternative approaches

This study focused on adaptation by employing a Gillespie-Orr model, which operates under the assumption of SSWM, characterized by $Ns \gg 1$ and $N\mu \ll 1$. We ensured adherence to the SSWM assumption by choosing an appropriate population size and mutation rate (see Table 1). In our simulations, $N\mu = 0.092$. However, Ns varies due to selection coefficients being dependent on individuals' genetic backgrounds and phenotypes. Nonetheless, $Ns > 1$ exists at $s > 0.0002$ for this model, so most beneficial mutations should meet the $Ns \gg 1$ criterion. In the instance they do not, these mutations are driven by drift (Lynch 2010). This also applies to slightly deleterious mutations. $N|s| = 1$ represents a "drift barrier", where the fitness effects of alleles become too small to drive allele frequency changes, leading to drift-dominated dynamics.

The drift barrier explains our finding that the DFE among fixations followed a gamma distribution as opposed to an exponential in both network and additive models (Supplementary Fig. 6). As populations approach the optimum, selection's efficiency is weakened by the drift barrier (Lynch 2010). At the barrier, the probability of fixation (π) is no longer predicated by $\pi \approx 2s$ when the population is close to the optimum. Instead, $\pi = 1/2N$ becomes more applicable (Charlesworth and Charlesworth 2010). The switch from selection-dominated to drift-dominated dynamics occurs when $N|s| = 1$ (Koonin 2016). In other words, slightly-deleterious and slightly-beneficial alleles are more likely to be driven by drift than larger effect alleles. In our simulations, this means that when $|s| < 0.0002$, drift is expected to dominate. Note that this is not a hard threshold: the probability of fixation is between $2s$ and $1/2N$, approaching the latter as $N|s| \to 1$. Hence, drift can still influence an allele's fixation probability when $N|s| > 1$ (Clatterbuck *et al.* 2013). Owing to the relatively large population size in our simulations ($N = 5000$), drift is unlikely to drive allele fixations early in the adaptive walk (i.e. when most beneficial alleles have $s > 0.0002$), however later steps might be more affected as the possible fitness improvement of any given allele faces diminishing returns. In a way, the SSWM conditions ensure the drift barrier hypothesis will exist as populations near the optimum: without environmental change, new mutations at advanced adaptive steps will inevitably approach the drift barrier.

Another assumption of the SSWM model is that only one mutation segregates at any given time point, i.e. before a new beneficial mutation arises, the previous beneficial mutation must have fixed. This was met most of the time (see Supplementary Fig. 8, 9), however, some simulations revealed instances where segregating variation contributed to adaptation, including the presence of alleles under balancing selection (Supplementary Fig. 12). These observations highlight the need for future research to explore the model using a polygenic adaptation approach. Such an approach would encompass the involvement of many loci, each contributing small amounts to the molecular components underlying complex traits. By adopting this perspective, alternative paradigms of adaptation, such as strong mutation and weak selection, can be investigated. In these scenarios, adaptation occurs through small frequency shifts at multiple loci simultaneously (Walsh and Lynch 2018c). The genetic architecture of polygenic traits, that is, the number of loci, frequencies, and the distributions of their additive and non-additive effects on phenotypes, can greatly affect the response to selection (Barghi *et al.* 2020). The non-additive effects of genetic networks on phenotypes in particular has been understudied. These interactions have ramifications for theories of the evolution of recombination (Ritz *et al.* 2017), and the persistence of linkage disequilibrium (LD) in natural populations (Ortiz-Barrientos *et al.* 2016).

Opportunities and challenges in modeling complex genetic networks

Our study serves as an introductory venture into the evolutionary modeling of genetic networks, elucidating how these networks guide evolutionary trajectories. The strong selection and weak mutation assumptions grant computational tractability but sideline rich polygenic dynamics (Barghi *et al.* 2020; Thornton 2019). The NAR motif offers initial insights but lacks empirical complexity, making future work with more complex architectures important. Moving forward, our focus will shift towards a spectrum of genetic architectures to better capture non-additive effects. This transition will require further statistical and computational innovations to maintain feasibility. We also aim to investigate a diversity of network structures beyond negative autoregulation, enriching our understanding of how topology influences evolutionary pathways (Siegal *et al.* 2007). Collaborations with molecular biologists will enrich our models with further mechanistic details. Rigorous validation against empirical data and existing theories will be a cornerstone in constructing more realistic models of adaptation (e.g. Bertheloot *et al.* 2020).

As we delve into more complex systems, computational hurdles become prominent. The complexity of the ODE system significantly impacts *in silico* performance (Agocs 2020; Städter *et al.* 2021). We are considering optimization strategies such as smarter caching of common genotypes, data-oriented design, and adaptive step-size algorithms. Pre-computed ODE solutions stored on disk are also under exploration as a means to tackle the computational load. A final point of consideration is the difference in allelic effect scaling between additive and network models, as shown in Eq. 3 and Eq. 11. This scaling difference could influence rates of adaptation, mandating caution in attributing observed differences solely to the network's inherent complexities. Hence, we do not make conclusions on the NAR network's effect on the rate of adaptation due to the scaling difference between the additive and network models. In summary, our study, while preliminary, opens up avenues for understanding the complex interplay between genetic networks

and evolutionary dynamics, leaving us optimistic about the future of this research area.

Conclusions

Our network model contributes methodologically to the field and holds broad applicability to many open questions in evolutionary biology, from quantitative genetics to the mechanisms underpinning adaptability. It offers a robust framework to investigate issues such as the maintenance of genetic variation, the evolution of recombination, and the dynamics of adaptation from both standing variation and *de novo* mutations. We have introduced how networks can generate negative epistasis, with implications on the evolution of recombination, maintenance of linkage, and the trade-offs between network complexity and adaptability. While promising, it is essential to acknowledge the model's limitations, particularly concerning its computational demands and simplifications, which future work should aim to address. Empirical validation through evolutionary experiments will be crucial in assessing our model's real-world applicability. Expanding into adaptation by standing variation will further illuminate the role of genetic interactions during polygenic adaptation. As detailed systems models of complex traits become increasingly available, our framework can be employed to simulate the evolution of traits based on empirically-described networks. This will provide a more nuanced understanding of how molecular mechanisms might constrain adaptation. The utility of our model also extends beyond evolutionary biology, offering valuable insights for interdisciplinary collaborations that aim to integrate molecular biology, computational science, and statistical genetics. Future work should focus on leveraging this interdisciplinary potential and on exploring more complex genetic networks to better understand how gene interactions limit evolution.

Author contributions

DO conceived the original idea. NO, JE and DO developed together the new theoretical approach, constructed the models and designed simulation experiments. NO implemented and ran the simulations. BH and JE provided mathematical guidance. NO wrote the paper with assistance from all authors.

Data availability

The data used in this study is available at https://doi.org/10.48610/f3850b0. Instructions and scripts for repeating the analysis and for running the SLiM simulations are available at https://github.com/nobrien97/NARAdaptiveWalk2023. SLiM modifications are available at https://github.com/nobrien97/SLiM/releases/tag/AdaptiveWalks2023: note that this modification requires the Ascent ODE solution library installed https://github.com/AnyarInc/Ascent/releases/tag/v0.7.0.

Acknowledgments

We would like to acknowledge members of the Ortiz-Barrientos lab for their feedback on earlier versions of this manuscript. This research was conducted using the Gadi HPC system maintained by the National Computational Infrastructure (NCI), which is supported by the Australian Government.

Funding

This work was supported by an Australian Research Council grant awarded to DO (FT200100169) and the Australian Research Council Centre of Excellence for Plant Success in Nature and Agriculture (CE200100015).

Conflicts of interest

The authors declare no conflicts of interest.

Literature cited

Agocs FJ. 2020. (py)oscode: Fast solutions of oscillatory ODEs. Journal of Open Source Software. 5:2830.

Alon U. 2019. *An Introduction to Systems Biology: Design Principles of Biological Circuits*. CRC Press. Boca Raton, FL, USA. second edition.

Álvarez-Castro JM, Carlborg Ö. 2007. A unified model for functional and statistical epistasis and its application in quantitative trait loci analysis. Genetics. 176:1151–1167.

Ang RML, Chen SAA, Kern AF, Xie Y, Fraser HB. 2023. Widespread epistasis among beneficial genetic variants revealed by high-throughput genome editing. Cell Genomics. 3:100260.

Aston E, Channon A, Belavkin RV, Gifford DR, Krašovec R, Knight CG. 2017. Critical mutation rate has an exponential dependence on population size for eukaryotic-length genomes with crossover. Sci Rep. 7:15519.

Barghi N, Hermisson J, Schlötterer C. 2020. Polygenic adaptation: A unifying framework to understand positive selection. Nat Rev Genet. 21:769–781.

Barua A, Mikheyev AS. 2021. An ancient, conserved gene regulatory network led to the rise of oral venom systems. PNAS. 118:e2021311118.

Beisel CJ, Rokyta DR, Wichman HA, Joyce P. 2007. Testing the extreme value domain of attraction for distributions of beneficial fitness effects. Genetics. 176:2441–2449.

Bendixsen DP, Østman B, Hayden EJ. 2017. Negative epistasis in experimental RNA fitness landscapes. J Mol Evol. 85:159–168.

Berry S, Walcott M, Avalos CG, Du R. 2021. Ascent. AnyarInc.

Bertheloot J, Barbier F, Boudon F, Perez-Garcia MD, Péron T, Citerne S, Dun E, Beveridge C, Godin C, Sakr S. 2020. Sugar availability suppresses the auxin-induced strigolactone pathway to promote bud outgrowth. New Phytologist. 225:866–879.

Bolton MD, Thomma BPHJ, Nelson BD. 2006. *Sclerotinia Sclerotiorum* (Lib.) de Bary: Biology and molecular traits of a cosmopolitan pathogen. Mol Plant Pathol. 7:1–16.

Bomblies K, Peichel CL. 2022. Genetics of adaptation. PNAS. 119:e2122152119.

Boyle EA, Li YI, Pritchard JK. 2017. An expanded view of complex traits: From polygenic to omnigenic. Cell. 169:1177–1186.

Bradshaw HD, Wilbert SM, Otto KG, Schemske DW. 1995. Genetic mapping of floral traits associated with reproductive isolation in monkeyflowers (*Mimulus*). Nature. 376:762–765.

Breen MS, Kemena C, Vlasov PK, Notredame C, Kondrashov FA. 2012. Epistasis as the primary factor in molecular evolution. Nature. 490:535–538.

Charlesworth B, Charlesworth D. 2010. *Elements of Evolutionary Genetics*. Roberts and Company. Greenwoord Village, Colorado, USA.

Charras-Garrido M, Lezaud P. 2013. Extreme value analysis: An introduction. Journal de la SFdS. 154:66–97.

Claringbould A, de Klein N, Franke L. 2017. The genetic architecture of molecular traits. Curr. Opin. Syst. Biol. 1:25–31.

Clatterbuck H, Sober E, Lewontin R. 2013. Selection never dominates drift (nor vice versa). Biol Philos. 28:577–592.

Costanzo M, Baryshnikova A, Bellay J, Kim Y, Spear ED, Sevier CS, Ding H, Koh JL, Toufighi K, Mostafavi S *et al*. 2010. The genetic landscape of a cell. Science. 327:425–431.

Cowperthwaite MC, Bull JJ, Meyers LA. 2005. Distributions of beneficial fitness effects in RNA. Genetics. 170:1449–1457.

Delignette-Muller ML, Dutang C. 2015. fitdistrplus: An R package for fitting distributions. J. Stat. Softw. 64:1–34.

Doane AS, Elemento O. 2017. Regulatory elements in molecular networks. WIREs Systems Biology and Medicine. 9:e1374.

Dorweiler J, Stec A, Kermicle J, Doebley J. 1993. *Teosinte Glume Architecture 1*: A genetic locus controlling a key step in maize evolution. Science. 262:233–235.

Draghi JA, Parsons TL, Plotkin JB. 2011. Epistasis increases the rate of conditionally neutral substitution in an adapting population. Genetics. 187:1139–1152.

Erwin DH, Davidson EH. 2009. The evolution of hierarchical gene regulatory networks. Nat Rev Genet. 10:141–148.

Eyre-Walker A, Keightley PD. 2007. The distribution of fitness effects of new mutations. Nat Rev Genet. 8:610–618.

Falconer DSM. 1996. *Introduction to Quantitative Genetics*. Pearson Education Limited. Longmans Green, Harlow, Essex, UK. fourth edition.

Felsenstein J. 1979. Excursions along the interface between disruptive and stabilizing selection. Genetics. 93:773–795.

Fisher RA. 1918. The correlation between relatives on the supposition of Mendelian inheritance. Trans. R. Soc. Edinburgh. 52:399–433.

Fisher RA. 1923. XXI.—On the dominance ratio. Proc. R. Soc. Edinb. 42:321–341.

Fisher RA. 1930. *The Genetical Theory of Natural Selection*. The Clarendon press. Oxford, UK.

Fisher RA. 1936. Statistical methods for research workers. Statistical methods for research workers.. .

Fisher RA, Tippett LHC. 1928. Limiting forms of the frequency distribution of the largest or smallest member of a sample. Math. Proc. Camb. Philos. Soc. 24:180–190.

Foll M, Poh YP, Renzette N, Ferrer-Admetlla A, Bank C, Shim H, Malaspinas AS, Ewing G, Liu P, Wegmann D *et al*. 2014. Influenza virus drug resistance: A time-sampled population genetics perspective. PLOS Genetics. 10:e1004185.

François P. 2014. Evolving phenotypic networks in silico. Semin. Cell Dev. Biol. 35:90–97.

Gillespie JH. 1983. A simple stochastic gene substitution model. Theor. Popul. Biol. 23:202–215.

Gillespie JH. 1984. Molecular evolution over the mutational landscape. Evolution. 38:1116–1129.

Gillespie JH. 1991. *The Causes of Molecular Evolution*. University Press.

Hahn MW, Conant GC, Wagner A. 2004. Molecular evolution in large genetic networks: Does connectivity equal constraint? J Mol Evol. 58:203–211.

Halfon MS. 2020. Silencers, enhancers, and the multifunctional regulatory genome. Trends in Genetics. 36:149–151.

Haller BC, Messer PW. 2019. SLiM 3: Forward genetic simulations beyond the Wright-Fisher model. Mol Biol Evol. 36:632–637.

Halligan DL, Kousathanas A, Ness RW, Harr B, Eöry L, Keane TM, Adams DJ, Keightley PD. 2013. Contributions of protein-coding and regulatory change to adaptive molecular evolution in murid rodents. PLOS Genetics. 9:e1003995.

Hansen TF. 2013. Why epistasis is important for selection and adaptation. Evolution. 67:3501–3511.

Hill WG, Robertson A. 1966. Effect of linkage on limits to artificial selection. Genet Res. 8:269–294.

Huang W, Mackay TFC. 2016. The genetic architecture of quantitative traits cannot be inferred from variance component analysis. PLOS Genetics. 12:e1006421.

Imhof M, Schlötterer C. 2001. Fitness effects of advantageous mutations in evolving *Escherichia coli* populations. PNAS. 98:1113–1117.

Jordan IK, Mariño-Ramírez L, Wolf YI, Koonin EV. 2004. Conservation and coevolution in the scale-free human gene coexpression network. Mol Biol Evol. 21:2058–2070.

Joyce P, Rokyta DR, Beisel CJ, Orr HA. 2008. A general extreme value theory model for the adaptation of DNA sequences under strong selection and weak mutation. Genetics. 180:1627–1643.

Karlebach G, Shamir R. 2008. Modelling and analysis of gene regulatory networks. Nat Rev Mol Cell Biol. 9:770–780.

Kassen R, Bataillon T. 2006. Distribution of fitness effects among beneficial mutations before selection in experimental populations of bacteria. Nat Genet. 38:484–488.

Kauffman S, Levin S. 1987. Towards a general theory of adaptive walks on rugged landscapes. J. Theor. Biol. 128:11–45.

Kemper KE, Visscher PM, Goddard ME. 2012. Genetic architecture of body size in mammals. Genome Biol. 13:244.

Keon J, Byrde R. 1987. Some aspects of fungal enzymes that degrade plant cell walls, In: *Fungal Infection of Plants: Symposium of the British Mycological Society*, Cambridge University Press. pp. 133–157.

Kimura M. 1983. *The Neutral Theory of Molecular Evolution*. Cambridge University Press. Cambridge.

Koch EN, Costanzo M, Bellay J, Deshpande R, Chatfield-Reed K, Chua G, D'Urso G, Andrews BJ, Boone C, Myers CL. 2012. Conserved rules govern genetic interaction degree across species. Genome Biology. 13:R57.

Koonin EV. 2016. Splendor and misery of adaptation, or the importance of neutral null for understanding evolution. BMC

Biology. 14:114.
Kopp M, Hermisson J. 2009. The genetic basis of phenotypic adaptation II: The distribution of adaptive substitutions in the moving optimum model. Genetics. 183:1453–1476.
Kosheleva K, Desai MM. 2018. Recombination alters the dynamics of adaptation on standing variation in laboratory yeast populations. Mol Biol Evol. 35:180–201.
Kousathanas A, Keightley PD. 2013. A comparison of models to infer the distribution of fitness effects of new mutations. Genetics. 193:1197–1208.
Lande R. 1976. Natural selection and random genetic drift in phenotypic evolution. Evolution. 30:314–334.
Lande R, Shannon S. 1996. The role of genetic variation in adaptation and population persistence in a changing environment. Evolution. 50:434–437.
Lebeuf-Taylor E, McCloskey N, Bailey SF, Hinz A, Kassen R. 2019. The distribution of fitness effects among synonymous mutations in a gene under directional selection. eLife. 8:e45952.
Lenth RV. 2023. *Emmeans: Estimated Marginal Means, Aka Least-Squares Means*.
Liu X, Li YI, Pritchard JK. 2019. Trans effects on gene expression can drive omnigenic inheritance. Cell. 177:1022–1034.e6.
Lynch M. 2010. Evolution of the mutation rate. Trends in Genetics. 26:345–52.
MacLean RC, Buckling A. 2009. The distribution of fitness effects of beneficial mutations in *Pseudomonas aeruginosa*. PLOS Genetics. 5:e1000406.
Mancera E, Nocedal I, Hammel S, Gulati M, Mitchell KF, Andes DR, Nobile CJ, Butler G, Johnson AD. 2021. Evolution of the complex transcription network controlling biofilm formation in *Candida* species. eLife. 10:e64682.
Martin G, Lenormand T. 2008. The distribution of beneficial and fixed mutation fitness effects close to an optimum. Genetics. 179:907–16.
Matuszewski S, Hermisson J, Kopp M. 2015. Catch me if you can: Adaptation from standing genetic variation to a moving phenotypic optimum. Genetics. 200:1255–1274.
Maynard Smith J. 1970. Natural selection and the concept of a protein space. Nature. 225:563–564.
Maynard Smith J, Burian R, Kauffman S, Alberch P, Campbell J, Goodwin B, Lande R, Raup D, Wolpert L. 1985. Developmental constraints and evolution: A perspective from The Mountain Lake Conference on Development and Evolution. Q. Rev. Biol.. 60:265–287.
Morimoto D, Walinda E, Fukada H, Sugase K, Shirakawa M. 2016. Ubiquitylation directly induces fold destabilization of proteins. Sci Rep. 6:39453.
Nevozhay D, Adams RM, Murphy KF, Josić K, Balázsi G. 2009. Negative autoregulation linearizes the dose–response and suppresses the heterogeneity of gene expression. PNAS. 106:5123–5128.
Ohta T. 2002. Near-neutrality in evolution of genes and gene regulation. PNAS. 99:16134–16137.
Orr HA. 1998. The population genetics of adaptation: The distribution of factors fixed during adaptive evolution. Evolution. 52:935–949.
Orr HA. 2000. Adaptation and the cost of complexity. Evolution. 54:13–20.
Orr HA. 2003a. The distribution of fitness effects among beneficial mutations. Genetics. 163:1519–1526.
Orr HA. 2003b. A minimum on the mean number of steps taken in adaptive walks. J. Theor. Biol.. 220:241–247.
Orr HA. 2010. The population genetics of beneficial mutations. Philos. Trans. R. Soc. Lond., B, Biol. Sci.. 365:1195–1201.
Ortiz-Barrientos D, Engelstädter J, Rieseberg LH. 2016. Recombination rate evolution and the origin of species. TREE. 31:226–236.
R Core Team. 2023. *R: A Language and Environment for Statistical Computing*. Vienna, Austria.
Rice SH. 1998. The evolution of canalization and the breaking of Von Baer's laws: Modeling the evolution of development with epistasis. Evolution. 52:647–656.
Rice SH. 2008. Theoretical approaches to the evolution of development and genetic architecture. Ann. N. Y. Acad. Sci.. 1133:67–86.
Ritz KR, Noor MAF, Singh ND. 2017. Variation in recombination rate: Adaptive or not? Trends in Genetics. 33:364–374.
Rokyta DR, Beisel CJ, Joyce P, Ferris MT, Burch CL, Wichman HA. 2008. Beneficial fitness effects are not exponential for two viruses. J Mol Evol. 67:368–376.
Rosenfeld N, Elowitz MB, Alon U. 2002. Negative autoregulation speeds the response times of transcription networks. J Mol Biol. 323:785–793.
Sanjuán R, Moya A, Elena SF. 2004. The distribution of fitness effects caused by single-nucleotide substitutions in an RNA virus. PNAS. 101:8396–8401.
Schenk MF, Szendro IG, Krug J, de Visser JAGM. 2012. Quantifying the adaptive potential of an antibiotic resistance enzyme. PLOS Genetics. 8:e1002783.
Schlitt T, Brazma A. 2007. Current approaches to gene regulatory network modelling. BMC Bioinformatics. 8:S9.
Schmidtke G, Kalveram B, Weber E, Bochtler P, Lukasiak S, Hipp MS, Groettrup M. 2006. The UBA domains of NUB1L are required for binding but not for accelerated degradation of the ubiquitin-like modifier FAT10 *. JBC. 281:20045–20054.
Siegal ML, Promislow DEL, Bergman A. 2007. Functional and evolutionary inference in gene networks: Does topology matter? Genetica. 129:83–103.
Simpson GG. 1944. *Tempo and Mode in Evolution.*. Columbia University Press, N.Y.
Slatkin M. 1987. Quantitative genetics of heterochrony. Evolution. 41:799–811.
Städter P, Schälte Y, Schmiester L, Hasenauer J, Stapor PL. 2021. Benchmarking of numerical integration methods for ODE models of biological systems. Sci Rep. 11:2696.
Thornton KR. 2019. Polygenic adaptation to an environmental shift: Temporal dynamics of variation under gaussian stabilizing selection and additive effects on a single trait. Genetics. 213:1513–1530.
Vale PF, Choisy M, Froissart R, Sanjuán R, Gandon S. 2012. The distribution of mutational fitness effects of phage Φx174

on different hosts. Evolution. 66:3495–3507.
Visscher PM, Brown MA, McCarthy MI, Yang J. 2012. Five years of GWAS discovery. The American Journal of Human Genetics. 90:7–24.
Visscher PM, Goddard ME. 2019. From R.A. Fisher's 1918 paper to GWAS a century later. Genetics. 211:1125–1130.
Wade MJ, Winther RG, Agrawal AF, Goodnight CJ. 2001. Alternative definitions of epistasis: Dependence and interaction. TREE. 16:498–504.
Wagner A. 1996. Does evolutionary plasticity evolve? Evolution. 50:1008–1023.
Wagner A. 2011. Genotype networks shed light on evolutionary constraints. TREE. 26:577–584.
Walsh B, Lynch M. 2018a. *Evolution and Selection of Quantitative Traits*. Oxford University Press. New York, NY, USA.
Walsh B, Lynch M. 2018b. Long-term response: 3. Adaptive walks, In: *Evolution and Selection of Quantitative Traits*, Walsh B, Lynch M, editors, Oxford University Press. pp. 991–1013.
Walsh B, Lynch M. 2018c. Maintenance of quantitative genetic variation, In: *Evolution and Selection of Quantitative Traits*, Walsh B, Lynch M, editors, Oxford University Press. pp. 1016–1078.
Walsh B, Lynch M. 2018d. Short-term Changes in the Mean: 1. The Breeder's Equation, In: *Evolution and Selection of Quantitative Traits*, Walsh B, Lynch M, editors, Oxford University Press. pp. 481–506.
Weinreich DM, Delaney NF, DePristo MA, Hartl DL. 2006. Darwinian evolution can follow only very few mutational paths to fitter proteins. Science. 312:111–114.
Wellenreuther M, Hansson B. 2016. Detecting polygenic evolution: Problems, pitfalls, and promises. Trends Genet. 32:155–164.
Wickham H. 2016. *Ggplot2: Elegant Graphics for Data Analysis*. Use R!. Springer-Verlag New York.
Williamson RJ, Josephs EB, Platts AE, Hazzouri KM, Haudry A, Blanchette M, Wright SI. 2014. Evidence for widespread positive and negative selection in coding and conserved noncoding regions of *Capsella grandiflora*. PLOS Genetics. 10:e1004622.
Wright S. 1931. Evolution in Mendelian populations. Genetics. 16:97–159.
Wright S. 1932. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. Proceedings of the XI International Congress of Genetics. 8:209–222.
Xiang Y, Gubian S, Suomela B, Hoeng J. 2013. Generalized simulated annealing for efficient global optimization: The GenSA package for R. The R Journal. 5/1.
Xu L, Chen H, Hu X, Zhang R, Zhang Z, Luo ZW. 2006. Average gene length is highly conserved in prokaryotes and eukaryotes and diverges only between the two kingdoms. Mol Biol Evol. 23:1107–8.
Zeng ZB. 1988. Long-term correlated response, interpopulation covariation, and interspecific allometry. Evolution. 42:363–374.
Zeyl C, DeVisser JAGM. 2001. Estimates of the rate and distribution of fitness effects of spontaneous mutation in *Saccharomyces cerevisiae*. Genetics. 157:53–61.

Tables

1

Table 1 Simulation parameters

Symbol	Parameter	Description	Value
N	Population size	Number of individuals in the population	5,000
μ	Mutation rate	Per-locus per-generation rate of mutation in the genome	9.1528e-06
τ	Mutational effect variance	Variance in the sampling distribution of mutational effects on molecular components, $a \sim \mathcal{N}(0, \tau)$	1.0
L_Q	Number of causal loci	Number of quantitative trait loci/molecular quantitative trait loci in the genome	2
n_l	Genome length	Total number of loci modelled in the genome. Non-mQTLs generate neutral non-trait related variation.	1000
r_{EU}	Polygenic mQTL ratio	Ratio of α_Z mQTLs to β_Z mQTLs in the genome	1:1
σ_w	Fitness function width	Width of the Gaussian fitness function (5% drop in fitness for individuals one unit away from the optimum).	$\sqrt{10}$
O_{shift}	Optimum shift amount	Magnitude of shift in the phenotypic optimum after the burn-in phase	1
$n_{burn-in}$	Number of burn-in generations	Number of generations of burn-in prior to shifting the optimum.	50,000
n_{gen}	Number of adaptation generations	Number of generations the model was run for following the optimum shift	10,000

Supplementary material

Supplementary Table 1 Model symbols and parameters.

Symbol	Parameter	Description
NAR molecular components		
α_Z	Z removal rate	The rate at which Z product is removed from the cell.
β_Z	Z production rate	The rate at which Z is produced.
K_Z	Repression coefficient	The Z product concentration at which further Z expression is reduced by half. Fixed at $K_Z = 1$.
K_{XZ}	Activation coefficient	The X product concentration at which Z expression is half the maximum. Fixed at $K_{XZ} = 1$.
n_Z	Hill repression coefficient	The steepness coefficient of the Z repression curve. Fixed at $n_Z = 8$.
n_{XZ}	Hill activation coefficient	The steepness coefficient of the Z activation curve. Fixed at $n_{XZ} = 8$.
Model hierarchy		
P	Phenotype	The quantitative trait being measured.
$\mathbf{M}$	Molecular traits	The vector of molecular traits contributing to the phenotype. M_i represents the value of molecular trait i.
$\mathbf{C}$	Molecular components	The vector of molecular components contributing to molecular traits. C_i represents the value of molecular component i.
G	Genotype	The genes contributing to molecular components.
Genotype to molecular component		
n_C	Number of molecular components	The number of molecular components contributing to the phenotype.
A^{ij}	An allele at locus i and chromosome j	The vector of allelic effects across the n_M molecular components at locus i and chromosome j.
$\mathbf{C}^{base}$	Baseline molecular component value	A vector of baseline rate/concentrations of products produced by wild-type molecular component values. C_i^{base} refers to the baseline molecular component value of molecular component i.
L_Q	Number of causal loci	The number of loci contributing to molecular components.
Molecular component to molecular trait		
n_M	Number of molecular traits	The number of molecular traits contributing to the complex trait.
$\mathbf{E_M}(t)$	Molecular trait expression functions	A vector of n_M functions describing molecular trait expression at time t.
$\mathbf{E_M}'(t)$	First derivative of $\mathbf{E_M}(t)$	The rate of change in $\mathbf{E_M}$ over t.
$\hat{\mathbf{E}}_\mathbf{M}$	Solution to the initial value problem $\mathbf{E_M}(0)$	The ODE solution.
$h(\hat{\mathbf{E}}_\mathbf{M})$	Functional for measuring molecular traits	A functional describing how to measure the solution to the molecular trait function/s. $\mathbf{M} = h(\hat{\mathbf{E}}_M)$.
NAR-specific molecular component to molecular trait		
t_{start}	X activation time	The left boundary where $X = 1$ during NAR evaluation. Before this point, $X = 0$.

Continued on next page

Symbol	Parameter	Description
t_{stop}	X deactivation time	The right boundary where $X = 1$ during NAR evaluation. After this point, $X = 0$.
t_{max}	Development time	The amount of time for which $Z(t)$ is evaluated for. Models the developmental time of the molecular trait in an individual.
Molecular trait to complex trait		
$P(\mathbf{M})$	Phenotype construction function	The function combining the molecular traits.
Complex trait to fitness		
$w(\mathbf{z})$	Fitness function	The relationship between an individual's phenotype, $\mathbf{P}$, and fitness, w.
$\mathbf{P_O}$	Phenotypic optimum	The phenotype where $w = 1$.
$\mathbf{z}$	Phenotypic deviation from the optimum	An individual's deviation from the optimum phenotype, $\mathbf{z} = \mathbf{P_O} - \mathbf{P}$.
Fitness to the next generation		
n_l	Genome length	Number of loci in an individual's genome.
l_i^Q	Molecular quantitative trait locus i	Causal locus i contributing to a molecular component.
a	Allelic effect	The allelic effect of an mQTL on a molecular component.
a_{old}	Previous allelic effect	The previous allelic effect at an mQTL.

Supplementary Table 2 Mean generalized Pareto distribution (GPD) parameters fit to mutant screen distributions of fitness effects over adaptive walks in additive and network populations. Brackets indicate 95% confidence intervals. Parameters were estimated by fitting a GPD to a random sample of mutations from the mutant screen experiments conducted on each replicate at each adaptive step. κ is the shape parameter of the GPD. $\bar{\kappa}$ is the mean κ value across n replicate simulations. The log-likelihood ratio tests the null hypothesis that a sample of beneficial mutations belongs to an exponential distribution (which meets $\kappa = 0$). P-values across the n replicates were combined using Fisher's method.

Additive model				
Adaptive Step	n	$\bar{\kappa}$	**Log-likelihood ratio**	**p-value**
Before optimum shift	1213	$-2.12(\pm 0.04)$	$0.73(\pm 0.003)$	< 0.001
1	1110	$-2.10(\pm 0.04)$	$0.83(\pm 0.004)$	< 0.001
2	333	$-2.02(\pm 0.08)$	$0.87(\pm 0.006)$	< 0.001
≥ 3	59	$-1.85(\pm 0.11)$	$0.89(\pm 0.013)$	< 0.001
Network model				
Adaptive Step	n	$\bar{\kappa}$	**Log-likelihood ratio**	**p-value**
Before optimum shift	950	$-1.88(\pm 0.03)$	$0.77(\pm 0.002)$	< 0.001
1	912	$-2.38(\pm 0.25)$	$0.83(\pm 0.004)$	< 0.001
2	270	$-4.34(\pm 1.44)$	$0.86(\pm 0.006)$	< 0.001
≥ 3	48	$-3.81(\pm 2.00)$	$0.88(\pm 0.012)$	< 0.001

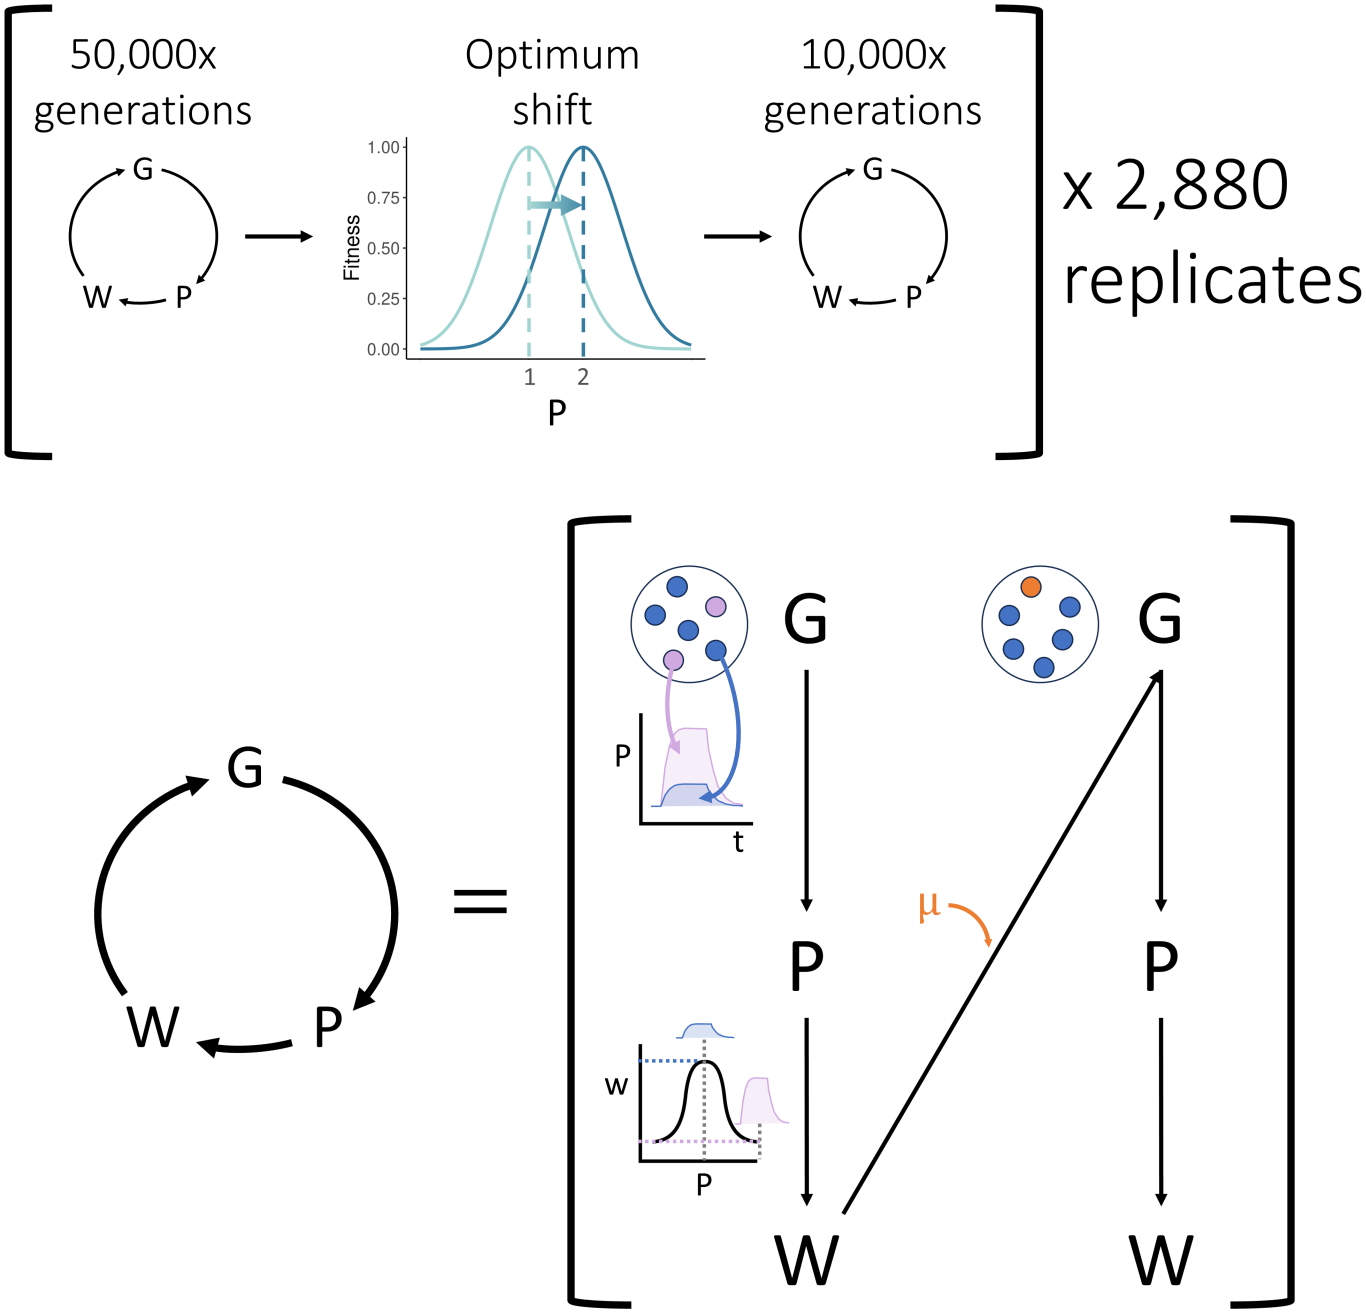

Supplementary Figure 1 Flowchart of SLiM simulation procedure. SLiM simulates a Wright-Fisher process to model evolutionary change through a genotype-phenotype-fitness (GPW) map. Our simulation began with 50,000 generations of burn-in to ensure populations were adapted to the environment. We then shifted the phenotypic optimum and adaptation was tracked for a further 10,000 generations. This process was repeated 2,880 times per model for replication purposes (i.e. there were 2,880 replicates of network and additive adaptive walks for a total 5,760 simulations). The GPW map consisted of several stages: first, the genotype was translated to phenotype via summing genetic effects at QTLs (additive models; not shown) or by solving a system of ordinary differential equations (network models, P/t figure shows solutions to the differential equation for the blue and purple genotypes). Phenotype was then translated to fitness by a stabilizing selection fitness function (shown by the w/P figure – the purple phenotype has lower fitness than the blue phenotype). The fitness value then influenced the chance that an individual was sampled as a parent for the next generation. After parents were chosen, random mutations (μ) could occur to introduce further genotypic and phenotypic variation in the next generation (shown in orange).

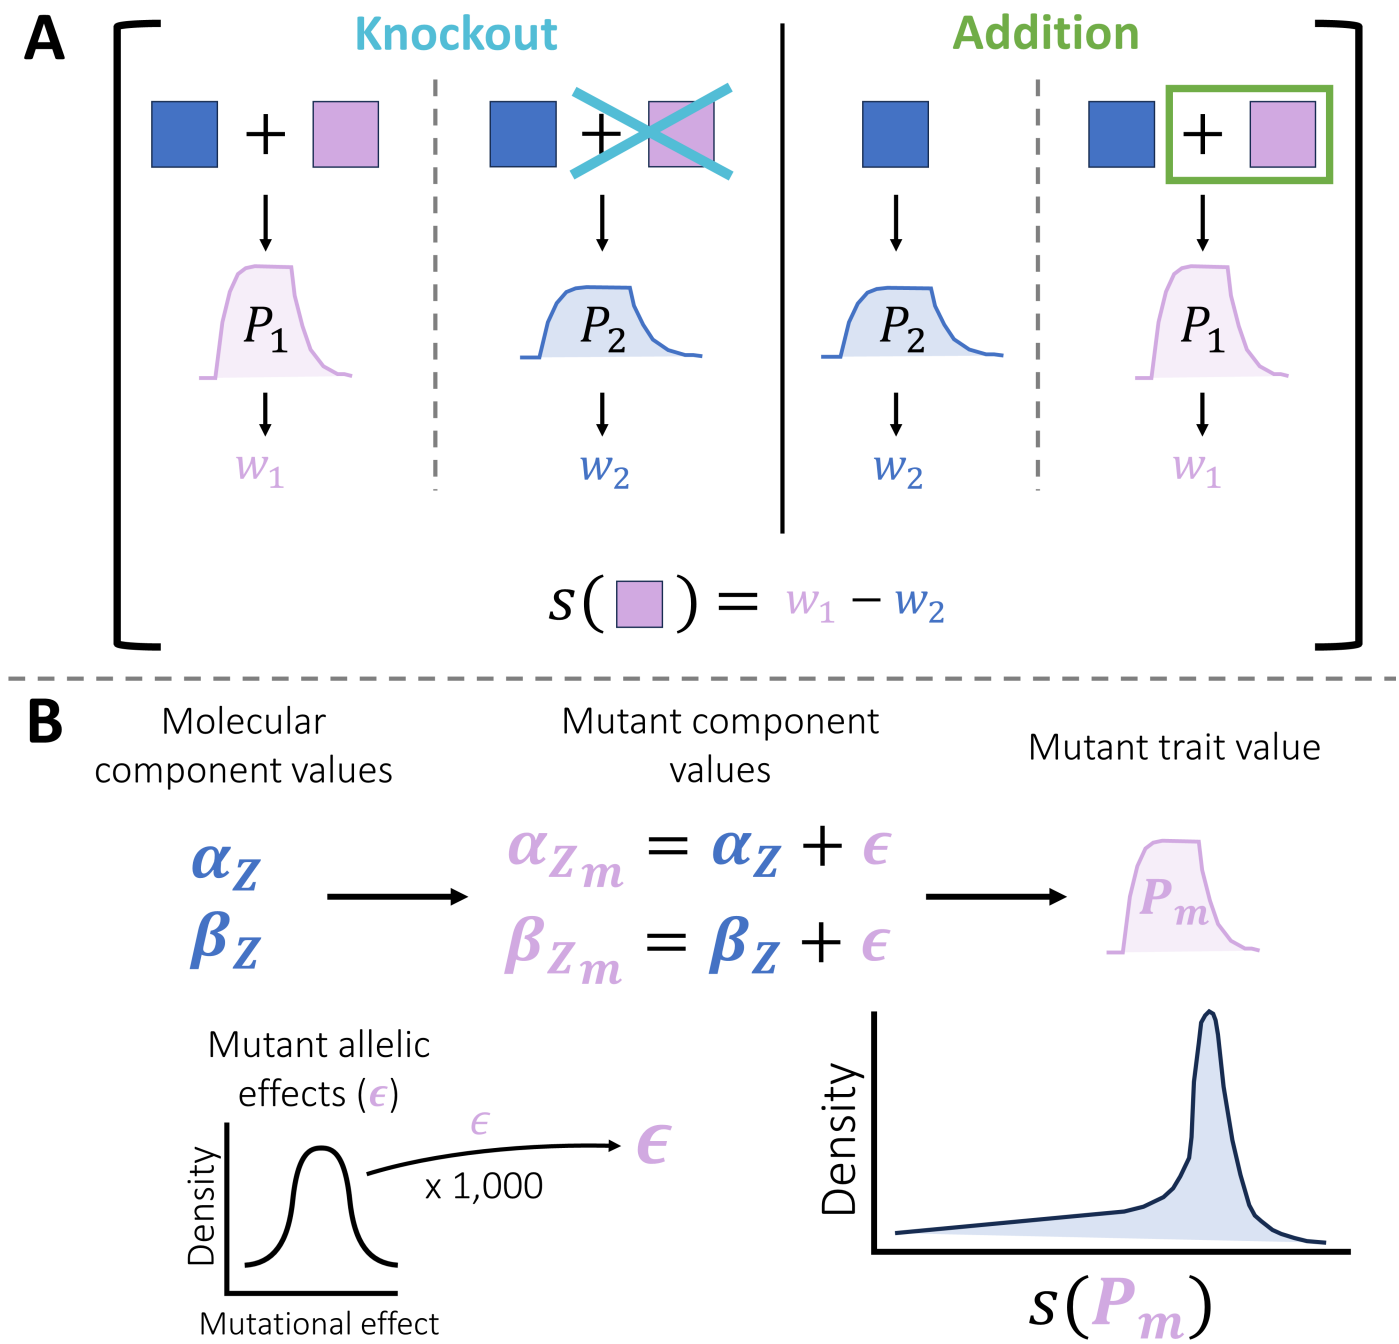

Supplementary Figure 2 Flowchart describing how selection coefficients and the distribution of fitness effects (DFE) among new mutations were estimated. (A) Consider two alleles contributing to the phenotype (blue and purple boxes) at either the same or different loci. Individuals with both alleles have some phenotype (P_1), which gives rise to a fitness (w_1). By removing the purple allele and recalculating the phenotype, we achieve a different phenotype (P_2) and fitness (w_2). The difference between w_1 and w_2 represents the selection coefficient (s) of the purple allele. This difference can be measured either through addition – adding the purple allele to the genotype, or by knockout (removing the purple allele from the genotype). (B) To estimate the distribution of fitness effects (DFE) among new mutations, we conducted a mutation screen experiment. We generated mutants by taking 1,000 samples (ϵ) from a standard normal distribution and adding those to the molecular component values from the SLiM simulations (α_Z and β_Z). Mutant phenotypes ($\mathbf{P_m}$) were calculated by inputting the mutant component values into an ordinary differential equation and solving it. The 1,000 samples were independently added to each molecular component to measure the DFE of both components. We then calculated the selection coefficients of ϵ by the addition method in (A). P_m represents P_1 in (A), whilst the phenotypes without the added ϵ represents P_2. We then plotted the joint distribution of s across all sampled ϵ.

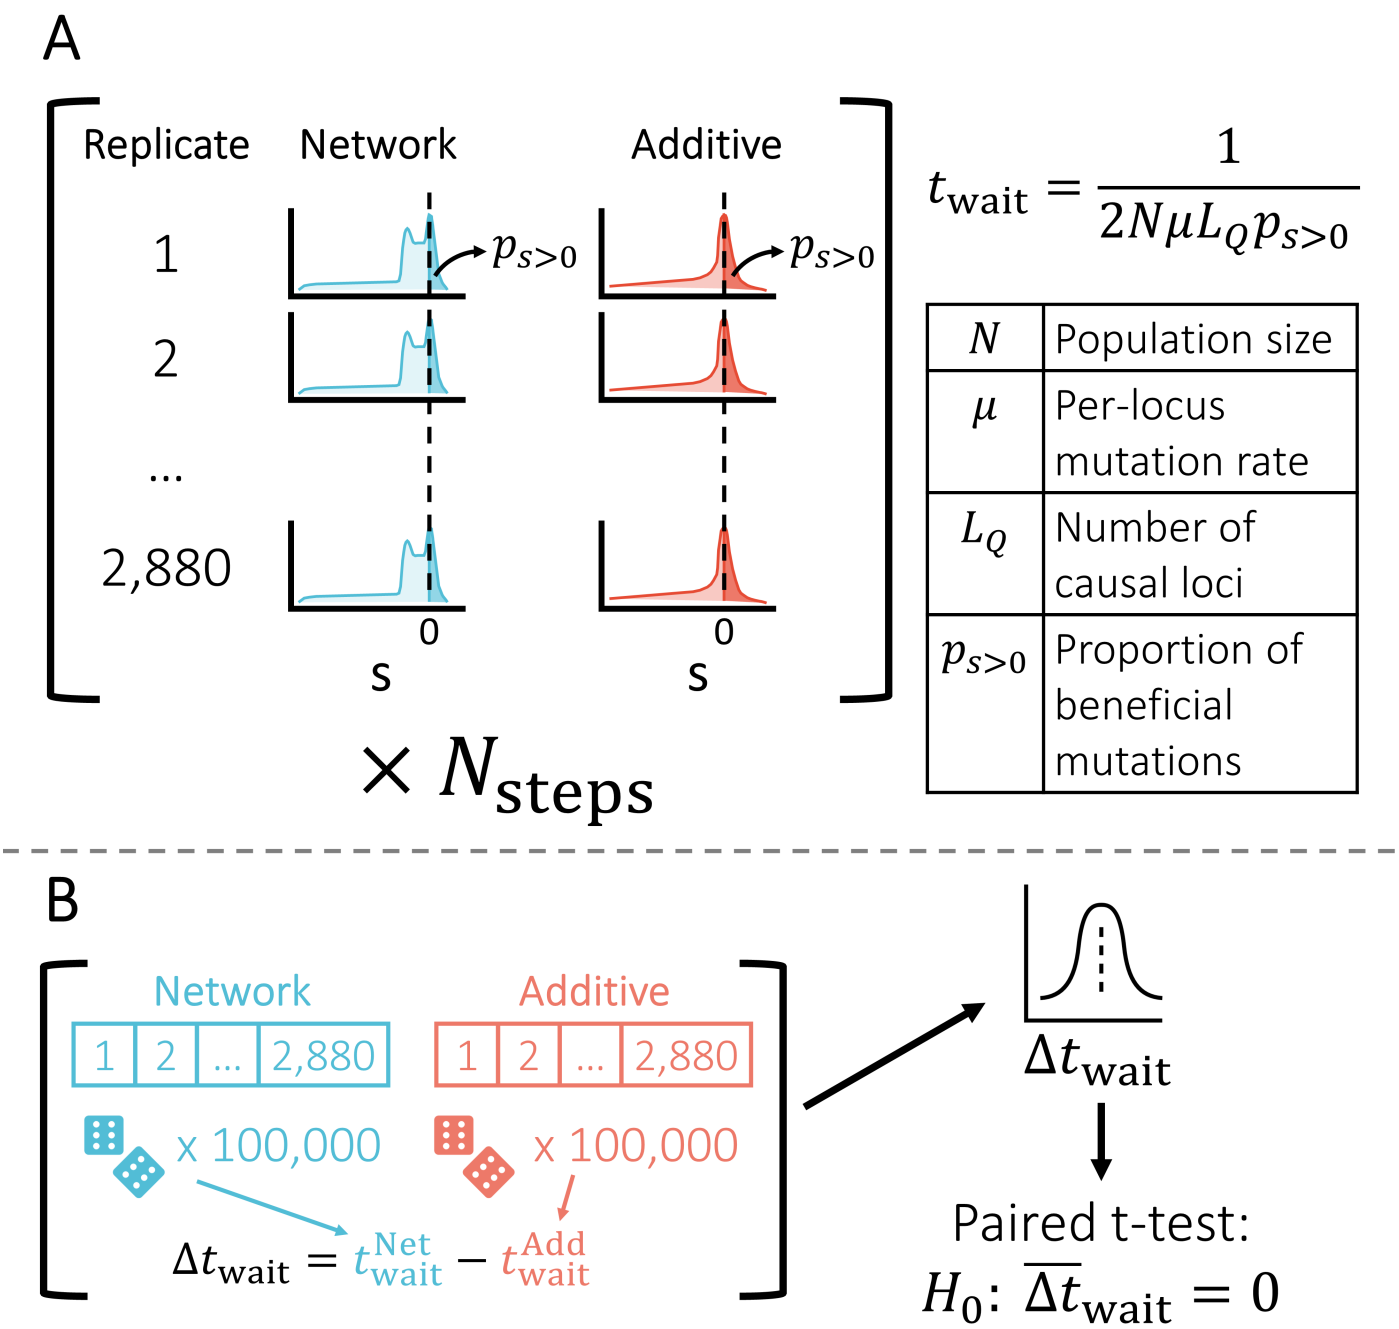

Supplementary Figure 3 Flowchart describing how waiting times to new beneficial mutations were calculated. (A) For each of the 2,880 replicates, we took the distribution of s calculated during the mutation screen experiments and extracted the proportion of new mutations with $s > 0$, $p_{s>0}$. This is the area under the curve shown in dark blue/red (to the right of the dashed line). This was repeated for each adaptive step (N_{steps} within a replicate (i.e. replicates with more than one adaptive step had $p_{s>0}$ calculated for each adaptive step). The waiting time to a new beneficial mutation for a given replicate at a given step was calculated as $t_{\text{wait}} = 1/(2N\mu L_Q p_{s>0})$, where N is the population size (multiplied by 2 because we simulate a diploid organism), μ is the per-locus, per-generation mutation rate, and L_Q is the number of causal loci. (B) To calculate the difference between models in the waiting time to a new beneficial mutation, we ran a bootstrap analysis. We sampled 100,000 random pairs of additive and network models calculating t_{wait} for each. We calculated the difference between their waiting times, Δt_{wait}. The random sampling gave a distribution of differences between network and additive waiting times. We used a paired t-test to determine if the mean difference in waiting times between models, $\bar{\Delta} t_{\text{wait}}$ was not zero (and hence there was a difference between models).

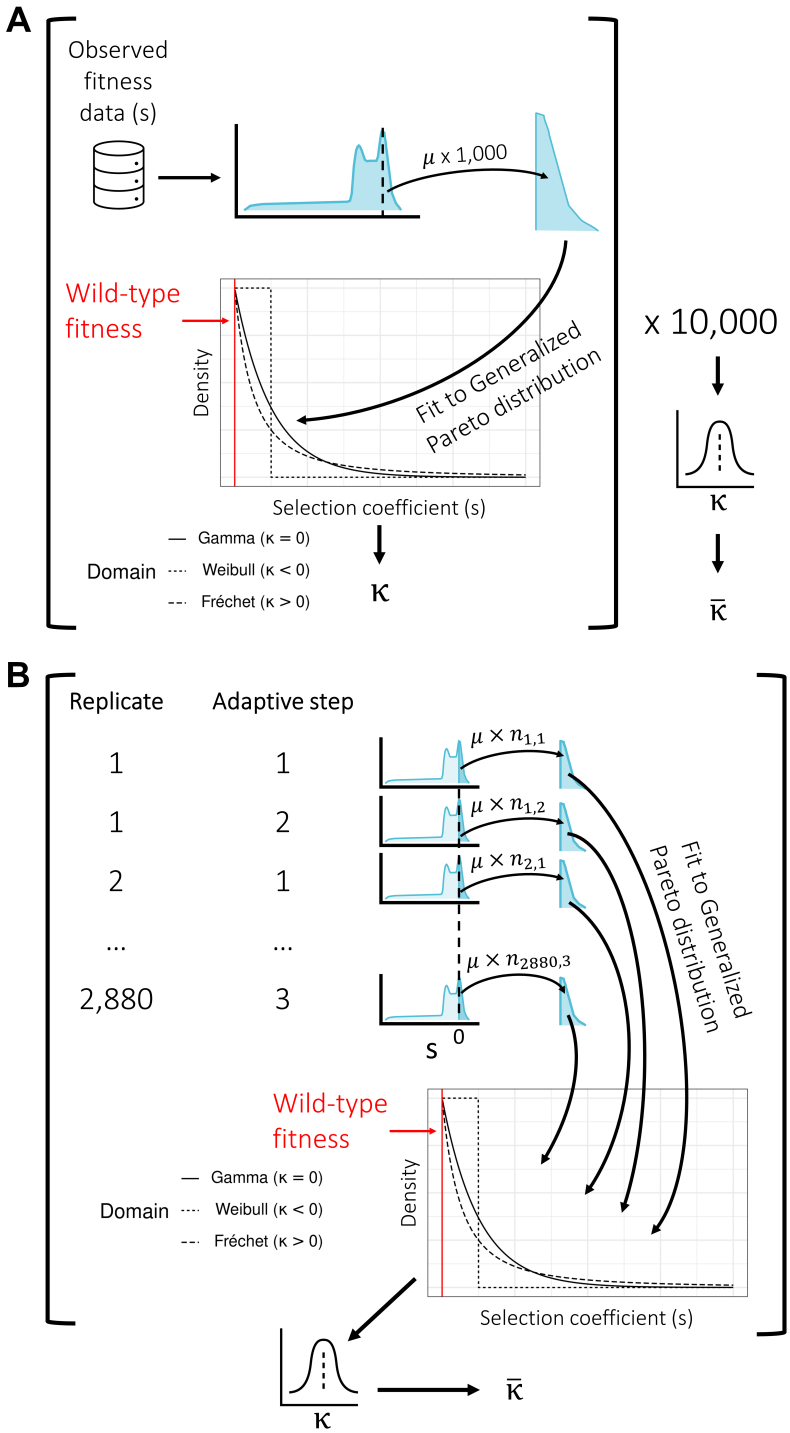

Supplementary Figure 4 Flowchart describing how Beisel *et al.*'s (2007) method was used to estimate the shape of the distribution of fitness effects (DFE) among beneficial mutations via two approaches. (A) Using the pooled distribution of s across all simulations and adaptive steps that we calculated during the mutation screen experiments, we sampled 1,000 beneficial mutations (where $s > 0$) to generate a distribution of beneficial mutations. We then fit this to a generalized Pareto distribution (GPD) using the method outlined by Beisel *et al.* (2007). From the fit, we extracted the shape parameter of the fit, κ. This process was repeated 10,000 times to generate a distribution of κ. We treated the mean κ, $\bar{\kappa}$ as the estimate for the shape parameter. (B) We sampled $n_{x,y}$ mutations from each simulation's s distribution at each adaptive step. The s distributions were generated during the mutation screen experiments. $n_{x,y}$ was $\min(n^{x,y}_{s>0}, 100)$, where $n^{x,y}_{s>0}$ was the number of mutant screen alleles with $s > 0$ in replicate x and adaptive step y. The sampled distributions were fit to a GPD as per (A), generating a distribution of κ from each replicate and adaptive step. The mean was again treated as the estimate for the shape parameter.

Supplementary Figure 5 Distribution of generations at which mutations fixed in adapting additive (red) and network (blue) populations at each adaptive step. There was no significant difference between populations at adaptive step 1. Generations are given relative to the optimum shift so that Generation 0 is when the optimum shifted.

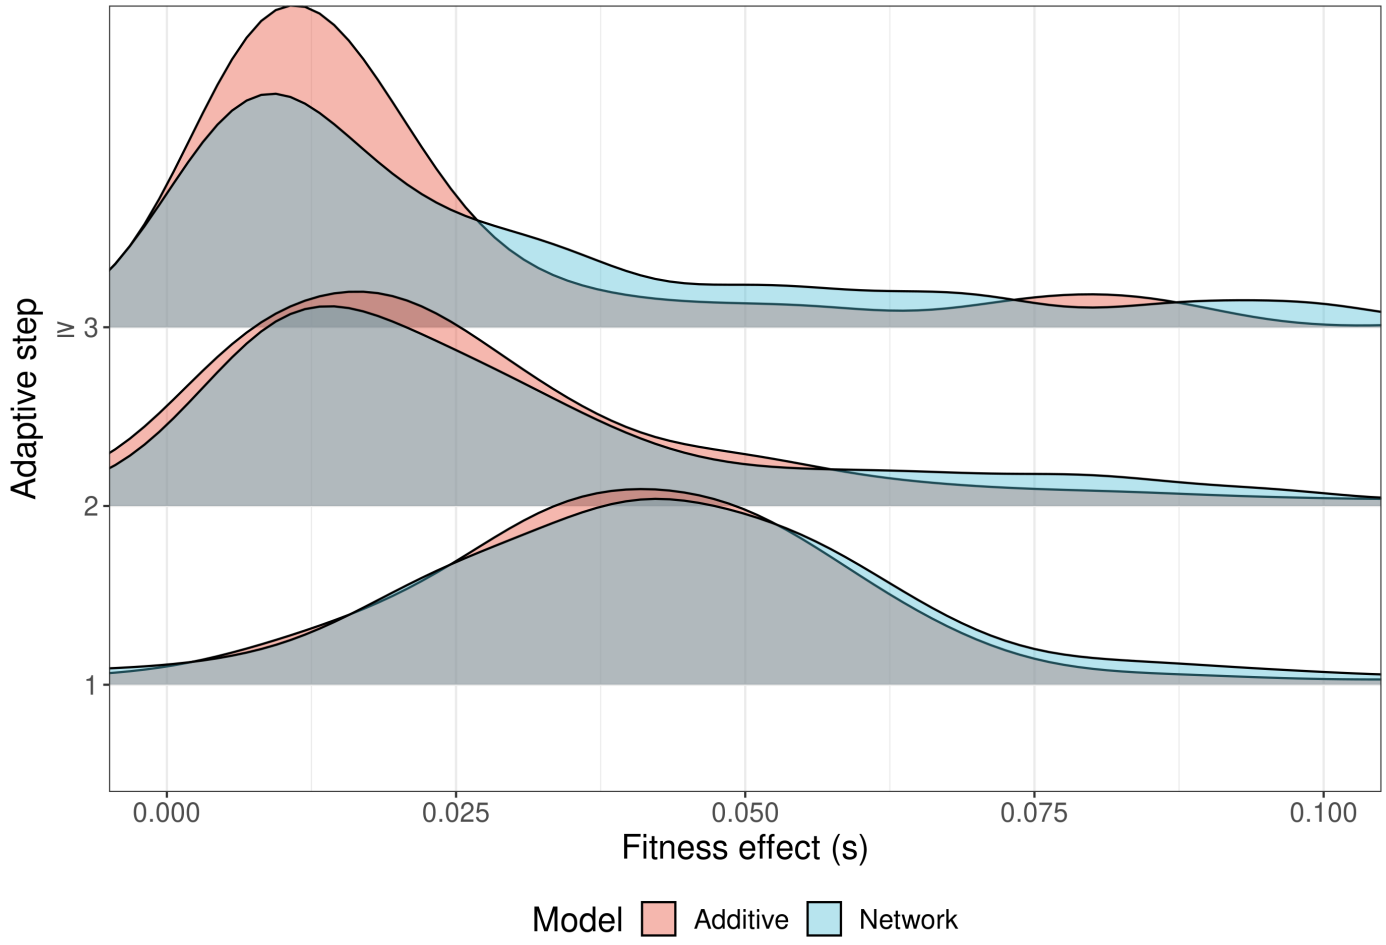

Supplementary Figure 6 Zoomed view of Fig 5B. The mode of the distribution is clearly greater than 0, possibly due to small effect alleles being more susceptible to loss by drift.

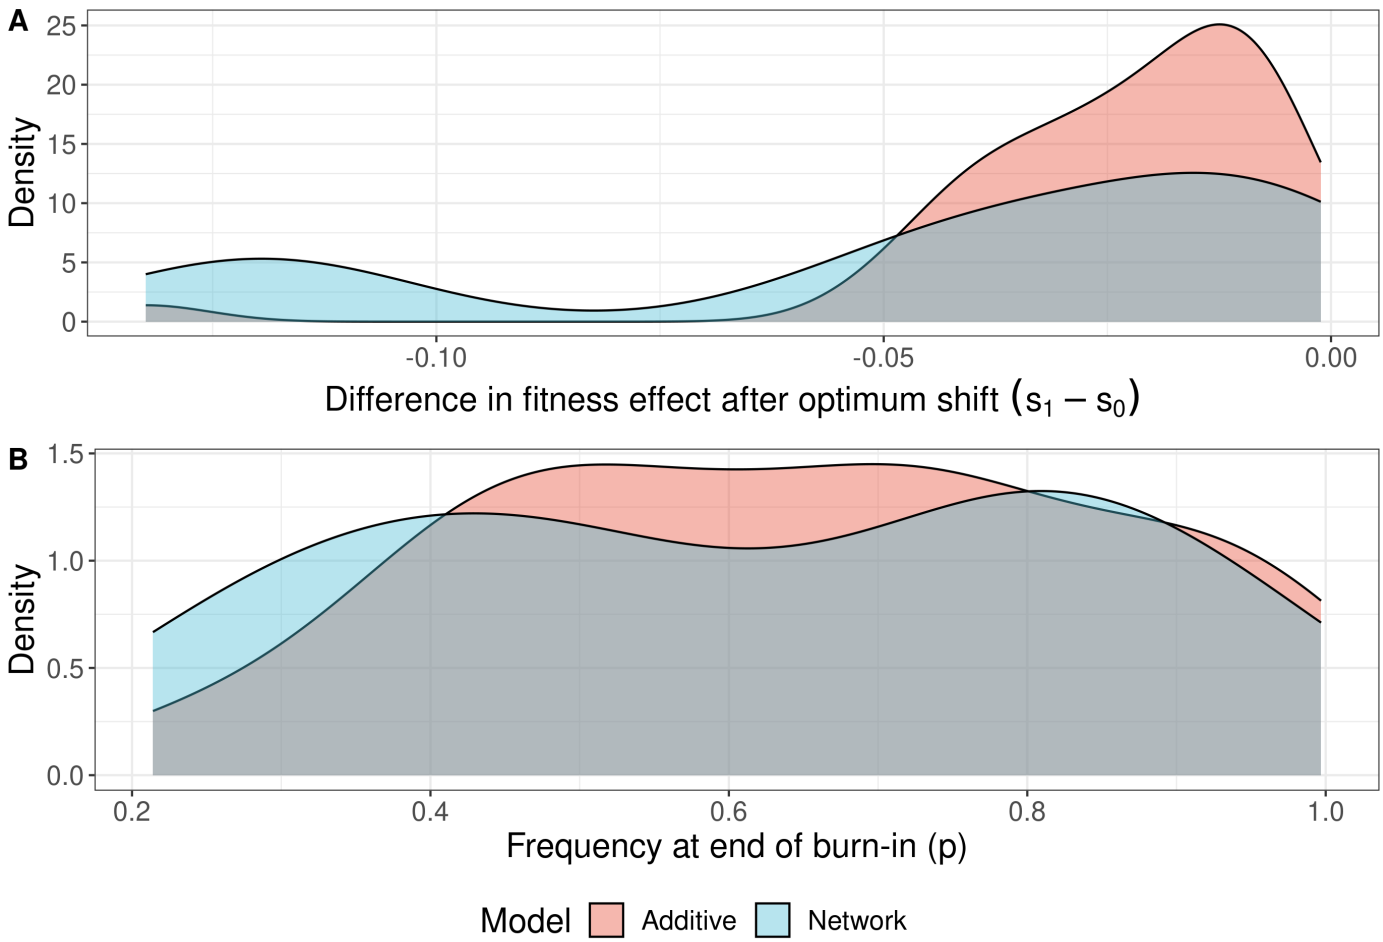

Supplementary Figure 7 Behavior of deleterious fixations in additive (red) and network (blue) models. (A) The selection coefficient, s, decreases after the optimum shift at generation 50,000. (B) The allele frequency, p, of the fixations at the optimum shift.

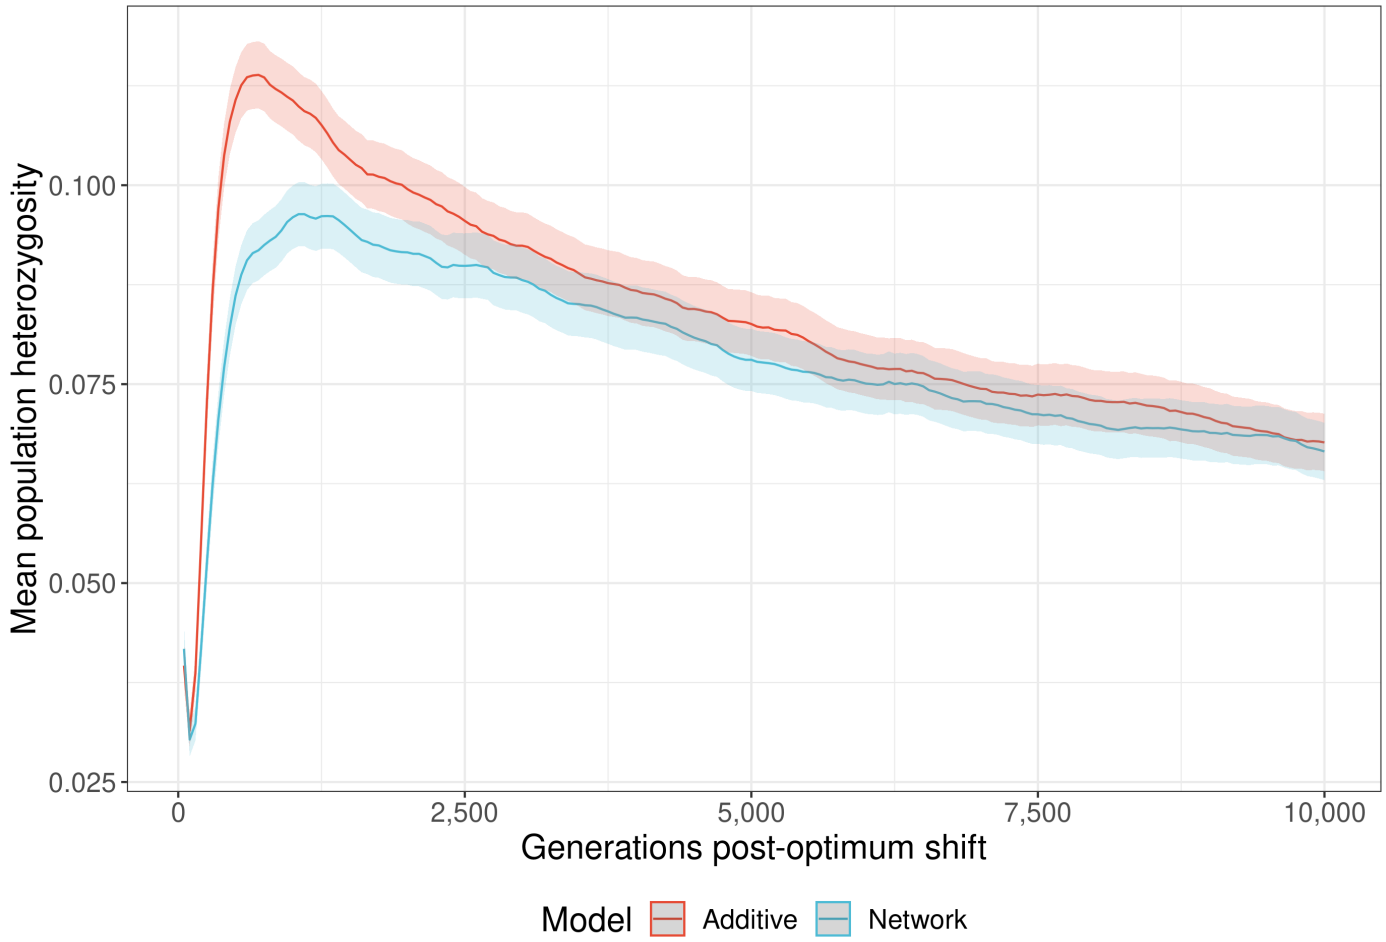

Supplementary Figure 8 Mean observed heterozygosity in additive (red) and network (blue) populations over the adaptation period. The mean is taken over 2,880 replicates per group. Heterozygosity was measured at the two QTLs (mQTLs in NAR models) in both models.

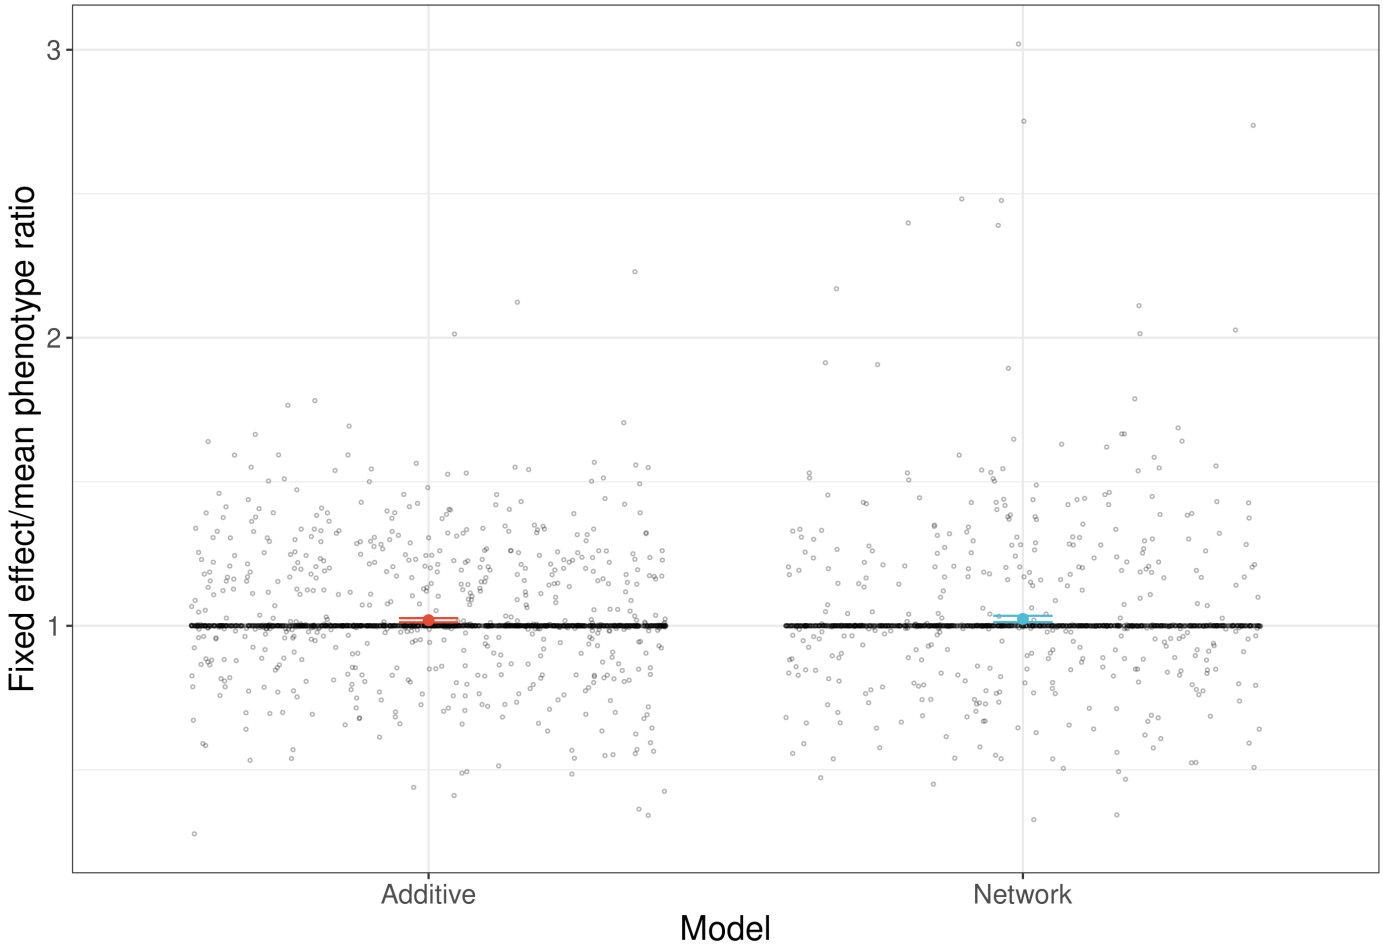

Supplementary Figure 9 Ratio of fixed effect phenotypes to mean population phenotype in additive and network models. Values > 1 indicate that segregating variation decreases the mean population phenotype, and vice-versa for values < 1. Values $= 1$ indicate no segregating variation contributing to trait variance in the population.

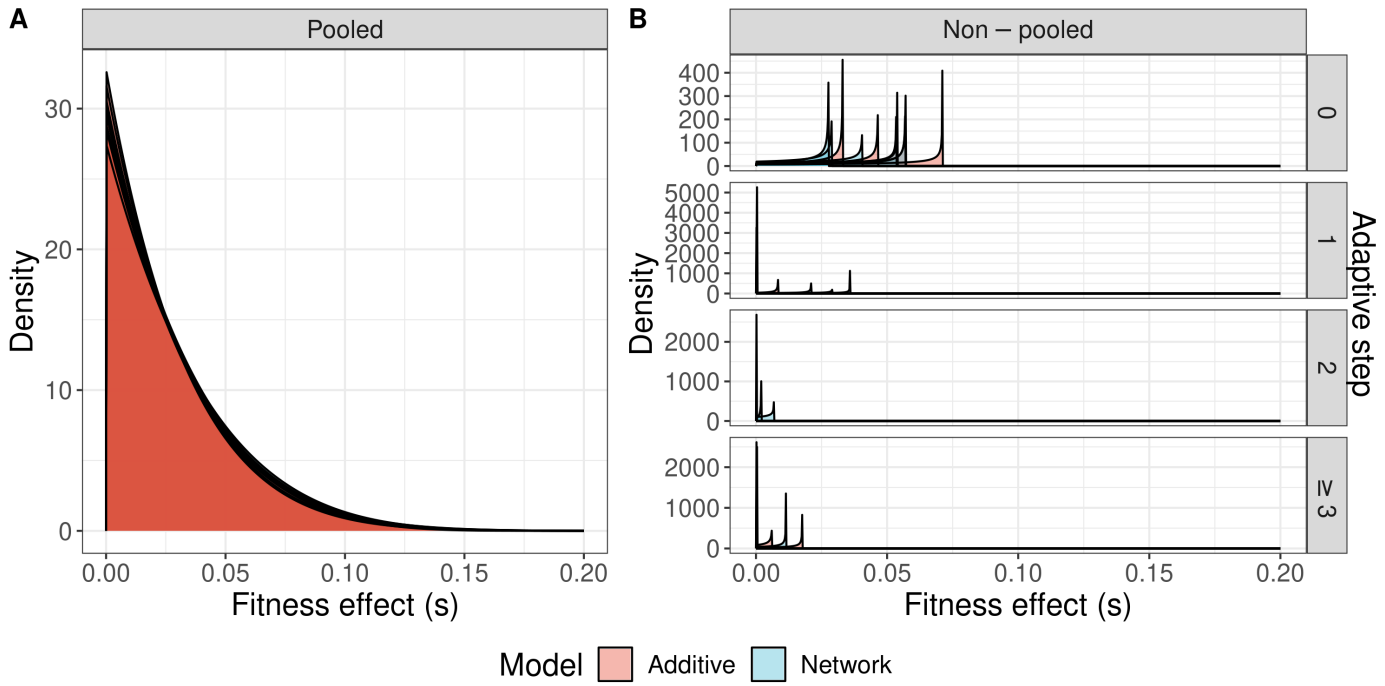

Supplementary Figure 10 Comparison of sampled generalized Pareto distribution (GPD) fits to the distributions of fitness effects (DFE) among beneficial mutations fit onto mutations sampled by two methods. Mutations were created by a mutant screen experiment where 1,000 mutations were randomly generated and their fitness effects measured relative to the phenotype of a population at a given adaptive step. The pooled method (A) pooled together the mutations from all mutant screens across simulation replicates and adaptive steps and used bootstrapping to repeatedly sample 1,000 alleles from the pooled distribution to fit a GPD onto. The non-pooled method (B) fit the GPD onto a sample of 100 mutant screen alleles from each replicate simulation at each of its adaptive steps. Each figure shows the resulting GPD fits of 5 randomly sampled replicates per model of these two methods (i.e. either 5 simulations in the non-pooled case or 5 repeated samples of 1,000 alleles for the pooled case). Note that fewer than 5 samples existed for adaptive steps > 1 in (B), owing to the rarity of walks that long.

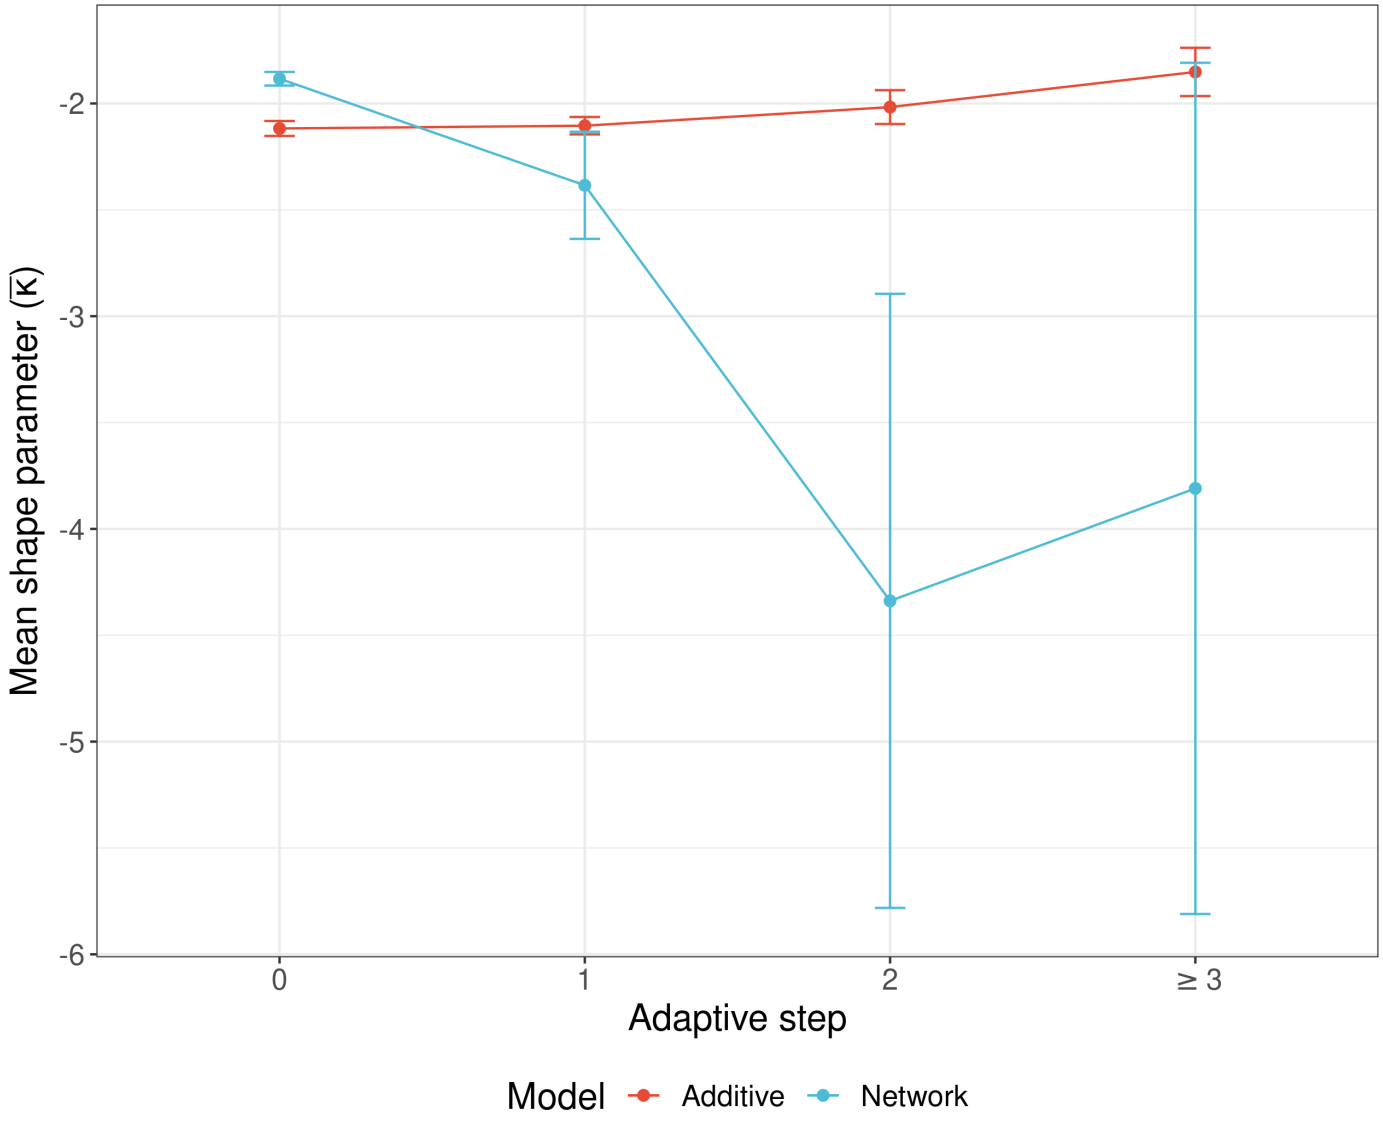

Supplementary Figure 11 Comparison of mean shape parameters, κ, from generalized Pareto distribution (GPD) fits to beneficial alleles randomly sampled from each simulation's mutant screen experiment between models and at each adaptive step. κ describes the shape of the GPD, with negative values (seen here) indicating a Weibull domain of attraction. As κ decreases, the maximum size of a beneficial mutation decreases. Error bars are 95% confidence intervals. Sample sizes for each group are given in Supplementary Table 2.

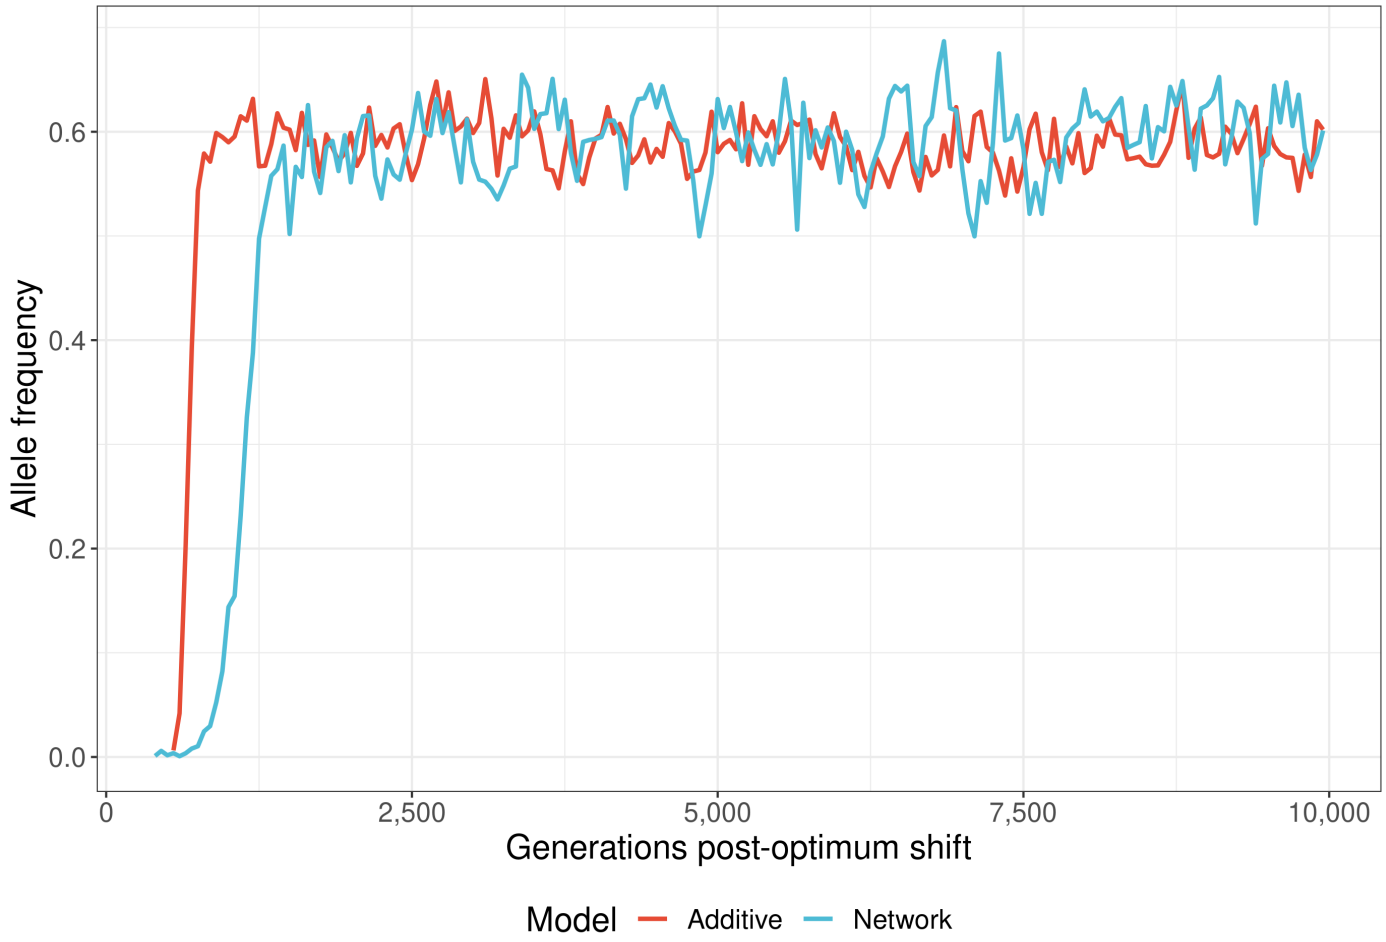

Supplementary Figure 12 Two examples of alleles under balancing selection during adaptation in both additive (red) and network (blue) models. The additive mutation had a phenotypic effect $\alpha = -0.618$ while the NAR allele had a phenotypic effect $\alpha = -0.689$

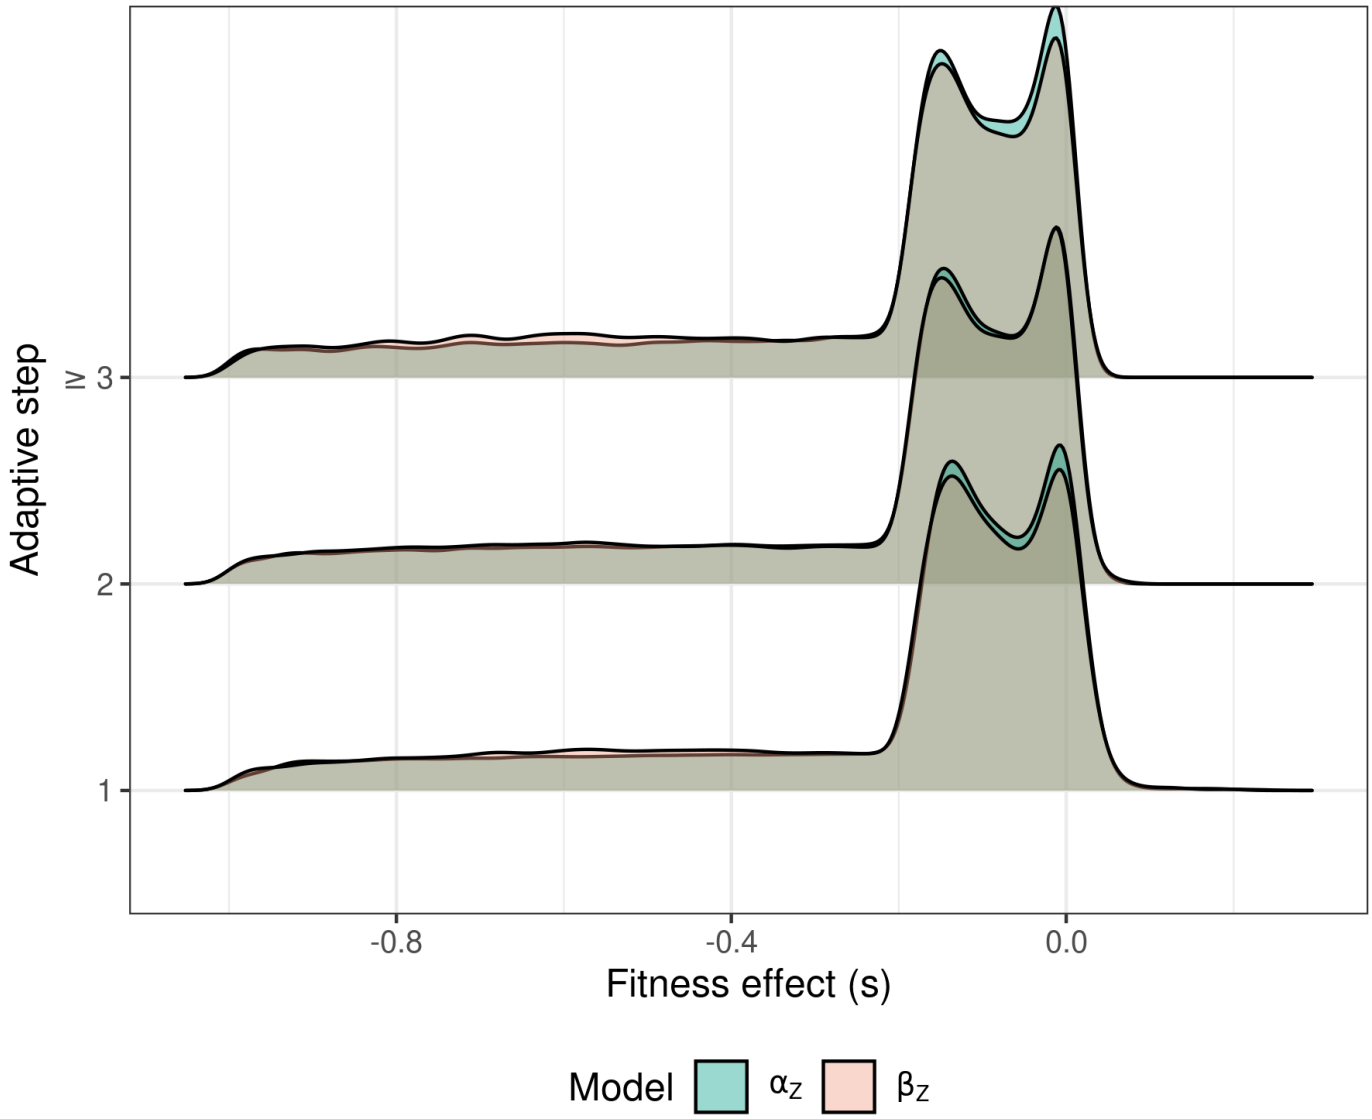

Supplementary Figure 13 The distribution of fitness effects among new mutations in network models for both α_Z and β_Z molecular components. Compared to β_Z mutations, α_Z mutants were less likely to have strongly deleterious effects and more likely to have slightly deleterious effects existing on one of the two modes of the distribution. The shape of the distribution was similar between the molecular components.