

Distributions of Beneficial Fitness Effects in RNA (Submitted to GENETICS)

Matthew C. Cowperthwaite*, J. J. Bull*[†], Lauren Ancel Meyers*^{†‡}

December 20, 2004

*Institute for Cellular and Molecular Biology, University of Texas, Austin, TX 78712

[†]Section of Integrative Biology, University of Texas, Austin, TX 78712

[‡]Santa Fe Institute, Santa Fe, NM 87501

Running Head: Distributions of Beneficial Mutation Sizes

Key Words: RNA, Beneficial Mutations, Extreme-Value Theory, Fitness, Fitness Landscape

Corresponding Author:

Lauren Ancel Meyers

Section of Integrative Biology
University of Texas
Austin, TX 78712
(512) 471-4950 (ph.)
(512) 471-3878 (fax)
laurenmeyers@mail.utexas.edu

Abstract

Beneficial mutations are the driving force of evolution by natural selection. Molecular biology has provided insights into the causes and types of mutations, as well as the biochemical and functional bases of specific beneficial mutations. Yet, relatively little is known about the distribution of the fitness effects of beneficial mutations in populations. Recent work of Gillespie and Orr suggested some of the first generalizations for the distributions of beneficial mutations and, surprisingly, they are partially biology independent. In particular, the theory suggests beneficial mutations obey an exponential distribution of fitness effects, with the same exponential parameter across different regions of genotype space, provided only that there are few possible beneficial mutations available to that genotype. Here we tested this hypothesis with a quasi-empirical model of RNA evolution in which fitness is based on the secondary structures of molecules and their thermodynamic stabilities. The theory was supported in local regions of genotype space when mildly beneficial mutations were ignored. However, the theory neither generalized to the full distribution of beneficial mutations in local regions of genotype space nor did the theory generalize across genotype space. Although in conflict with the current theory, these results suggest that more complex statistical generalizations about beneficial mutations may be possible.

The distribution of the fitness effects of beneficial mutations is of special interest in evolutionary biology, as it profoundly influences the rate and course of adaptation. In turn, adaptive dynamics influences competition, the propensity toward extinction and maintenance in communities, speciation and a plethora of other macro-evolutionary processes. It seems almost a truism that the array of beneficial fitness effects must depend idiosyncratically on the biological details of an organism and its environment. Nonetheless, population geneticists have begun to derive generalities describing these distributions that may be at least partly independent of biology.

GILLESPIE (1983) offered the beginnings of a general theory for the distribution of beneficial fitness effects with the following argument: if the wild-type allele is sufficiently fit, then it resides far in the right-hand tail of the distribution of mutational effects. Any beneficial mutations lie further in the tail, hence their distribution falls in the domain of extreme-value theory (EVT) from statistics. Extreme-value theory tells us that if the underlying distribution of allelic fitnesses is “well-behaved” (See LEADBETTER *et al.* (1983) for a detailed treatment) in several respects, then the spacings between the highest fitnesses in an appropriately large random sample are independent, exponentially-distributed random variables (GUMBEL 1958; WEISSMAN 1978). Therefore, if one assumes that the few beneficial mutants of a high-fitness wildtype allele are a random sample from an underlying distribution of allelic fitnesses, then, when the mutant alleles are rank ordered by size, the spacings between the consecutive beneficial alleles should be approximately exponential.

ORR (2003, 2002) expanded upon Gillespie’s work and derived 2 potentially important corollaries: (1) the distribution of beneficial fitness effects (that is, the difference between the mutant fitness and the wildtype fitness) is exponential and (2) wildtype genotypes differing in the number beneficial mutations accessible by a single mutation (henceforth, “one-step beneficial mutations”) have nearly identical distributions of beneficial fitness effects. These properties were proposed to be general for all evolving systems, provided that the fitness function falls under the purview of EVT and the fitness of the wildtype genotype is greater than almost all mutant alleles. Gillespie and Orr proposed that these are reasonable assumptions for populations that have recently experienced an environmental shift, which has caused the previously optimized wildtype to become slightly suboptimal.

Experimental tests of these theories are extremely difficult to conduct, because one must measure the fitness of all beneficial mutations for a large number of genotypes. Nonetheless, several groups have recently put forth significant efforts to characterize the distributions of beneficial fitness effects in experimental populations of bacteria and viruses (IMHOF and SCHLOTTERER 2001; ROZEN *et al.* 2002). However, as pointed out by ORR (2003), these experiments are not able to test the theory in a comprehensive. The approach used by SANJUAN *et al.* (2004) is perhaps the most promising method to directly test the Orr-Gillespie theory because known point mutations of a viral clone were constructed *in vitro*. Yet, despite an incredible empirical effort by the respective groups, all of these studies utilize a relatively small number of genotypes, which limits their statistical power.

Given the potential generality of the Gillespie-Orr theory, it seemed important to conduct a rigorous test of its predictions. Here we describe a quasi-empirical approach in which we computationally estimate the fitness of RNA molecules based on the similarity of predicted secondary structures to target structures. This system is a computationally tractable and biologically grounded model system that has previously provided insights into evolutionary

dynamics and fitness landscapes (HUYNEN *et al.* 1996; FONTANA and SCHUSTER 1998; VAN NIMWEGEN *et al.* 1999; ANCEL and FONTANA 2000; WILKE and ADAMI 2001; MEYERS *et al.* 2004). In this study, we measure the fitnesses of thousands of genotypes. We find that the full distribution of fitness effects deviates from exponential and varies across sequence space. Yet, consistent with the theory, large regions of sequence space exhibit similar distributions of fitness effects for large-effect beneficial mutations.

MODEL

RNA folding: In many systems, molecular shape is the most important component of function, and hence fitness. Single-stranded RNA molecules carry electrostatic charges that cause them to fold into functional, three-dimensional shapes (tertiary structure). RNA three-dimensional folding is still poorly understood. Yet, the secondary structure of an RNA molecule, which provides the primary scaffold for tertiary structure, is relatively well understood and can be rapidly predicted. Secondary structure results from the formation of complementary base pairs and can be reliably predicted for arbitrary short molecules based on free-energy minimization (WATERMAN 1978; NUSSINOV and JACOBSON 1980; ZUKER and STIEGLER 1981). Two limitations of this approach must be noted: (1) free-energy minimization may not be the only force driving secondary structure formation, and (2) pseudo-knots, a common secondary structural motif, are disallowed because their formation is poorly understood. For this study, we used the dynamic programming implementation in the Vienna RNA package for these calculations (HOFACKER *et al.* 1994).

In particular, we estimated the set of lowest free energy structures of an RNA molecule using an extension (WUCHTY *et al.* 1999) of standard thermodynamic prediction algorithms (WATERMAN 1978; NUSSINOV and JACOBSON 1980; ZUKER and STIEGLER 1981; ZUKER 1989). We refer to this ensemble of low free energy shapes as the *suboptimal repertoire* of a molecule, which is estimated by suboptimal folding. Suboptimal folding ignores energy barriers among alternative states and assumes that a molecule equilibrates among all shapes with free energy within $5kT$ of the ground state, where k is the Boltzmann constant and T is the temperature. This is approximately equivalent to 3 kcal/mol at 37° C, and corresponds to the breaking of 2 G-C/G-C stacking interactions (base pairs). We used the Boltzmann factor to estimate the probability of any particular shape in the suboptimal repertoire of an RNA molecule. For any specific shape σ , the Boltzmann probability of σ , $p_\sigma = e^{-\Delta G_\sigma/kT}/Z$, measures the relative stability of σ with respect to the entire repertoire. Z is the partition function (MCCASKILL 1990) of a molecule and is computed thus:

$$Z = \sum_{\sigma} e^{-\Delta G_\sigma/kT} \tag{1}$$

where ΔG_σ is the free energy of σ and the sum includes all shapes in the suboptimal repertoire. Assuming equilibration, p_σ estimates the probability of finding σ in a large sample of identical RNA molecules and approximates the amount of time any given molecule spends in σ . The minimum free energy conformation is the most probable shape in any suboptimal repertoire.

For any sequence, we can thereby rapidly compute its suboptimal repertoire and the approximate probability of each shape in the repertoire. This constitutes a biologically-

grounded map from genotype (sequence) to phenotype (shape ensemble).

Measuring fitness: Our computational RNA genotype-to-phenotype model is able to accommodate a variety of biologically realistic fitness functions. For example, RNA molecules have been selected experimentally to bind a ligand with high affinity (ELLINGTON 1994). We cannot yet explicitly model such binding interactions, but we can approximate such systems by assuming there exists an ideal secondary structure and the nearer the shape ensemble of a molecule is to that ideal the better it will bind (SCHUSTER *et al.* 1994; ANCEL and FONTANA 2000). In our model, at equilibrium, a fraction p_σ of a large number of identical sequences assumes shape σ and binds to a ligand with a corresponding constant.

For each shape in the suboptimal repertoire, we used a hyperbolic decaying function $f(\sigma)$ to calculate a selective value based on how well σ matched a target shape:

$$f(\sigma) = \frac{1}{\alpha + (d(\sigma, t)/L)^\beta} \quad (2)$$

where α and β are scaling constants, $d(\sigma, t)$ is the Hamming distance between the current shape and the target shape, and $L = 76$ is the length of the sequence. The value $\alpha = 1$ parameter was chosen to scale the fitness values between ≈ 1 and 100; $\beta = 1$ was chosen to produce the hyperbolic decaying shape of the selective-value function and maintain consistency with prior work (FONTANA and SCHUSTER 1998; ANCEL and FONTANA 2000). By scaling the distance with a hyperbolic decaying function, we modeled strong selection for target structure.

We chose a nucleotide sequence length of 76 for several reasons. This length has 228 one-step mutants, which should be sufficiently large for EVT to apply GILLESPIE (1983). The free-energy minimization algorithms are most accurate for short sequences, thus our results will not be confounded by folding errors. This length also gives us computational tractability - we can measure the fitness of every one-step mutant sequence in a reasonable time. Finally, most tRNA molecules in natural organisms are approximately 76 nucleotides in length.

The overall fitness, W , of a molecule is the average of the selective values of the shapes in its suboptimal repertoire, each weighted by its probability, $W = \sum_\sigma f(\sigma)p_\sigma$. The range of fitness values (W) possible given our choice of parameters is 0.99 - 100. This function simultaneously considers secondary structure and thermodynamic stability such that the highest fitness molecules will be those that fold stably into minimum free energy shapes that look much like the target shape. This fitness function is essentially continuous because no two sequences have identical suboptimal repertoires.

Obtaining low-rank genotypes: The rank of a wildtype allele (i) is simply its position in a set of fitnesses that are rank ordered from 1 (most fit) to $m + 1$ (least fit), where m is the number of single-mutant sequences (ORR 2003). The Gillespie-Orr theory depends on the fitness of the wildtype allele being higher than nearly all of its 1-step mutants, that is, it is based on genotypes with very few 1-step mutations of higher fitness ($i \ll m + 1$). Our results are based on large samples of low-rank genotypes, which are relatively rare and therefore difficult to find by random sampling. Thus, we generated samples of low-

rank genotypes using adaptive walks. We emphasize that these walks were not intended to simulate biological evolution, but simply served as a heuristic for locating appropriate sequences for our study.

Adaptive walks were initiated with random sequences with no base-composition bias. We refer to the sequence at the current step of an adaptive walk as the wildtype sequence. At each step of the walk, the fitness of every one-step mutant of the current wildtype sequence was measured as described above. We randomly selected a single one-step *beneficial* mutant sequence to be the next wildtype sequence. The process was repeated until the wildtype sequence arrived at a local optimum (i.e. no mutations were beneficial). A single wildtype allele of each rank class was selected at random from each adaptive walk to obtain a set of suitable low-rank genotypes. The shape of the distributions of beneficial fitness effects we obtained was robust to two different types of adaptive walks (randomly selected beneficial mutant vs select the best mutant, data not shown), thus our results appear robust to the choice of an adaptive walk model.

Generating high-fitness sequences: We generated sets of high-fitness molecules using an algorithm that produces sequences which specifically fold into a particular secondary structure. The program, ‘RNAinverse’ in the ViennaRNA package, initially divides the target shape into several smaller regions and the starting sequence into segments, which each correspond to a small region of the target structure (HOFACKER *et al.* 1994). Each segment of the starting sequence is individually optimized through single base changes or compatible base-pair changes. Once all of the separate regions of the starting sequence have been individually optimized the full sequence is created and further optimized. This results in molecules that fold into the specified minimum free energy structure, but may or may not have a high degree of thermodynamic stability.

Estimation of exponential parameters: GILLESPIE (1983) and ORR (2003) proposed that the distribution of absolute fitness differences among the few fittest alleles will follow an exponential distribution (density $\lambda e^{-\lambda x}$, where λ is the exponential parameter characterizing the distribution). The mean μ of the distribution is $1/\lambda$ and can be estimated by maximum likelihood as

$$\hat{\mu} = \frac{\sum_n (W_j - W_i)}{n} \quad (3)$$

where W_j is the fitness of the j -th beneficial mutation, W_i is the fitness of the wildtype allele and n is the number of alleles, such that $W_j > W_i$. Since the estimate of the exponential parameter as $1/\hat{\mu}$ is biased, we work with the estimate of the mean, which has the advantage of being more interesting biologically than its reciprocal. Another useful property of an exponential distribution is that a log-linear plot of the total observations greater than x yields a straight line, and deviations from exponential are thus easily observed in such a plot [CDF: $P\{X > x\} = e^{-\lambda x}$; $\ln(e^{-\lambda x}) = -\lambda x$].

ORR (2002) claimed that distributions of *fixed* beneficial fitness effects in actual biological systems may deviate from exponentiality at the left end (small benefit mutations) but obey exponentiality in the right end (large benefit mutations). Left truncation of the distributions of *new* beneficial fitness effects may therefore yield the exponential property even though the

full distribution may not. To estimate μ of the full distribution from a truncated distribution, we first compute the mean μ_T of the truncated distribution

$$\mu_T = E(x) = \frac{\int_T^\infty x \lambda e^{-\lambda x} dx}{\int_T^\infty \lambda e^{-\lambda x} dx} = \mu + T \quad (4)$$

where T is the truncation threshold. Thus, an unbiased estimate of the mean $\hat{\mu}$ of the full distribution is

$$\hat{\mu} = \mu_T - T \quad (5)$$

If the full distribution is exponential, then $\hat{\mu}$ is unaffected by truncation when corrected in this fashion.

RESULTS

Gillespie and Orr proposed that EVT could be applied to describe the distribution of fitness effects of beneficial mutations to high-fitness genotypes. The use of EVT rests on several assumptions: (1) allelic fitnesses are drawn from an underlying well-behaved distribution of allelic fitnesses, (2) the one-step mutants of a genotype are an i.i.d random sample from this distribution, and (3) the wildtype allele lies well into the right-hand tail of the underlying distribution, thus the fitness effects of beneficial one-step mutations of the wildtype allele will be further in the tail. We therefore set out to rigorously test the fundamental predictions of the theory.

The RNA landscape obeys EVT: The underlying distribution of fitness values was analyzed to determine if the fitness space is “well-behaved” with respect to EVT. We produced a set of approximately 3.6 million random sequences and measured the fitness of each sequence in the set. The distribution of fitnesses in this set of genotypes shows a strong peak at $W \approx 1.2$ and the fraction of sequences with $W > 3.0$ is less than 10^{-4} (Figure 1). Any sequences with $W > 3.0$ would be expected to be sufficiently far into the tail to be in the domain of EVT.

A fitness difference Δ_i is the absolute fitness difference between the alleles of rank i and $i + 1$ in a set of allelic fitnesses ranked from 1 (most fit) to N (least fit) (ORR 2003). For the top few i , EVT predicts Δ_i to be asymptotically exponentially distributed and $E(\Delta_i) = E(\Delta_1)/i$. The set of 3.6 million sequences was randomly divided into 15,880 subsets of 229 sequences. The number 229 was chosen because it is the number of one-step mutants plus the wildtype allele of a 76-nucleotide sequence, which we consider in the adaptive walks discussed below. We measured Δ_1 , Δ_2 and Δ_3 for each set of 229 random sequences. The inset to Figure 1 confirms the exponential distribution of Δ_1 , Δ_2 and Δ_3 . Δ_1 was found to be 1.99 and 2.87 times the size of Δ_2 and Δ_3 , respectively, which are close to the expected values 2 and 3, respectively.

Importantly for this study, therefore, the fitness landscape is “well-behaved” with respect to EVT. Thus, if one-step mutational neighborhoods are essentially random samples of sequences, the distribution of beneficial fitness effects should be similar to that found for sets of random sequences.

Distribution of fitness effects with random starting points: The rank (i) of a genotype is defined as the position of that genotype in a set of allelic fitnesses ranked from 1 (most fit) to m (least fit), where m is the number of single-mutant sequences (ORR 2003). An allele of rank i has $i - 1$ one-step beneficial mutations. The distribution of beneficial mutations was analyzed using sequences of rank $i \leq 4$, to be confident that they would be in the domain of EVT. The wildtype genotypes were generated with adaptive walks beginning with random starting genotypes. The mean starting fitness of the random sequences was $1.1 (\pm 0.002 \text{ SE})$. The mean ending fitness was $3.4 (\pm 0.01 \text{ SE})$. The distributions of beneficial effects were produced using genotypes from 5721 adaptive walks that attained a final absolute fitness between 3 and 9. The average walk accrued $84.4 (\pm 0.28 \text{ SE})$ substitutions before reaching a local optimum.

The set of genotypes used to estimate the distribution of beneficial fitness effects was produced by randomly selecting a single sequence for every rank class $i \leq 4$ from each adaptive walk. This produced a unique data set for each rank class and ensured the statistical independence of the observations within each data set. By measuring $S = W_j - W_i$, the difference between the fitnesses of each high-fitness mutant genotype and the wildtype on an absolute scale, we estimated the distribution of beneficial fitness effects for each rank class. The distributions of beneficial effects deviate from exponential by having an excess of small-sized mutations (Figure 2). For each wildtype rank examined, at least 80% of the beneficial mutations increase fitness by less than 0.01, on an absolute scale.

Inspection of Figure 2 suggests that the distributions may be nearly exponential for the larger S values. Indeed, when the distribution of effects is truncated to $S > 0.2$, this class of mutations appears approximately exponential (Figure 2, inset). However, it must be emphasized that the genotypes with $S > 0.2$ are a very small fraction of the full distribution ($< 0.5\%$). Consistent with Orr’s assertions, the ML estimate of the means of the truncated distributions are statistically indistinguishable for the different rank classes [$p = 0.71 (i = 2, 3)$ and $p = 0.88 (i = 3, 4)$, Wilcoxon-Mann-Whitney]. In contrast to the theory, the means of the full distributions for different ranks are significantly different [$p < 2.2 \times 10^{-16} (i = 2, 3)$ and $p < 2.2 \times 10^{-16} (i = 3, 4)$, Wilcoxon-Mann-Whitney].

Distribution of fitness effects in high-fitness space: The fitnesses attained at the end of the adaptive walks started from random genotypes were low relative to the maximum possible fitness of 100, thus these data do not represent regions of sequence space with high-fitness genotypes. To evaluate mutational effects in high-fitness regions of sequence space, we used inverse folding to generate a large set of sequences with secondary structures that nearly or perfectly matched the target structure. These sequences were used to start 8390 adaptive walks (referred to as “high-fitness walks”). We only considered walks in which the final sequence attained a fitness greater than 20, giving a subset of 6959 adaptive walks. The mean final fitness attained in this subset of walks was $56.71 (\pm 0.21, \text{ SE})$.

Using a single sequence of each rank $i \leq 4$ from each walk, we generated the distribution of beneficial effects as described for the random walks (Figure 3). The entire spectrum of beneficial effects is not exponentially distributed, again because of an excess of small fitness effect mutations. Furthermore, the fitness effects for the mutants of the high-fitness genotypes are on average greater than for the mutants of the random-walk genotypes. The

larger effect mutations resulted in the much faster rate of adaptation of the sequences in high-fitness space: the average size of a fixed mutation was significantly larger for the high-fitness walks than the random walks (high-fitness walks: 1.034 ± 0.004 ; random walks: 0.034 ± 0.0001). The rate of adaptation, however, does not correspond to the rate obtained in a truly evolutionary process, but the comparison of relative rates is meaningful nonetheless.

Deviation from exponential behavior: To compare the distributions of beneficial fitness effects in the different regions of sequence space, we progressively truncated the distributions to determine the minimum threshold required to achieve exponentiality. The motivation for this approach is that the appropriate truncation value is not obvious by inspection of the plots in Figures 2 and 3. Thus, for each distribution, we plotted the estimated mean S value for progressively larger truncation values and identified the point at which the estimate of the mean asymptotes. Any systematic variation in the mean and, by extension, the exponential parameters will appear as variation in the asymptotic values.

The truncated distributions of the two sets of sequences are vastly different (Figure 4). The high-fitness sequences maintain a significantly higher mean fitness effect than the random sequences and become exponential at a significantly higher threshold ($S = 0.20$ for random walks; $S = 10.0$ for high-fitness walks). We have included the comparison between S (top panel Figure 4) and s (bottom panel Figure 4). Neither fitness measure removes the non-exponentiality nor the difference in fitness effects in the two regions of genotype space. Therefore, the distributions of beneficial fitness effects differ among the two regions of sequence space.

Decline in mean s during walks: So far we have considered the absolute difference in fitness (S) between the wildtype and its mutants. Now we consider the relative fitness difference between the genotypes (s), which is defined as the absolute fitness difference between the mutant and wildtype alleles normalized to the absolute fitness of the wildtype allele. We monitored the change in s during an adaptive walk by measuring the mean size of all new beneficial effects in the one-step neighborhood of the wildtype genotypes (Figure 5). We only considered wildtype alleles with $i \leq 4$ for this analysis. The average s in the mutant neighborhood declined during the course of an adaptive walk, demonstrating that small benefit mutations come to dominate the landscape with the approach toward a local optimum. Not surprisingly, the mutants of the "high-fitness" genotypes had higher s values during most stages of the adaptive walks because of the large S values observed.

The important but perhaps obvious result is that the mean of the distribution of beneficial effects declines as the adaptive walk progresses. Early in the course of a walk, large s mutations exist permitting adaptation to proceed quickly. As the adapting sequence approaches a local optima, the possible s values become progressively smaller thus slowing adaptation. Thus, the distribution of beneficial fitness effects changes during the course of an adaptive walk. This is in agreement with the theoretical predictions of Fisher's geometric model of adaptation (FISHER 1930) and empirical studies of viral adaptation (BURCH and CHAO 1999)

DISCUSSION

GILLESPIE (1983) pioneered a theory of adaptation for populations that are displaced from a fitness optima by an environmental change. He argued that the wildtype allele would remain sufficiently far in the extreme right-hand tail of the distribution of allelic fitnesses that the fitness of any beneficial mutants would be within the domain of extreme-value theory. This theory tells us that the differences between consecutive rank-ordered extreme values from a randomly selected set of values, should follow an exponential distribution. ORR (2002, 2003) then used this theory to argue that the distribution of beneficial fitness effects for genotypes in the extreme right-hand tail of the fitness distribution would be exponential, with a single exponential parameter governing all such genotypes (ORR 2003).

This theory is potentially very important for both artificial and natural evolution. It offers a framework for predicting the outcome of adaptation in response to environmental challenges such as pharmaceuticals, pesticides, and herbicides. To date, the predictions of the model have received mixed experimental support (IMHOF and SCHLOTTERER 2001; ROZEN *et al.* 2002; SANJUAN *et al.* 2004), but these types of studies are generally based on small sample sizes of limited statistical power and have additional limitations, as discussed by ORR (2003). These studies are mentioned not to diminish their importance, but rather to illustrate the difficulty in testing the theory.

We have tested this theory using a quasi-empirical model of RNA evolution. RNA secondary-structure prediction by free-energy minimization gives a biologically-realistic map from a genotype (sequence) to phenotype (shape ensemble). We assigned fitnesses to individual RNA molecules based upon biologically-motivated properties of their shape ensembles. No a priori assumptions were made regarding an underlying distribution of allelic fitnesses or fitness correlations among similar sequences.

We found two fundamental departures from the predictions of the theory. First, the distribution of beneficial effects depends on the fitness of the un-mutated genotype; the average size of a beneficial effect increases with the fitness of the parental sequence. Second, for the two fitness classes we evaluated, the distributions of beneficial effects are non-exponential. The distributions appeared exponential only after left truncations that discarded over 99% of the observations. The appropriate truncation thresholds also differed for these two classes of sequences.

A priori, two possible explanations for the discrepancy between the theory and our observations can be proposed. First, the fitness landscape in our quasi-experimental system may not satisfy the prerequisites of extreme value theory. This explanation was rejected through a large random survey of sequences in which the tail of the fitness distribution was shown to have the essential characteristics predicted by EVT (Figure 1). Thus, we turn to a second possible explanation: correlations among closely related sequences defy Gillespie and Orr's assumptions that the one-step mutations of any given genotype have i.i.d. random fitnesses from the distribution of all allelic fitnesses. In other words, the theory assumes that fitness values are distributed completely randomly throughout genotype space.

In our model, most point mutations are nearly neutral because they alter the structural repertoire, and therefore the fitness, of a molecule only slightly. Thus, the fitnesses of a sequence and its one-step mutants are correlated, implying that, on average, the fitness differences between beneficial mutants and their parent sequences will be smaller than expected

if the fitnesses of the beneficial mutants had been i.i.d. random samples from an overall fitness distribution, as the theory assumes. Therefore, the correlation between the fitnesses of parental genotypes and their one-step mutants produced at least part of the discrepancy between our observations and the Orr-Gillespie theory - the excess of small effect mutations.

Fitness correlations among closely related genotypes are certainly not unique to our model. In a closely related model in which the fitness of an RNA sequence is determined by thermostability alone, FONTANA *et al.* (1993) measured fitness correlations among genotypes and their mutants. In particular, they measured the correlation length (ρ), which is the distance d at which the fitnesses of a reference sequence and a d -mutant sequence become essentially statistically independent. They estimated $\rho = 6.25$ for a 70 nucleotide sequence suggesting that, on average, one-step mutants will have similar fitnesses to their parental genotypes, although it does not specifically address correlations between high fitness sequences and their beneficial mutants. Fitness correlations are evident in many other biological systems. For example, if stable RNA structures are important to fitness, then the interactions between the paired bases violate the assumption of independence of separate mutations: a beneficial base pairing could be restored by either of two mutations that would each achieve correlated fitness effects (PARSCH *et al.* 2000). In β -helix-loop-helix (bHLH) proteins, structure-function relationships may produce correlation among the fitness effects of mutations (ATCHLEY *et al.* 2000) and, by extension, produce distinct fitness distributions in local regions of genotype space.

The observed association between high fitness genotypes and large beneficial effects arises from both the correlation structure of the fitness landscape and the shape of the selective-value function and thus may be specific to our model. We used a hyperbolic decaying function because it models the realistic scenario in which most molecular structures are essentially unviable and very few have high fitness. This assumption is supported by remarkable structural conservation for many different classes of RNA (DOUDNA 2000). By virtue of fitness correlations, a sequence with low to medium fitness (like those from the random adaptive walks), will lie near its one-step mutants in the shallow region of the selective value function. Thus any beneficial effects will likely be quite small. In contrast, high-fitness sequences (like those from the high-fitness walks) and their one-step mutants will occupy the steep region of the function, where beneficial effects may be relatively large. A preliminary survey of mutations under a linear selective-value function suggests that a single (non-exponential) distribution of beneficial effects applies to all sequences regardless of wildtype fitness (data not shown). This suggests that the observed anisotropy in beneficial fitness effects is closely linked to hyperbolic shape of the selective-value function.

We offer a few caveats to this study. First, our model includes no selection for function per se, but rather for characteristics known to be important for functional RNA molecules *in vitro* and *in vivo*. Second, our fitness function is bounded both below and above, yet the theory of Orr and Gillespie was developed for unbounded distributions (GILLESPIE 1983; ORR 2003). Since our general survey of RNA sequence space was consistent with EVT (Figure 1) and even the fittest sequences in our study are quite far from the upper bound, we believe that the fitness bounds do not pose a problem. Third, our results are for one model system and other systems may yield different results.

Although this study demonstrates that the current theory might not withstand the complexity of all biological systems, some generality was evident. In particular, the distributions

of fitness effects were monotonically decaying and the general shape of the distributions of beneficial fitness effects was invariant across genotype space, as predicted by the Orr-Gillepsie theory. Furthermore, after discarding the nearly-neutral beneficial mutations, the distributions of the remaining large effects were approximately exponential. This suggests that a more flexible theoretical framework may be possible in the future.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the Center for Computational Biology and Bioinformatics for their support of this research. The authors thank Santiago Elena, David Hall, Paul Joyce, H.A. Orr and Claus Wilke for helpful discussions on the manuscript. The Santa Fe Institute (SFI) and the David and Lucile Packard Foundation program on Robustness at SFI have provided invaluable support to M.C. and L.A.M. (a member of the SFI external faculty). This research was supported in part by a fellowship from an NSF IGERT graduate training grant in Computational Phylogenetics and Applications to Biology to M.C. and research grants to J.J.B. (NIH GM-57756) and to L.A.M. (NSF DEB-0303636).

LITERATURE CITED

- ANCEL, L. and W. FONTANA, 2000 Plasticity, Modularity and Evolvability in RNA. *Journal of Experimental Zoology* **288**(3): 242–283.
- ATCHLEY, W. R., K. R. WOLLENBERG, W. M. FITCH, W. TERHALLE, and A. W. DRESS, 2000 Correlations Among Amino Acid Sites in bHLH Protein Domains: An Information Theoretic Analysis. *Mol Biol Evol* **17**(1): 164–178.
- BURCH, C. L. and L. CHAO, 1999 Evolution by Small Steps and Rugged Landscapes in the RNA Virus phi6. *Genetics* **151**(3): 921–927.
- DOUDNA, J. A., 2000 Structural genomics of RNA. *Nature Structural Biology* **7**(11): 954–6.
- ELLINGTON, A., 1994 RNA selection. Aptamers achieve the desired recognition. *Current Biology* **4**(5): 427–429.
- FISHER, R. A., 1930 *The Genetical Theory of Natural Selection*. Oxford University Press, Oxford.
- FONTANA, W. and P. SCHUSTER, 1998 Continuity in evolution: on the nature of transitions. *Science* **280**(5368): 1451–1455.
- FONTANA, W., P. F. STADLER, E. BAUER, T. GRIESMACHER, I. L. HOFACKER, M. TACKER, P. TARAZONA, E. D. WEINBERGER, and P. SCHUSTER, 1993 RNA Folding and Combinatory Landscapes. *Physical Review E* **47**: 2083–2099.
- GILLESPIE, J., 1983 A simple stochastic gene substitution model. *Theoretical Population Biology* **23**(2): 202–215.
- GUMBEL, E. J., 1958 *Statistics of Extremes*. Columbia University Press, New York.
- HOFACKER, I. L., W. FONTANA, P. F. STADLER, L. S. BONHOEFFER, M. TACKER, and P. SCHUSTER, 1994 Fast Folding and Comparison of RNA Secondary Structures. *Monatshefte fur Chemie* **125**: 167–188.
- HUYNEN, M., P. STADLER, and W. FONTANA, 1996 Smoothness within ruggedness: the role of neutrality in adaptation. *Proceedings of the National Academy of Sciences* **93**(1): 397–401.
- IMHOF, M. and C. SCHLOTTERER, 2001 Fitness effects of advantageous mutations in evolving *Escherichia coli* populations. *PNAS* **98**(3): 1113–1117.
- LEADBETTER, M., G. LINDGREN, and H. ROOTZEN, 1983 *Extremes and Related Properties of Random Sequences and Processes*. Springer-Verlag.
- MCCASKILL, J., 1990 The equilibrium partition function and base pair binding probabilities for RNA secondary structure. *Biopolymers* **29**(6-7): 1105–1109.
- MEYERS, L. A., J. F. LEE, M. COWPERTHWAIT, and A. D. ELLINGTON, 2004 The Robustness of Naturally and Artificially Selected Nucleic Acid Secondary Structures. *Journal of Molecular Evolution* **58**(6): 618–625.
- NUSSINOV, R. and A. JACOBSON, 1980 Fast algorithm for predicting the secondary structure of single-stranded RNA. *Proceedings of the National Academy of Sciences* **77**(11): 6309–6313.

- ORR, H. A., 2002 The Population Genetics of Adaptation: The Adaptation of DNA Sequences. *Evolution* *56*(7): 1317–1330.
- ORR, H. A., 2003 The Distribution of Fitness Effects Among Beneficial Mutations. *Genetics* *163*(4): 1519–1526.
- PARSCH, J., J. M. BRAVERMAN, and W. STEPHAN, 2000 Comparative Sequence Analysis and Patterns of Covariation in RNA Secondary Structures. *Genetics* *154*(2): 909–921.
- ROZEN, D. E., J. A. G. M. DE VISSER, and P. J. GERRISH, 2002 Fitness Effects of Fixed Beneficial Mutations in Microbial Populations. *Current Biology* *12*(12): 1040–1045.
- SANJUAN, R., A. MOYA, and S. F. ELENA, 2004 The distribution of fitness effects caused by single-nucleotide substitutions in an RNA virus. *PNAS* *101*(22): 8396–8401.
- SCHUSTER, P., W. FONTANA, P. F. STADLER, and I. L. HOFACKER, 1994 From sequences to shapes and back: a case study in RNA secondary structures. *Proceedings of the Royal Society of London B* *255*(1344): 279–284.
- VAN NIMWEGEN, E., J. CRUTCHFIELD, and M. HUYNEN, 1999 Neutral evolution of mutational robustness. *Proceedings of the National Academy of Sciences* *17*(76): 9716–9720.
- WATERMAN, M., 1978 Secondary structure of single-stranded nucleic acids. In *Studies on foundations and combinatorics, Advances in mathematics supplementary studies*, Volume 1, pp. 167–212. Academic Press N.Y.
- WEISSMAN, I., 1978 Estimation of Parameters and Larger Quantiles Based on the k Largest Observations. *Journal of the American Statistical Society* *73*(364): 812–815.
- WILKE, C. O. and C. ADAMI, 2001 Interaction between directional epistasis and average mutational effects. *Proceedings of the Royal Society of London: Series B* *268*(1475): 1469–1474.
- WUCHTY, S., W. F. I. HOFACKER, and P. SCHUSTER, 1999 Complete suboptimal folding of RNA and the stability of secondary structures. *Biopolymers* *49*(2): 145–165.
- ZUKER, M., 1989 On Finding All Suboptimal Foldings of an RNA Molecule. *Science* *244*(4900): 48–52.
- ZUKER, M. and P. STIEGLER, 1981 Optimal computer folding of large RNA sequences using thermodynamics and auxiliary information. *Nucl. Acids. Res.* *9*(1): 133–148.

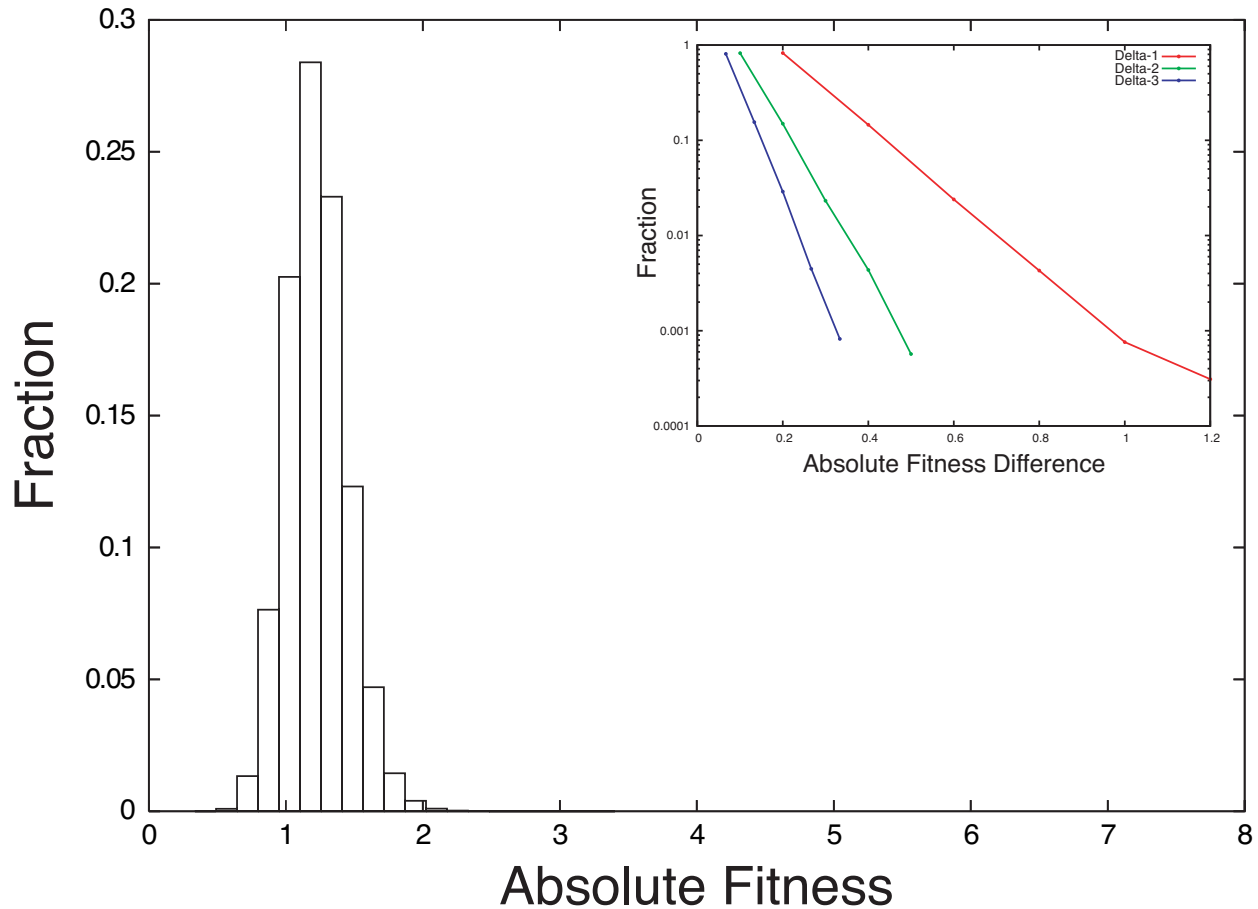


Figure 1: The distribution of absolute fitness of 3,636,520 random sequences. The data was divided into 10 equal-width bins and plotted so that the center of the column on the x-axis is at the upper bin bound. The y-axis is the fraction of sequences falling into a particular bin. Inset: the distribution of Δ_1 , Δ_2 and Δ_3 (see text) for 15,880 sets of 229 absolute fitness values. The x-axis is the fitness effect and the y-axis is the fraction of fitnesses falling a particular bin on a log scale. The bin width is 0.2 for Δ_1 , 0.1 for Δ_2 , 0.67 for Δ_3 .




Figure 2: The cumulative distribution of beneficial fitness effects of wildtype alleles from random walks. Data are from 5721 adaptive walks starting from random sequences - one wildtype genotype per rank per walk. The x-axis is the size of the beneficial fitness effect and the y-axis is the fraction of mutants with fitness greater than the x-axis value on a log scale. The red curve is $R = 2$ ($n = 5004$), the green curve is $R = 3$ ($n = 9908$) and the blue curve is $R = 4$ ($n = 14871$). Inset: exponential behavior when truncated at $S = 0.20$. Color of curves matches the main figure.

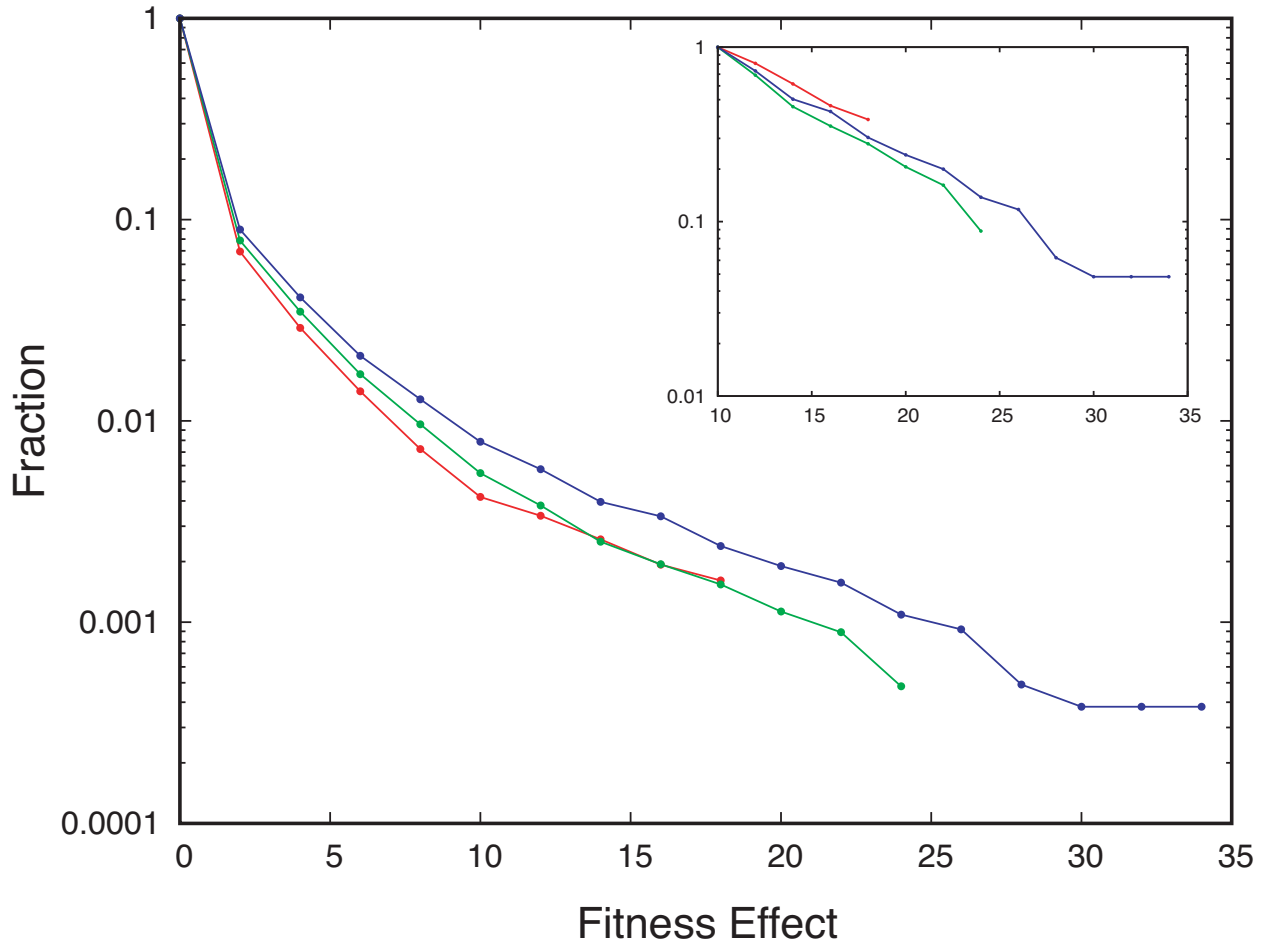


Figure 3: The cumulative distribution of all one-step beneficial fitness effects of wild-type alleles from the high-fitness walks. Data are from 6959 adaptive walks starting near fitness optima - one wildtype genotype per rank per walk. The x-axis is the size of the beneficial fitness effects and the y-axis is the fraction of mutants with fitness greater than the x-axis value on a log scale. The red curve is $R = 2$ ($n = 6204$), the green curve is $R = 3$ ($n = 12374$) and the blue curve is $R = 4$ ($n = 18432$). Inset: exponential behavior when truncated at $S = 10.0$. Color of curves matches the main figure.




Figure 4: The effect of truncation on the estimated mean beneficial effect. The distributions become approximately exponential when the curve shown here asymptotes. The upper plot is the estimate of the mean using absolute fitness differences. The lower plot is the estimate based on s values.




Figure 5: Mean s for all beneficial mutations in the neighborhood of $R \leq 4$ wildtype sequences across the length of an adaptive walk. Data are from 5721 random and 6959 high-fitness adaptive walks. The x-axis is the number of substitutions and the y-axis is the mean s of the one-step beneficial mutations from all low-rank wildtype alleles at that step.