

Plasticity, Evolvability, and Modularity in RNA

LAUREN W. ANCEL¹ AND WALTER FONTANA^{2,3*}

¹*Department of Biological Sciences, Stanford University, Stanford, California 94305*

²*Santa Fe Institute, Santa Fe, New Mexico 87501*

³*Institute for Advanced Study, Program in Theoretical Biology, Princeton, New Jersey 08540*

ABSTRACT RNA folding from sequences into secondary structures is a simple yet powerful, biophysically grounded model of a genotype–phenotype map in which concepts like plasticity, evolvability, epistasis, and modularity can not only be precisely defined and statistically measured but also reveal simultaneous and profoundly non-independent effects of natural selection. Molecular plasticity is viewed here as the capacity of an RNA sequence to assume a variety of energetically favorable shapes by equilibrating among them at constant temperature. Through simulations based on experimental designs, we study the dynamics of a population of RNA molecules that evolve toward a predefined target shape in a constant environment. Each shape in the plastic repertoire of a sequence contributes to the overall fitness of the sequence in proportion to the time the sequence spends in that shape. Plasticity is costly, since the more shapes a sequence can assume, the less time it spends in any one of them. Unsurprisingly, selection leads to a reduction of plasticity (environmental canalization). The most striking observation, however, is the simultaneous slow-down and eventual halting of the evolutionary process. The reduction of plasticity entails genetic canalization, that is, a dramatic loss of variability (and hence a loss of evolvability) to the point of lock-in. The causal bridge between environmental canalization and genetic canalization is provided by a correlation between the set of shapes in the plastic repertoire of a sequence and the set of dominant (minimum free energy) shapes in its genetic neighborhood. This statistical property of the RNA genotype–phenotype map, which we call plastogenetic congruence, traps populations in regions where most genetic variation is phenotypically neutral. We call this phenomenon neutral confinement. Analytical models of neutral confinement, made tractable by the assumption of perfect plastogenetic congruence, formally connect mutation rate, the topography of phenotype space, and evolvability. These models identify three mutational regimes: that corresponding to neutral confinement, an exploration threshold corresponding to a breakdown of neutral confinement with the simultaneous persistence of the dominant phenotype, and a classic error threshold corresponding to the loss of the dominant phenotype. In a final step, we analyze the structural properties of canalized phenotypes. The reduction of plasticity leads to extreme modularity, which we analyze from several perspectives: thermophysical (melting—the RNA version of a norm of reaction), kinetic (folding pathways—the RNA version of development), and genetic (transposability—the insensitivity to genetic context). The model thereby suggests a possible evolutionary origin of modularity as a side effect of environmental canalization. *J. Exp. Zool. (Mol. Dev. Evol.)* 288:242–283, 2000. © 2000 Wiley-Liss, Inc.

Biological evolution is the transformation of heritable phenotypes through time. Evolution is fueled by the introduction of novel phenotypes and steered by population-level interactions including natural selection and genetic drift. The predominant route to heritable phenotypic change originates with genetic mutation. The processes that translate genetic variation into phenotypic variation give rise to an association between genotype and phenotype which we represent as a map that is sensitive to environmental conditions. Concepts such as canalization, epistasis, and modularity

underlie our understanding of phenotypic variability [for a sweeping perspective see Wagner and Altenberg ('96) and Schlichting and Pigliucci ('98)]. Yet, a microfoundation of these concepts and of their interconnections in terms of the relation between genotype and phenotype is largely missing. Our goal here is to initiate such a foundation in the specific context of a conceptually, computa-

*Correspondence to: W. Fontana, Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501. E-mail: walter@santafe.edu
Received 4 November 1999; accepted 20 May 2000

tionally, and empirically simple yet powerful genotype–phenotype model based on the folding of RNA sequences (genotypes) into shapes (phenotypes).

RNA folding seems, at first, unlikely to be able to address canalization, epistasis, and modularity. Although we shall mostly use language appropriate for RNA, an RNA sequence could be viewed as a metaphor for a genome, and a position along the sequence as a metaphor for a locus with four possible alleles (nucleotides). A phenotype (RNA shape) then is a simple pattern of gene–gene interactions (base pairs, see section 1.2). As this paper weaves together several seemingly diverse concepts, we begin with an overview.

Our study provides a molecular illustration for the Simpson–Baldwin effect (Baldwin, 1896; Simpson, '53; Ancel, '99a) using RNA as an example. Central to the initial “discovery” stage of the Simpson–Baldwin effect is phenotypic plasticity, that is, the genetically influenced capacity of an individual to develop into one among a range of phenotypes. In an evolutionary context, a fixed environment will convey a selective advantage to those individuals that can access an improved phenotype within their plastic repertoire over those who cannot. The “assimilation” stage of the Simpson–Baldwin effect arises from the fitness costs of plasticity. Among the individuals selected in the first stage, those that can still access the improved phenotype while reducing their range of phenotypic plasticity will have a selective advantage. The Simpson–Baldwin effect describes the genetic determination of a phenotype that previously seemed to be acquired anew in each generation. A frequently considered mechanism of plasticity is learning (Hinton and Nowland, '87; Maynard-Smith, '87).

Although molecules do not learn, biopolymers like RNA are plastic in the sense that a given sequence can realize a repertoire of alternative structures, rather than being frozen in its minimum free energy configuration (section 1). An RNA sequence samples a variety of energetically low lying structures by wiggling among them under thermal fluctuation. The overall time a sequence spends in a shape reflects the thermodynamic stability of that shape. Consider now a population of replicating and mutating RNA sequences that are subject to selection for structural proximity to a constant target shape (section 2). Assume further that each shape attainable by a sequence contributes to that sequence's overall fitness in proportion to the time the sequence spends in it.

Sequences with an advantageous but energeti-

cally suboptimal shape will be selected over those that lack that shape. Plasticity entails a fitness cost because the more alternative shapes an RNA molecule can fold into, the less time it will spend in each shape, including advantageous ones. Sequences with an advantageous but energetically suboptimal shape σ will therefore be replaced by mutants with σ at lower free energy until σ becomes the minimum free energy structure. Subsequently, natural selection will fine tune the thermodynamic stability of σ by favoring variants with few alternative structures in the energetic vicinity of σ . We show that in our RNA model such genetic assimilation occurs extremely rapidly and covers several orders of magnitude in thermodynamic stability (section 2).

Some models link plasticity to a speed-up in evolution (Hinton and Nowlan, '87). This is not the case in our RNA model. In fact, the reduction of plasticity in a constant environment leads to a slow down of evolution to the point of a phenotypic dead-end. In section 2 we describe this dynamic, and in section 3 we offer an explanation for this behavior in terms of features that are intrinsic to the RNA genotype–phenotype map. Several threads come together as we argue that genetic assimilation (the reduction of plasticity) requires a genotype–phenotype map in which plasticity mirrors variability. In other words, the shapes appearing in a sequence's repertoire of energetically favorable structures correlate significantly to the minimum free energy structures of the one-error mutants of that sequence. This turns out to be a general property of the RNA genotype–phenotype map (section 3). We call this phenomenon “plastogenetic congruence.”

Phenotypic variability describes the extent of phenotypic variation accessible to a genotype through mutation. The evolvability of an individual is the likelihood of reaching a phenotype with improved fitness through mutation (Altenberg, '95). As such, evolvability is linked to variability via the fitness function. Plasticity, on the other hand, captures the phenotypic variation at a fixed genotype, typically induced by environmental heterogeneity. In this sense, the impacts of the environment on plasticity are analogous to those of genetic mutation on phenotypic variability. Plastogenetic congruence means that plasticity and variability mirror each other: low plasticity (that is, strong genetic determination) implies low variability (that is, strong mutational buffering), and vice versa.

Plastogenetic congruence implies that an evolutionary reduction of plasticity has a flip side: a decline in variability. This link results in a self-

defeating process in which the loss of plasticity through natural selection leads to the loss of phenotypic variability, to the extent of evolutionary lock-in.

Plastogenetic congruence also sheds light on Waddington's theories for the evolution of organismal development (Waddington, '57). He introduced two modes of evolution: environmental canalization is the honing of developmental pathways to reduce environmental noise, and genetic canalization is the integration of genetic factors to reduce the deleterious effects of mutation. Under plastogenetic congruence then, genetic canalization will ensue as a byproduct of selection for environmental canalization. This yields a mechanistic explanation of a hypothesis put forward by Wagner et al. ('97).

Pairwise epistasis is the influence that an alteration of gene i has on the phenotypic consequences of a subsequent alteration of gene j (Wagner et al., '98). This definition of epistasis captures the genetic control of variability. In RNA, low plasticity coincides with low variability, maintained by the fixation of epistatic interactions that buffer the phenotype against mutations. Remarkably, epistatic interactions in RNA can eliminate variability almost completely. This phenomenon, which we call "neutral confinement," contributes to the evolutionary lock-in mentioned earlier.

Analytical models built on a stylized version of plastogenetic congruence predict that, for certain parameter regimes, the mutation rate needed to escape this exploration catastrophe is so high as to result in the loss of the dominant phenotype (error catastrophe). We discuss this issue in sections 4 and 5.

The RNA folding genotype–phenotype map enables not only an exploration of evolutionary dynamics but also a characterization of the morphological endpoints of evolution. In section 5 we compare three classes of sequences that share the same dominant (that is, minimum free energy) structure. One set is derived from a random sample of sequences with the given dominant structure, another set has evolved on a neutral network (see section 1.3), and the third set results from genetic assimilation under the plastic genotype–phenotype map. The characteristic which best distinguishes among the three classes is modularity. We study modularity of shape characters from a variety of perspectives, all contributing to a definition and quantification of modular traits as "transposable characters" that maintain their structural

integrity in different sequence and environmental contexts. Although the three classes share the same minimum free energy structure, that structure is not even remotely modular on the random sequences while it is extremely modular on sequences that have experienced genetic assimilation.

Genetic assimilation leaves us with sequences that possess modular shapes and are at the same time evolutionarily locked in. This seems to contradict the hypothesized evolutionary advantage of modularity: modularity partitions quantitative traits into independently and easily evolvable units (Wagner and Altenberg, '96). While modularity may indeed facilitate the quantitative polishing of a trait, it leads to an evolutionary lock-in with respect to significant structural modifications of that trait. Resistance to structural change is the hallmark of a module. Once modules are available, the generation of further evolutionary novelty (or plasticity) then may shift from locked-in modules to the combinatorial arrangement of modules into new units.

1. RNA FOLDING AS A GENOTYPE-PHENOTYPE MAP

1.1. Why RNA?

RNA combines genotype and phenotype into a single molecule. This makes RNA folding in many respects a limited, but also a simple, model of a genotype–phenotype map. As a model system, RNA has the advantages of both computational tractability (Waterman, '95) and suitability as a substrate for test tube evolution experiments (Joyce, '89; Landweber, '99). The RNA sequence–structure relation also occupies a rare intermediate level of abstraction bridging the empirical and the formal. RNA folding algorithms are sufficiently realistic for computational discoveries to suggest worthwhile empirical investigations. At the same time, RNA folding is sufficiently abstract to provide insight and to suggest axioms for the construction of simplified models that are analytically tractable.

1.2. Secondary structure

RNA molecules are heteropolymers of (predominantly) four units called ribonucleotides. Ribonucleotides have a ribose phosphate in common but differ in the base attached to the sugar. The essence of an RNA sequence is therefore captured by a string over a four letter alphabet, each letter representing a particular base: A for adenine, U for uracil, C for cytosine, and G for guanine. Hydrogen bonds give rise to stereoselective recogni-

tion between certain base pairs, specifically A·U and G·C. This base pairing enables an RNA sequence to be copied into a negative and back again into a positive. The pairing is not always exact. Error rates depend on the molecular machinery that assists in pairing and ligating the bases. For example, the per nucleotide error rate is 7×10^{-5} to 2.7×10^{-4} for Influenza and 3×10^{-4} for Coliphage Q β (Eigen et al., 1989). In this way, base pairing enables heredity in RNA viruses. We therefore treat an RNA sequence as a genotype. The same base pairing mechanism, however, also enables segments of a sequence to pair with other segments within the same sequence, causing it to fold back on itself into a three-dimensional structure. (For the formation of an intramolecular structure G·U is also an admissible pair.) This structure conveys chemical behavior to the sequence and constitutes its phenotype.

The rapid replication time and simplicity of the phenotype make RNA a tractable laboratory model. RNA molecules can be evolved in the test tube using a variety of techniques for amplification, variation, and selection. Evolutionary optimization of RNA properties in the test tube occurs readily and effectively (Mills et al., '67; Spiegelman, '71; Ellington and Szostak, '90; Tuerk and Gold, '90; Beaudry and Joyce, '92; Bartel and Szostak, '93; Ellington, '94; Eklund et al., '95; Landweber and Pokrovskaya, '99). For a recent review see Landweber ('99).

Molecular structure in RNA can be characterized at many levels of resolution. One empirically well-established notion is the secondary structure, which is the topology of binary contacts as they result from base pairing (Figure 1). The secondary structure is a useful abstraction, since the pattern of base pairs provides both a geometric and thermodynamic scaffold for the tertiary structure of the molecule. This puts the secondary structure in correspondence with some functional properties of the tertiary structure.

A secondary structure on a sequence of length n can be represented as a graph of base pair contacts (Fig. 1). The nodes of the graph stand for bases at positions $i = 1, \dots, n$ along the sequence. The set of edges includes the unspecific covalent backbone connecting node i with node $i + 1$, for $i = 1, \dots, n - 1$, and those indicating pairings between non-adjacent positions. The set of such non-adjacent pairings P has to satisfy two conditions: (i) every edge in P connects a node to at most one other node, and (ii) if both $i < j$ and $k < l$ are in P , then $i < k < j$ implies that $i < l < j$. Failure to




Fig. 1. RNA secondary structure graph. A secondary structure is a graph consisting of structural elements called cycles or loops: a hairpin loop occurs when one base pair encloses a number of unpaired positions, a stack consists of two base pairs with no unpaired positions, while an interior loop has two base pairs enclosing unpaired positions. An internal loop is called a bulge if either side has no unpaired positions. Finally, multiloops are loops delimited by more than two base pairs. A position that does not belong to any loop type is called external, such as free ends or joints. Components are shape features delimited by external bases.

meet condition (ii) results in pseudoknots which are interactions that belong to the next—the tertiary—level of structure. Both conditions distinguish RNA structure from protein structure, in particular condition (i) which builds RNA secondary structure exclusively from binary contacts. We use a picture of the graph (Fig. 1) as the visually most immediate representation of a secondary structure. We sometimes use a more convenient line-oriented representation of nested parentheses, such as “(((((((...))).((...))).))))),” in which a dot stands for an unpaired position, and a pair of matching parentheses indicates positions that pair with one another.

The elements of a secondary structure graph are certain types of cycles or loops, see Fig. 1. Two contiguous base pairs constitute the smallest loop. We make the reasonable assumption that the overall energy of a secondary structure is the sum of its loop energies. These have been measured and tabulated (Freier et al., '86; Turner et al., '88; Jaeger et al., '89; He et al., '91) as a function of loop size and delimiting base pairs. A stack of base

pairs—a double-stranded region of several contiguous base pairs—is the major stabilizing element. The formation of an energetically favorable stacking region, however, implies the formation of an energetically unfavorable loop constraining unpaired bases (for example, a hairpin loop). This “frustrated” energetics leads to a vast combinatorics of stack and loop arrangements constituting the structural repertoire of an individual RNA sequence. RNA is an excellent model system for sequence–structure relations in biopolymers precisely because of our ability to rapidly compute the set of all structures realized by a sequence. In particular, we use dynamic programming to compute the minimum free energy secondary structure and statistical mechanics quantities, such as the partition function (Nussinov et al., '78; Waterman and Smith, '78; Nussinov and Jacobson, '80; Zuker and Stiegler, '81; Zuker and Sankoff, '84; McCaskill, '90). The work in this paper makes use of the Vienna RNA folding package (Hofacker et al., '94-'98), a state-of-the-art library of utilities and RNA folding programs routinely applied in the laboratory.

1.3 Robust properties of RNA folding

RNA folding algorithms vary considerably in the accuracy of secondary structure predictions for individual instances (Huynen et al., '97). Our main concern, however, is with statistical properties of the sequence to structure map as a whole, rather than with specific cases. We review two statistical descriptions of the RNA folding map: typical shapes and neutral networks (Fontana et al., '93a,b; Schuster et al., '94; Grüner et al., '96a,b).

Sequence space and shape space are both very high dimensional, and the sequence space is substantially larger than the shape space. Analytical tools developed in Stein and Waterman ('78) yield an upper bound of only $S_n = 1.48 \times n^{-3/2} 1.85^n$ shapes vis-à-vis 4^n sequences, where n is the sequence length (Hofacker et al., '99). The mapping from sequences to minimum free energy shapes is significantly many-to-one.

Typical shapes

Relatively few shapes are realized with very high frequency, contrasting many rarely occurring shapes. More precisely, as sequence length goes to infinity, the fraction of such “typical shapes” tends to zero (their number grows nevertheless exponentially), while the fraction of sequences folding into them tends to one. Consider a numerical example: In the space of GC-only sequences of

length $n = 30$, 1.07×10^9 sequences fold into 218,820 shapes; 22,718 shapes (10.4%) are typical in the sense of being formed more frequently than the average number of sequences per shape, and 93.4% of all sequences fold into these 10.4% shapes (Grüner et al., '96a,b; Schuster, '97).

Neutrality and neutral networks

Many sequences have the same (typical) shape α as their minimum free energy structure. We call such sequences “neutral” with respect to α . A structure α thereby identifies an equivalence class of sequences. A one-error mutant of a sequence that shares the same minimum free energy structure as that sequence is called a “neutral neighbor.” By “neutrality” of a sequence we mean the fraction of its $3n$ one-error mutants that are neutral neighbors. This notion of neutrality pertains to the minimum free energy structures of RNA sequences, and should not be confused with fitness-based neutrality.

Any given sequence has a significant fraction of neutral neighbors, and the same holds for these neighbors. In this way, jumping from neighbor to neighbor, we can map an extensive mutationally connected network of sequences that fold into the same minimum free energy structure (Schuster et al., '94; Reidys et al., '97). We termed such networks “neutral networks” (Schuster et al., '94). The ability to change a sequence while preserving the phenotype is critical for evolutionary dynamics. The role of neutrality has been typically viewed as a conservative one, buffering the phenotypic effects of mutations. Neutrality, however, also enables phenotypic change, because it permits phenotypically silent mutations to influence the phenotypic consequences of other mutations.

The boundary of a neutral network is the set of sequences that differ by one mutation from a sequence in the network, but are not themselves contained in the network. Transitions between structures are transitions between adjacent neutral networks. This suggests a measure of nearness between RNA structures based on the fraction of common boundary shared between their corresponding networks in sequence space (Fontana and Schuster, '98a,b), rather than on shape similarity. RNA shape space is then organized by an accessibility relation based on the adjacency of neutral networks in genotype space. Such a topology enables a formal definition of continuous and discontinuous phenotypic change independently of fitness criteria (Fontana and Schuster, '98a). Figure 2 describes the RNA shape

transformations that are discontinuous in this accessibility sense.

1.4. Plasticity in RNA

Plasticity is the genetic determination of a range of phenotypic possibilities, where the realized phenotype depends on the environmental context of the organism. If the environment is relatively consistent or the trait is fixed in development, then plasticity may only be evident over several generations. In other cases, a plastic individual may

change phenotypes during its lifetime through learning or by reacting to environmental stimuli. For a single RNA molecule, unspecific contact with a heat bath triggers transitions between molecular configurations, provided the energy barriers between configurations are sufficiently low. The range of configurations realized by an RNA molecule at a constant temperature depends on the energies of the configurations. The kinetic process of folding into the minimum free energy configuration is the RNA analogue of development. Ran-




Fig. 2. Continuous and discontinuous RNA shape transformations. The figure illustrates transformations between RNA secondary structure motifs. Solid (dashed) arrows indicate continuous (discontinuous) transformations in the topology of Fontana and Schuster ('98). Three groups of transformations are shown. **Top:** The loss and formation of a base pair adjacent to a stack are both continuous. **Middle:** The opening of a constrained stack (e.g., closing a multiloop) is continuous while its creation is discontinuous. This reflects

the fact that the formation of a long stack upon mutation of a single position is a highly improbable event, whereas the unzipping of a random stack is likely to occur as soon as a mutation blocks one of its base pairs. **Bottom:** Generalized shifts are discontinuous transformations in which one or both strands of a stack slide past one another, ending up with or without an overlap. Accordingly, generalized shifts are divided into the four classes shown.

dom transitions between kinetically accessible structures that are energetically close to the minimum free energy configuration are an unavoidable consequence of the stochasticity of molecular motion, akin to “developmental noise.”

A microscopic kinetic folding model for RNA has been developed recently (Flamm et al., '99), but at present it is too involved for the goals of this paper. Instead, we exploit an extension of the standard thermodynamic minimum free energy folding algorithm which permits the computation of all secondary structures within some energy range above the minimum free energy (Wuchty et al., '99). This suboptimal folding algorithm yields fast access to the low-energy portion of the secondary structure space of a given sequence. We neglect energy barriers and assume that a sequence equilibrates among all structures whose free energy is within $5kT$ from the minimum free energy configuration. The $5kT$ choice amounts to approximately 3 kcal at 37°C and corresponds to the loss of two G·C/C·G stacking interactions. Under thermodynamic equilibration, the Boltzmann probability of a structure σ , $\exp(-\Delta G_\sigma/kT)/Z$, corresponds to the overall fraction of time that the molecule spends in σ , where ΔG_σ is the free energy of structure σ , k is the Boltzmann constant, T the absolute temperature, and $Z = \sum_\tau \exp(-\Delta G_\tau/kT)$, the partition function. The latter is computed by an algorithm described in McCaskill ('90).

This defines a genotype–phenotype map that takes a sequence to a set of structures and their occupation times. We shall refer to this map as the “plastic map,” to distinguish it from the “simple map” where a sequence is associated with its minimum free energy structure only (Fig. 3).

Under natural selection toward a target structure (section 2), the obvious advantage for a plastic sequence that covers a broad spectrum of structures is the increased likelihood that some of them are structurally similar to the target. As pointed out in Scheiner ('93) the cost of plasticity may be in terms of maintaining the cellular machinery required for plasticity, rather than the direct impact on fitness due to the realization of a particular plastic trait. In the stripped down world of RNA there is no such machinery beyond the molecule itself. Yet, the cost of plasticity is evident: the broader the range of structures, the less time the sequence will spend in any one of them. Thus, even if some structures constitute an improvement, this can easily be undermined by a small occupancy time. Incidentally, this is analogous to Schmalhausen’s argument ('49) that one

cost of plasticity is given by “erroneous” phenotypic changes (Ancel, '99a). The erroneous changes in RNA are residencies in detrimental structures.

Biologists have drawn a distinction between two kinds of plasticity: phenotypic plasticity proper often refers to beneficial responses to macroenvironmental variation, while environmental variance refers to flexible responses to microenvironmental parameters (Waddington, '57; Gavrillets and Hastings, '94). Our work in RNA considers the latter. A given sequence assumes a range of structures in response to its microenvironment—energy fluctuations at a constant temperature.

1.5. Epistasis in RNA

Epistasis is the extent to which the phenotypic consequences of a mutation at position i depend on the genetic background provided by the remaining sequence. One common estimate for the epistasis of a genome measures the deviation of pairwise gene interactions from additivity (Wagner et al., '98).

Some epistatic interactions increase the neutrality of a sequence by making mutations at other sites inconsequential. Consider the sequence and its minimum free energy structure at the center of Fig. 4. We call a sequence position “neutral” if at least one out of three possible mutations at that locus leave the structure invariant. The neutral positions, like the one labeled x , are marked with grey bullets. The neutral mutation from C to G at position x yields the (same) structure shown at the top left of Fig. 4. The “+” symbols indicate positions that have become neutral as a result of this mutation, while the “−” symbol marks a position that has lost its neutrality. This illustrates that even if a point mutation does not affect the structure, it can alter the extent of neutrality across the sequence. Epistasis as the genetic control of neutrality plays an important role in the evolutionary dynamics we describe in section 2.

1.6. Neutrality: a note in terminology

In section 1.3 two sequences are called “neutral” if they share the same minimum free energy structure. The term neutral, however, may be used to indicate equal fitness. Under the simple map, the phenotype of an RNA sequence is just its minimum free energy structure (Fig. 3). In this case, neutrality with respect to minimum free energy structures implies neutrality with respect to




Fig. 3. Genotype–phenotype maps. The lower part illustrates the simple map that takes a sequence to its minimum free energy structure as the only phenotype. The upper part schematizes the plastic map that takes a sequence to the thermodynamic spectrum of shapes within $5kT$ (typically $T = 310.15$ K). If shape σ has free energy ΔG_σ , the sequence is

assumed to spend a fraction of time in σ that is given by its Boltzmann factor, $\exp(-\Delta G_\sigma/kT)/Z$ (values indicated on the left). The minimum free energy structure is the dominant phenotype, in the sense that the sequence spends the largest fraction of time in it.

both phenotype and fitness. Under the plastic map, however, the phenotype of an RNA sequence is a repertoire of energetically favorable structures. In this case, neutrality with respect to minimum free energy structure is not equivalent to neutrality with respect to phenotype or fitness. Two sequences that share a minimum free energy structure, will differ in their remaining repertoire. Nevertheless, we maintain the use of minimum

free energy neutrality for two reasons. First, it is relevant to the phenotype of sequences under the plastic map, as they spend a majority of time in the minimum free energy structure. Second, evolutionary dynamics under the simple map—which constitute our baseline for comparison—have historically been characterized by this notion (Huyen et al., '96; Fontana and Schuster, '98a; Reidys et al., '98).




Fig. 4. Epistasis.

2. SIMULATED EVOLUTION IN PLASTIC AND NON-PLASTIC RNA POPULATIONS

2.1. Model setup

In this section we describe simulations of evolving RNA populations. Under the simple genotype-phenotype map (section 1.4) a sequence S has a single minimum free energy structure σ_0 as its phenotype. To determine the fitness of S , we compare σ_0 to a prespecified target structure τ .

This fitness function is motivated by an experimental protocol in which RNA sequences are evolved, for example, to optimally bind a ligand. RNA sequences are artificially selected by running the RNA sample through a column that has the ligand tethered to the filling material. The desirable portion of the sample remains bound in the column and is eluted by a solvent with suitable ionic strength (Ellington and Szostak, '90; Tuerk and Gold, '90). The selected portion is then

subject to a further cycle of evolution by amplification through replication at a controlled error rate. The evolutionary end product is typically unpredictable in the laboratory situation. Its possible shape(s) are, however, implicitly specified by the choice of the ligand, like a simple lock specifies its key. We are not seeking RNA shapes with particular chemical properties, since little is known about the link between an RNA structure and its binding properties or catalytic behaviors. Instead, we specify the optimal shape directly at the outset, shortcutting the role of ligands in the laboratory. We then study the evolutionary dynamics, the evolutionary histories, and the thermodynamic properties of evolved sequences.

We define the selective value $f(\sigma)$ of shape σ as a hyperbolic function of the Hamming distance $d(\sigma, \tau)$ between the parenthesized representations (section 1.2) of σ and the target τ :

$$f(\sigma) = \frac{1}{0.01 + d(\sigma, \tau)/n}, \quad (1)$$

where n is the length of the sequence. The results reported here are robust to changes in the functional form of the selective value. In particular, they also hold under linear and exponential forms.

The fitness r_j of a sequence S_j is its replication rate constant, and for the non-plastic case we simply take it to be:

$$r_j = f(\sigma_0^j). \quad (2)$$

Our population evolves in a model chemical flow reactor whose outflow is regulated to maintain a nearly constant total sequence population size (Eigen, '71). The entire system is simulated in terms of continuous time stochastic chemical reaction kinetics. [For a description see Fontana and Schuster ('87) and Huynen et al. ('96); for a general simulation technique see Gillespie ('76, '77).] The number of replication events per time unit in the flow reactor depends on the replication rate constants of the individual sequences comprising a population, and changes over time as the population evolves to larger rate constants. When comparing different runs, we therefore plot statistics along replication events rather than the time shown by an external clock.

Point mutations are the sole source of genetic variation in our simulations. Unless otherwise stated, the replication accuracy per nucleotide position is 0.999, the sequence length is 76 (which corresponds to a short tRNA sequence), and the reactor capacity is 1,000 sequences. Simulations begin with a homogenous population of a randomly generated sequence species.

The plastic fitness function is a simple extension of the one above. For a sequence S_j , we consider all of its suboptimal structures σ_i^j whose free energy $E(\sigma_i^j)$ is within an interval of size Δ above the minimum free energy $E(\sigma_0^j)$. The σ_i^j are indexed with increasing energy, where index 0 refers to the ground state, and the set of all suboptimal structures accessible within $[E(\sigma_0^j), E(\sigma_0^j) + \Delta]$ is denoted with $\sigma^j(\Delta)$. Unless otherwise specified, $\Delta = 5kT$, where T is the absolute temperature, which is fixed at 310.15 K (37°C), and k is the Boltzmann constant, $k = 1.98717 \times 10^{-3}$ kcal/K. Thus, $\Delta = 3.08$ kcal. The macroscopic environment of our model consists in the target shape τ and the temperature T . Throughout this paper we only consider fixed τ and T .

The selective value of a suboptimal structure

σ_i^j is given by $f(\sigma_i^j)$ (eq. 1), exactly as in the non-plastic case. For a plastic sequence, however, σ_i^j contributes this value to the overall fitness of the sequence in proportion to the time the sequence spends in it. In the laboratory, a sequence flowing through the selection column will switch among its alternative structures, such that each structure contributes to the overall binding probability of the sequence in proportion to its Boltzmann factor. Hence,

$$r_j = \sum_{\sigma_i^j \in \sigma^j(\Delta)} \overbrace{f(\sigma_i^j)}^{\text{selective value of } \sigma_i^j} \frac{\overbrace{\exp(-E(\sigma_i^j)/kT)}^{\text{probability of } \sigma_i^j}}{Z}, \quad (3)$$

where Z is the partition function. Here we assume that a sequence achieves thermodynamic equilibration among its alternative structures. High-energy barriers between the structures, however, may render this assumption kinetically unrealistic on time scales shorter than the experiment. Our understanding of the kinetic process of RNA folding is too immature at present to assess this claim.

2.2. Discussion of sample runs

RNA shape space topology is invariant to plasticity

Figures 5 and 6 juxtapose the progress of a simple and a plastic population that evolve toward a tRNA target shape. We will denote population averages with $\langle \cdot \rangle$. The black curves trace $\langle d(\sigma_\delta^j, \tau) \rangle$, the average distance between the target and the minimum free energy shapes realized by the sequences in the population.

The plastic and simple populations exhibit qualitatively similar behavior: initial relaxation followed by punctuated dynamics. Steps toward the target shape correspond to difficult transformations from one dominant shape to another. Recall from section 1.3 that the difficulty of discovering minimum free energy shapes through mutation is independent of the fitness assigned to shapes. An evolutionary transformation of one shape into another corresponds to a transition between their neutral networks in sequence space. In the simulations, flat periods correspond to populations that are mutationally isolated from higher fitness phenotypes and are genetically diffusing along a shape-neutral network or along several shape-neutral networks that are fitness-neutral with respect to each other. Steps, or difficult shape transformations, are transitions between neutral networks with a small common border. Such shape trans-




Fig. 5. Evolution under plastic and simple genotype–phenotype maps (example 1). These graphs depict the evolution of the population average of the distance between the minimum free energy structures in the population and the target (a), as well as the average weighted shape distance to target (b). (See text for definitions.) The inset magnifies the trajec-

tories through the first few million replication events. In the plastic case, fitness improvements correspond to transitions manifest in the average weighted distance to the target. The target is a tRNA cloverleaf structure, the population size is logistically constrained to fluctuate around 1,000 sequences, and the replication accuracy is 0.999 per position.

formations involve the simultaneous shift of several base pairings in a single mutational event.

The gray curves in Figs. 5 and 6 monitor $\langle d(\sigma^j, \tau) \rangle$, the population average of the weighted shape distance. For any given sequence, $\langle d(\sigma^j, \tau) \rangle$ is the Boltzmann weighted average dis-

tance between its suboptimal shapes and the target:

$$\overline{d(\sigma^j, \tau)} = \frac{1}{Z} \sum_{\sigma_i^j \in \sigma^j(\Delta)} d(\sigma_i^j, \tau) \cdot \exp(-E(\sigma_i^j)/kT). \quad (4)$$

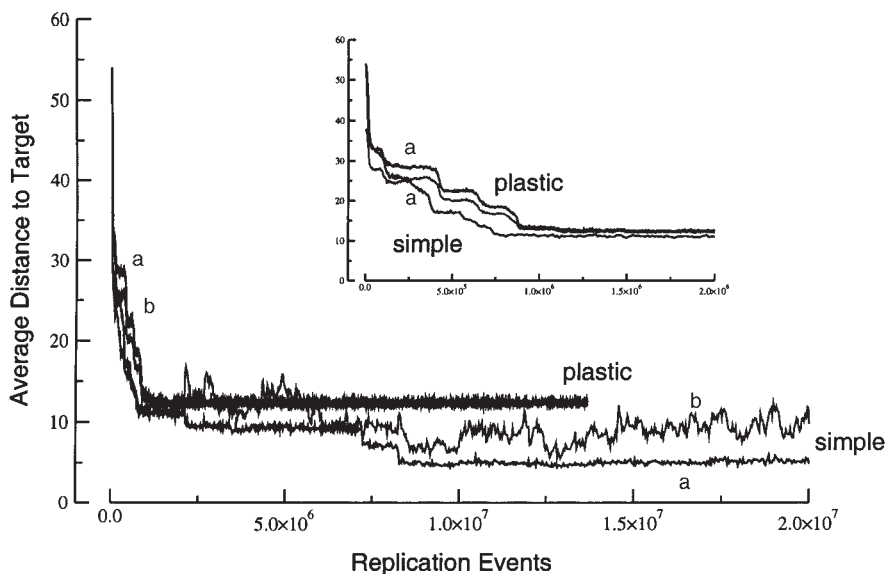


Fig. 6. Evolution under plastic and simple genotype–phenotype maps (example 2). See caption to Fig. 5.

In the plastic case, $\langle d(\sigma_0^j, \tau) \rangle$ dictates the trajectory of $\langle \overline{d(\sigma^j, \tau)} \rangle$. The minimum free energy structure of a sequence S is usually closer to the target than any other suboptimal structure of S . When this is the case, the Boltzmann weighted average distance to target over *all suboptimals*, not just those within $\sigma^j(\Delta)$, would necessarily be larger than the average minimum free energy distance to target, $\langle d(\sigma_0^j, \tau) \rangle$. But $\langle \overline{d(\sigma^j, \tau)} \rangle$, which considers only the shapes in $\sigma^j(\Delta)$, remains less than $\langle d(\sigma_0^j, \tau) \rangle$. This indicates that the suboptimal structures are not much worse than the minimum free energy structure or that their probabilities are relatively small. Below we show that both of these factors play a role. This is in contrast to the statistics for evolution under the simple map where the $\langle \overline{d(\sigma^j, \tau)} \rangle$ fluctuates wildly at values much higher than $\langle d(\sigma_0^j, \tau) \rangle$.

No Baldwin expediting effect

Figures 5 and 6 represent the outcomes of a large set of similar simulations: evolution proceeds much more slowly under the plastic map, and plastic populations never approach the target as closely as do simple populations, which often attain the target. We defer the explanation of this striking contrast until section 3.

Here we take issue with a generalization often drawn from the Hinton and Nowlan class of models (Hinton and Nowlan, '87; Maynard-Smith, '87) that phenotypic plasticity expedites evolution by effectively smoothing the fitness function. More specifically, they claim that plastic individuals scan a wider range of phenotypes, and therefore have a better chance of locating advantageous ones than non-plastic individuals who are blind to other possibilities. Plasticity will accelerate genetic evolution, however, only if the phenotypes realized through plasticity correspond to the phenotypes that can be reached through mutation. In section 3 we show that such a correlation exists in RNA (plastogenetic congruence), and that it ultimately leads populations into an evolutionary dead-end. In this section, we counter the claim that plasticity speeds up the early stages of the evolutionary process. On average, plasticity does not shorten the early periods of stasis preceding phenotypic innovations. To see this, we compare the early evolutionary trajectories under the plastic map and the simple map (see the inset of Fig. 5, which magnifies the first few million replications).

Recent work specifies the restrictive population-genetic conditions for a Baldwin expediting effect (Ancel, '99b). In the present case, the absence of

an expediting effect results from features intrinsic to the RNA genotype–phenotype map. Consider the evolutionary transition from a structure α to a higher fitness structure β . The simple population must wait for a mutation that takes a sequence from minimum free energy structure α to minimum free energy structure β . In the plastic population, this transition may be mediated by a sequence that has minimum free energy structure α and suboptimal structure β . Call this sequence S . The opportunity to encounter β as a suboptimal may expedite the production of a mutant with minimum free energy structure β , *if* there is a correlation between the suboptimal structures of a sequence and the minimum free energy structures produced through mutations on that sequence.

Assume that β is difficult to access from α in the shape space topology sensu Fontana and Schuster ('98a) (section 1.3). This is typically the case at the major phenotypic transitions in the simulations. We claim that the discovery of individual S under the plastic map is approximately as unlikely as the first appearance of an individual with β as its minimum free energy shape in a non-plastic population dominated by α . That is, plasticity does not significantly expedite the search for a higher fitness structure.

First we note that the individual S is always possible in principle. If we relax conditions by asking if there is a sequence S' that can assume both shapes α and β irrespective of their energies, that is, dropping the requirements that β is in the $5kT$ band and α is the minimum free energy structure, then such a sequence S' exists for any choice of α and β . This was proved by Reidys et al. ('97). A sequence is said to be compatible with a structure if it can form that structure, regardless of where it ranks energetically. A given structure α , then, determines a set consisting in all sequences compatible with α . Our individual S' must therefore be an element of the intersection between the compatibility set of α and the compatibility set of β . The theorem of Reidys et al. ('97) states that the intersection of any two compatibility sets is non-empty.

One can compute the size of these intersections Weber ('97). Figure 7 provides a nontechnical illustration. A sequence is compatible with a structure α if positions that are supposed to pair in α are occupied by nucleotides that are actually able to pair. Satisfying two structures at once entails additional constraints. If the two structures are “independent,” meaning that positions that pair




Fig. 7. Satisfying two structures at once (after Weber '97). **Top:** Two structures that differ in a “shift”—a transformation in which one strand slides past the other (see Fig. 2). **Middle:** The same structures represented as circles with the base pairs drawn as chords. **Bottom:** The two circles superimposed to illustrate the additional constraints on a sequence that arise from having to satisfy both structures simultaneously. Position 1 must have a nucleotide that can pair with both position 13 and 12. But the nucleotide choices at 12 must take into account its complementarity to position 2 as well, and so on until the end of this chain is reached at position 9. Any such chain can be satisfied, and hence there exists at least one sequence compatible with any two structures. [A secondary structure can be thought of as a permutation acting on sequence positions. For example, the permutation corresponding to the structure on the left specifies that positions {5, 6, 7, 8, 9} are fixed points, and that 4 is assigned to 10, 10 to 4, 3 to 11, 11 to 3, etc. The two permutations corresponding to the two structures generate a dihedral group, and the chain of constraints in the bottom circle is one orbit obtained by the action of this group on the sequence positions (Reidys and Stadler, '96; Reidys et al., '97).]

in one structure are distinct from positions pairing in the other, the constraints simply add up. If, however, the two structures are non-independent, additional constraints arise because certain

positions must now satisfy two pairings simultaneously, and, depending on how the two structures are related, such constraints propagate along chains of dependencies. The constraints on nucleotides for a sequence that must satisfy two particular structures increase with the number and length of such chains. The most severe constraints arise from structure pairs that differ in shifts and their generalization (see Figs. 2 and 7). These highly constrained sequences are precisely those that enable the difficult shape transformations in the accessibility topology of Fontana and Schuster ('98a,b). These transformations correspond to the evolutionary steps of Figs. 5 and 6. Plasticity expedites these transformations only after the population finds these highly constrained sequences, that is, the intersection of the compatibility sets of α and β . Yet, the difficult transformations correspond to the smallest overlap regions. These regions are, in essence, as hard to find as the boundary between the neutral networks of α and β . The benefit gained once the population is in the overlap regions is negligible compared to the time required to find these regions in the first place. The search for these regions is made even more challenging by the requirement that the intermediary sequence must not only be compatible with α and β , but it must have α as its minimum free energy structure and β within $5kT$. Therefore, if the transformation from α to β is hard to achieve in the simple case, it is approximately as difficult to achieve in the plastic case. In this way, the failure of plasticity to expedite the discovery of new structures stems from the organization of the RNA genotype–phenotype map.

Environmental canalization (reduction of plasticity)

In Fig. 8 we monitor the time course of five population plasticity statistics for the plastic case. We discuss three of them here. Curve i in Fig. 8 shows the average fraction of the partition function realized within $5kT$, which is

$$z^j = \frac{1}{Z} \sum_{\sigma_i^j \in \sigma_i^j(\Delta)} \exp(-E(\sigma_i^j)/kT). \quad (5)$$

This indicates the total probability of the structures in $5kT$ relative to all structures attainable by sequence S_j , that is, how much of the partition function is accounted for by the $5kT$ neighborhood of the minimum free energy structure. Curve iv shows the average number of structures realized




Fig. 8. Plastic evolution. The time evolution of five plasticity statistics (i–v) are shown in conjunction with the major shape transitions (curves vi and vii). See text

for the definitions of the quantities monitored. The inset enlarges a population-level Simpson–Baldwin effect discussed in the text.

in $5kT$, $\langle |\sigma^j(\Delta)| \rangle$, while curve v tracks the average structure variance, $\langle \text{var } \sigma^j(\Delta) \rangle$, with

$$\text{var } \sigma^j(\Delta) = \overline{d(\sigma_i^j, \sigma_0^j)^2} - \overline{d(\sigma_i^j, \sigma_0^j)}^2, \quad (6)$$

where the bar indicates an average over the structures in $\sigma^j(\Delta)$. In eq. (6) we make use of the fact that the notion of variance can be generalized to sets (here of structures) for which a mean does

not exist, but for which one can compute pairwise distances between the elements. Finally $\text{var } \sigma^j(\Delta)$ conveys information about the structural diversity in the plastic repertoire of sequence S_j , while $|\sigma^j(\Delta)|$ simply counts the structures. These are three quantitative measures of RNA plasticity. We juxtapose these to adaptive events indicated by the evolutionary trajectories (curves vi and vii).

After an initial period, $\langle |\sigma^j(\Delta)| \rangle$ levels off to between 126 and 254 structures for the simple case

(not shown). In contrast, the plastic curve (iv) attains much lower, fairly constant levels of about 30, 13, and 20 structures. This shows a population-wide reduced plasticity compared to the simple case and the initial condition (204 structures in $5kT$). The loss of plasticity occurs very rapidly. The reduction is also evidenced by the average fraction of the partition function, $\langle z^j \rangle$ (curve i), which rapidly climbs to 0.9 and settles at 0.95, while the simple case remains in the range [0.65, 0.75] (not shown). The structure variance, curve v (range in the plastic case [7, 15], versus range in the simple case [46, 92]), indicates that alternative structures in the $5kT$ band become more similar to the ground state.

Transitions between evolutionary epochs are marked by dashes running vertically through all graphs. The dynamics surrounding a transition break down into four episodes: (a) the population is dominated by a structure with no indications of a better structure; (b) a sequence with a better suboptimal structure arises and dominates the population; (c) natural selection favors mutations that move that structure to the minimum free energy; (d) it then favors subsequent mutations that stabilize the new minimum free energy structure by eliminating variation in the remaining plastic repertoire. As described above, populations spend long periods of time in stage a. Stage b is rapid and typically does not transpire at the population level. Stage c, during which the new structure has become the minimum free energy structure but has not yet been stabilized, is well-illustrated by the dip in $\langle z^j \rangle$ immediately after transitions 2, 3, and 6 in Fig. 8. The subsequent rapid assimilation to pre-transition levels reflects the stage d. We see even more dramatic evidence for the fourth stage at the beginning of the evolutionary trajectory, immediately after transition 1, with the drop in the number of structures within $5kT$, $\langle |\sigma^j(\Delta)| \rangle$ (curve iv) in Fig. 8.

Example of a population-level Simpson–Baldwin effect

The four episodes surrounding transitions constitute the Simpson–Baldwin effect. As discussed in the Introduction, there is a discovery phase (stages a and b) and an assimilation phase (stages c and d). Transitions 4 and 5 in Fig. 8 are somewhat unique in that the entire Simpson–Baldwin effect is manifested at the population-level, although not strictly chronologically. At transition 4, a new advantageous sequence species S with a suboptimal structure β that is closer to the target

than its minimum free energy shape α begins to invade the population. After 10^6 replication events the first sequences S' appear with β as their minimum free energy structure. The population retains α , however, as a frequent suboptimal shape with high Boltzmann weight. Mutation shifts around weights between α and β . Although α is selectively worse than β , this can be offset by a higher Boltzmann probability. Shape α thereby succeeds in remaining the most prevalent minimum free energy structure. This shows up as a discrepancy between the population averages of the weighted distance to target (curve vii), $\langle \overline{d(\sigma^j, \tau)} \rangle$, and the distance based only on the minimum free energy structures (curve vi), $\langle d(\sigma_0^j, \tau) \rangle$; see the enlarged inset at the bottom of Fig. 8. Upon 2×10^5 further replications, sequences with β as the minimum free energy structure take over the population, and $\langle d(\sigma_0^j, \tau) \rangle$ consequently drops toward $\langle \overline{d(\sigma^j, \tau)} \rangle$. The situation reverses, however, as Boltzmann probabilities shift again. After 6.5×10^5 replications, β disappears completely as a minimum free energy structure and is observed only as a suboptimal. Finally, 3×10^5 replications later the final transition back to sequences with β as minimum free energy structure occurs, accompanied by a realignment of $\langle \overline{d(\sigma^j, \tau)} \rangle$ with $\langle d(\sigma_0^j, \tau) \rangle$. The plot of $\langle z^j \rangle$ indicates a significant subsequent reduction in plasticity that stabilizes β (label “f” in graph i of Fig. 8). In terms of minimum free energy structures, events 4 and 5 taken together constitute one single transition that occurs in two stages. Event 4 marks the onset of the Simpson–Baldwin discovery phase which lasts until event 5 when the assimilation phase begins. This is the only example we encountered of both Simpson–Baldwin phases being revealed at the population level.

2.3. Summary

A quantitative characterization of simulated RNA evolution under the plastic and the simple genotype–phenotype maps, shows the following points:

- The punctuated character of the evolutionary dynamics holds for RNA populations evolving under both the simple and the plastic genotype–phenotype map.
- The populations evolving under the plastic map halt evolutionary progress much farther away from the target structure than those evolving under the simple map.
- Plastic populations show no evolutionary speed-up (or Baldwin expediting effect); overall they exhibit a strong slow-down.

- A strong reduction in plasticity occurs in the plastic case. The Simpson–Baldwin effect is observed along individual founding lineages causing the phenotypic transitions. The discovery phase, however, typically occurs rapidly and is not evident in population-level statistics. We observed one example of a Simpson–Baldwin effect that was entirely apparent at the macroscopic population level.

3. PLASTOGENETIC CONGRUENCE AND NEUTRAL CONFINEMENT

3.1. Plastogenetic congruence

Section 2 demonstrates that a plastic population evolving toward a constant target undergoes a drastic narrowing of plastic repertoires and ultimately hits an evolutionary dead-end. In the cur-

rent section, we construct a causal bridge from the loss of plasticity to the loss of evolvability by arguing a significant correlation between the structures available to a sequence in its plastic repertoire and the minimum free energy structures present in the mutational vicinity of that sequence. First we use a concrete example to provide intuition for the mechanistic underpinnings of such a correlation. Then through statistical evidence, we claim that it is a fundamental property of the genotype–phenotype map.

The top left corner of Fig. 9 shows a short sequence *A* and its minimum free energy structure α together with the list of all suboptimal configurations within $5kT$ of the minimum free energy. While the sequence spends 52% of the time in α , its relative stability is only marginal, as it competes with a number of energetically close subop-




Fig. 9. Plasticity and minimum free energy shapes of genetic neighbors. The figure illustrates how the plasticity of a sequence correlates with the minimum free energy structures of one-error mutants. Arrows show point mutations. The mu-

tation from *B* to *B'* involves the same nucleotide substitution at the same position as the mutation from *A* to *A'*. The only difference is the slight change in context due to the neutral mutation from *A* to *B*. See text for details.

timal structures. These alternative structures are energetically easy to access because base pair interactions holding together the stem in α are weak. For example, the entire stem can switch into a different position, such as in the 4th suboptimal configuration β .

Significantly, α 's marginal stability coincides with an increased sensitivity to mutations. For example, mutating a position inside α 's loop (arrow) tips the thermodynamics in favor of β . The result is a change in the minimum free energy structure from α to β (upper right corner of Fig. 4).

Now consider a neutral mutation—one that does not displace α as the minimum free energy structure—transforming A into B (lower left corner of Fig. 4). The mutation strengthens the thermodynamic stability of the stacking region of α . The Boltzmann weight of α increases from 52% in sequence A to 88% in B . Notice that the set of alternative structures available to B within $5kT$ has become considerably smaller compared to A . In addition, the few alternative structures that occur with appreciable probability have become more similar to the ground state. The relative stability of β (which is still compatible with the mutated sequence) has dropped by one order of magnitude and disappeared from the set of $5kT$ suboptimal structures in B . This coincides with an increased robustness of α towards mutations. In fact, the same nucleotide change that led from A to A' , changing its minimum free energy structure, leaves α invariant when applied to sequence B .

The upper arrow from A to A' illustrates the correlation between structures in the plastic repertoire and structures in the mutational vicinity. In particular, this point mutation changes a ($5kT$) suboptimal structure into the minimum free energy structure. It is not difficult to imagine that a sequence with β in its plastic repertoire is much more likely to have a genetic neighbor with β as the minimum free energy structure, then a sequence that lacks β in its plastic repertoire. The arrows from A to A' and further to B' illustrate the epistatic control of neutrality (section 1.5). This is a special case of the correlation just described: the more time a sequence spends in its minimum free energy structure, the higher the fraction of neutral neighbors in its genetic vicinity.

To summarize, we call a shape σ “plastically accessible” to a sequence S if σ is in the plastic repertoire of S . Likewise, σ is “genetically accessible” from S , if there exists a one-error mutant S' of S , such that σ is the minimum free energy struc-

ture of S' . The alignment of plastic accessibility with genetic accessibility means that the set of shapes into which a particular sequence can fold (plastic accessibility) strongly correlates to the minimum free energy shapes realized by its one-error mutants (genetic accessibility). We call this property of the genotype–phenotype map “plastogenetic congruence.”

The genetic accessibility of phenotypes constrains an evolutionary trajectory under the simple map (Fontana and Schuster, '98a,b). In the plastic simulations, however, plastic accessibility, through its alignment with genetic accessibility, has an equally profound impact on the evolutionary dynamics. We return to this after a quantitative demonstration of plastogenetic congruence for the RNA folding map.

Ideally we would perform statistics on a broad sampling of genotypes to show the extent of overlap between plastic repertoires and minimum free energy structures of genetic neighbors. In lieu of this computationally prohibitive approach, we present three pieces of partial evidence for plastogenetic congruence:

1. The frequency of a structure β as a minimum free energy structure among the one-mutant neighbors of a sequence S is significantly larger for sequences that have β in their plastic repertoires than for sequences that do not.
2. The minimum free energy structure α of a sequence S is present at high frequency in the plastic repertoires of one-mutant neighbors of S .
3. For any advantageous shape β in the plastic repertoire of any sequence S , S can typically evolve to another sequence S' with β as its minimum free energy structure, in only (on average) five point mutations.

In the first approach we generate random sequences S_α with a particular structure α as the minimum free energy structure [“inverse folding,” see Hofacker et al. ('94)]. Since two sequences that share the same minimum free energy structure often share other suboptimal structures, this simple procedure yields sets of sequences S_α^β that share the same ground state α and that also have a particular shape β in their $5kT$ -plastic repertoire. We simultaneously obtain control samples of sequences S_α that specifically lack the shape β from their plastic repertoire. We then compute the minimum free energy shapes of all one-mutant neighbors of sequences in the sample S_α^β and in

the control sample S_α , and compare the frequencies of sequences with β as the minimum free energy shape. Although a systematic exploration of the possible shape combinations α and β is unfeasible, Table 1 provides anecdotal evidence for our proposition: A sequence (with minimum free energy structure α) that has a structure β among its $5kT$ suboptimals is much more likely to have a one-error mutant with β as its minimum free energy structure than a sequence (with minimum free energy structure α) that lacks β among its $5kT$ suboptimals.

Table 1 shows ratios of likelihoods. The top seven structures β are in the neighborhood of α in the sense of the shape topology developed in Fontana and Schuster (98a,b). We are able to attain larger samples for the top structures as they appear more frequently as suboptimals of sequences with minimum free energy structure α . The likelihood ratios indicate that these suboptimal structures are easily converted into minimum free energy structures by a single point mutation.

The last row of Table 1 shows the only counterexample we found when α is the tRNA cloverleaf. We observe similar patterns for other ground states α , and for sequence lengths other than 76.

In a second approach, we generate a sample of sequences with minimum free energy structure α (a “neutral set” of α), and check for the presence of α in the plastic repertoires of sequences obtained by one point mutation from the neutral set. We search and calculate statistics over only the one-mutant neighbors that are compatible (see section 2.2) with α , since only these can have α as a suboptimal in the first place. (Given a sequence S , the fraction of one-error mutants compatible with α is $1 + (n_{GU} - 5n_{bp})/3n$, where n is the sequence length, n_{bp} is the number of base pairs in α and n_{GU} is the number of GU pairings that would occur when S folds into α .) While sequences that are compatible with α have, by definition, α among their suboptimals at some energy, we are only interested in the limited plasticity range of $0 \leq \Delta \leq 10kT$ from the minimum free

TABLE 1. Plastogenetic congruence I^1

Ground state (α)		
((((...(((.....))))).((((.....)))).....((((.....))))).))))....		
Suboptimal shape (β)	o-ratio	n-ratio
((((...(((.....))))).((((.....)))).....((((.....))))).))))....	9.0	8.0
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	7.4	6.7
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	6.6	4.9
.((((.....(((.....))))).((((.....)))).....((((.....))))).))))....	6.2	4.8
.....(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	14.0	3.9
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	4.0	3.3
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	3.5	2.8
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	15.8	14.0
..(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	9.7	10.8
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	7.5	7.5
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	10.8	7.1
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	7.0	6.2
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	7.3	5.7
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	6.3	5.5
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	7.0	4.7
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	5.2	4.2
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	2.8	3.2
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	2.7	2.4
(((.....(((.....))))).((((.....)))).....((((.....))))).))))....	0.5	0.4

¹Samples of sequences with a tRNA cloverleaf as the minimum free energy structure α (bottom line) and various structures β as suboptimal configurations were generated as described in the test. The frequency of β becoming the minimum free energy structure upon one point mutation was computed for the sample of sequences having β as a suboptimal and for a control sample of sequences lacking β within a $5kT$ range from α . The ratio of the former to the latter is tabulated. Two kinds of frequencies were computed: the frequency with which β occurs as a ground state among all one-error neighbors (“occurrence frequency” and corresponding o-ratio) and the frequency with which it occurs at least once in a 1-error neighborhood (“neighborhood frequency” and corresponding n-ratio). For example, row 1 states that the shown shape occurs 8 times more frequently as a minimum free energy structure in the one-error neighborhood (n-ratio) of a sequence that has that shape as a suboptimal compared to one that lacks it. The first part of the table is based on samples with more than 200 sequences with β as a suboptimal, while the second part is based on sample sizes of more than 100 but less than 200 sequences. A given sequence can have several suboptimal configurations listed in this table.

energy. Structures outside this range have Boltzmann probabilities too low to influence fitness. Figure 10 shows the fraction of sequences that have α , the tRNA cloverleaf, among their suboptimal configurations and that are located in the compatible one-error boundary of a tRNA neutral set. This fraction is shown as a function of the energy interval Δ . The fraction of sequences in the sample that have at least one one-mutant neighbor with the tRNA as a suboptimal (curve *i*) hits 1 at only $0.6kT$, while the fraction of all one-error mutants (curve *ii*) is above 0.8 at $5kT$, and reaches 1 at $9.6kT$ (for $T = 37^\circ\text{C}$).

The first numerical observation demonstrates that the occurrence of α among the plastically accessible configurations of a sequence indicates the immediate vicinity of α 's neutral network. The second method shows that α 's neutral network casts a "shadow" into the energetically low lying suboptimal configurations realized by sequences in its one-error compatible boundary. Together these facts suggest that plastic accessibility and genetic accessibility mirror each other in our RNA model. From a biophysical point of view, this matches intuition. It means that the thermodynamic stability of a suboptimal structure α that is realized with non-negligible probability can frequently be improved (even to the point of making α the




Fig. 10. Plastogenetic congruence II. For each point in the graph, we generate a new sample of 100 sequences with the tRNA cloverleaf as their minimum free energy structure. All one-error mutants were scanned for the presence of the tRNA cloverleaf as a suboptimal configuration. The graph shows the frequency of such sequences as a function of the energy range Δ that defines the plastic repertoire. Curve *i* refers to the frequency with which a sequence in the sample has at least one one-error mutant with the desired suboptimal structure ("neighborhood frequency"), while curve *ii* shows the "occurrence frequency." See text for details.

ground state) by one suitably placed point mutation. This occurs "positively" through mutations in the base pairing positions of α to nucleotides that yield better stacking energies; and through similar modifications to mismatches at the termini of helical regions. It also occurs "negatively" through mutations that outright eliminate energetically competing structures by making the sequence incompatible to them. Intuitively, the mutational stabilization of structures toward which a sequence is already predisposed occurs more readily than the construction of a minimum free energy structure from scratch.

The evolutionary reduction of plasticity proceeds if plasticity entails a fitness cost *and* mutation produces similar but less plastic alternatives to existing phenotypes. The latter relies on plastogenetic congruence. Figures 11 and 12 emphasize this evolutionarily enabling function of plastogenetic congruence. We generate a sample of sequences with a given α as a $5kT$ suboptimal configuration and no constraints on the minimum free energy structure. The likelihood of obtaining sequences with a predefined α among their $5kT$ repertoires can be tuned by enriching those sequence segments that should fold into the stacking regions of α with stabilizing GC pairs (see inset of Fig. 11).

We let each sample sequence be the starting point of a gradient walk. At each step of a walk, the one-error mutant is chosen that most increases the Boltzmann probability $P(\alpha)$ of α . Figure 11 shows the distribution of the walk lengths until α becomes the ground state and until $P(\alpha)$ cannot further be improved. The graphs at the left of Fig. 11 confirm that there is a large probability of making α the ground state in a single point mutation. Furthermore, this one-step probability is lower for samples that permit α to lie higher in the energy interval Δ . The difference is one order of magnitude as Δ increases from $5kT$ to $8kT$. A higher GC content in stacking regions also increases the one-step probability, since any bias toward α helps its evolution to lower free energy.

As shown in Fig. 11, α usually becomes the minimum free energy structure in only two to three mutations, yet it takes many more mutations for the thermodynamic stability $P(\alpha)$ to attain its local maximum.

Figure 12 shows the distributions of $P(\alpha)$ when α has become the minimum free energy structure for the first time and when $P(\alpha)$ is a local maximum. Two aspects are worth noting. First, the degree to which the thermodynamic sharpness of

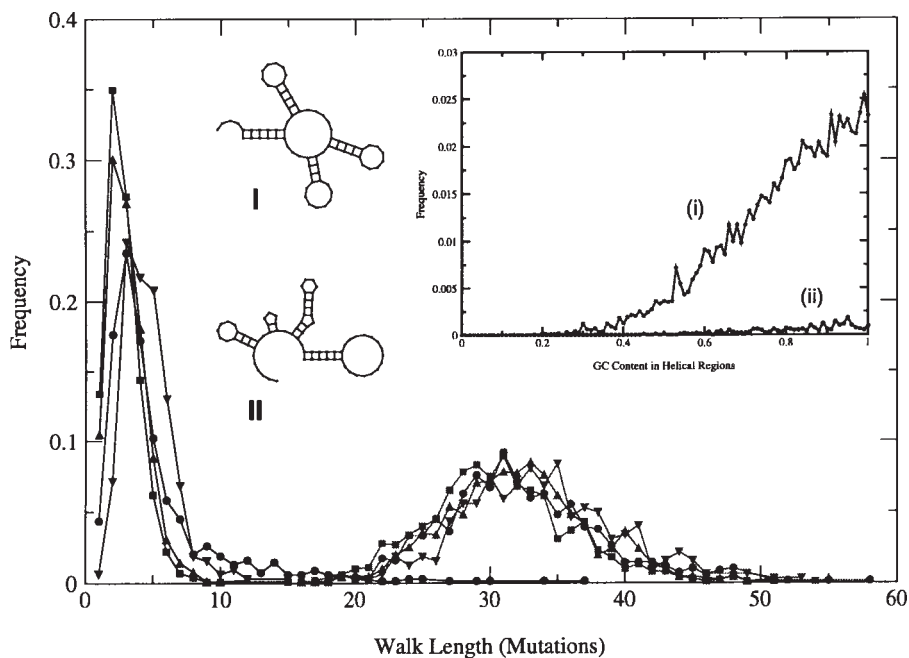


Fig. 11. Plastogenetic congruence III. The inset shows the frequency with which a random sequence compatible with a tRNA cloverleaf (structure **I**) has that structure among its $5kT$ suboptimal configurations (curve *i*). This frequency is shown as a function of the fraction of GC or CG pairs in the sequence segments that fold into the stacking regions of the cloverleaf shape. Curve *ii* is the frequency with which such sequences have structure **I** as a minimum free energy structure. The main graph shows the distribution of gradient walk lengths as described in the text. There are two sets of curves, black (left hump) and gray (right hump). The left set pertains to walk lengths up to the first appearance of the given struc-

ture (**I** or **II**) as a minimum free energy structure. The right set is the distribution of walk lengths until the Boltzmann probability of that structure could not be further improved. Only walks that terminated within 100 steps are considered. Circles pertain to structure **II**, $\Delta = 5kT$, GC fraction in helical regions is 0.5, 1,000 distinct walks were performed of which 686 terminated. Squares pertain to structure **I**, $\Delta = 5kT$, GC fraction is 0.5, 1,000 walks of which 999 terminated. Up-triangles pertain to structure **I**, $\Delta = 5kT$, GC fraction is 0.33 (i.e., no bias), 916 walks of which 909 terminated. Down-triangles pertain to structure **I**, $\Delta = 8kT$, GC fraction is 0.33 (i.e., no bias), 322 walks of which 322 terminated.




Fig. 12. Plastogenetic congruence III. The frequency distributions of Boltzmann probabilities for structures **I** and **II** are shown. See Fig. 11 for details. The left set of curves shows

the Boltzmann probabilities when the structures first appear as the ground state, and the right set shows the Boltzmann probabilities at local optima of the gradient walk.

tRNA structures (for example) can be improved is large. Second, the final $P(\alpha)$ distribution is much narrower and the average $P(\alpha)$ attained is substantially closer to 1 for a tRNA structure (I) than for another structure (II) chosen randomly. Only 68% of the gradient walks found within 11 steps a sequence with the shape II as the minimum free energy structure. This suggests structure dependent limits to canalization.

In sum, this third perspective on plastogenetic congruence looks beyond the one-error neighborhood. If a conversion of a suboptimal shape into a minimum free energy structure cannot be achieved in a single step, it can occur gradually, over several steps. Once a structure has become the ground state, there is still room for a significant reduction of plasticity.

3.2. Evolutionary downside of plastogenetic congruence

We have seen that plastogenetic congruence enables genetic assimilation, in that mutations easily move a good structure from a minor position in a plastic repertoire to the minimum free energy position. There is, however, a flip side to plastogenetic congruence. When the reduction of plasticity is linked to the reduction in genetic variability, plastogenetic congruence stalls evolution.

We now consider the neutrality of the sequences obtained through our gradient walks. Recall that neutrality is the fraction of single base mutations that preserve the minimum free energy structure. The top graph of Fig. 13 shows four examples of gradient walks, chosen for their diversity of approaches to a local optimum. The walk criterion optimizes the Boltzmann probability of the structure depicted at the bottom of Fig. 13. (This structure dominated the population when the evolutionary process of Fig. 6, section 2.2, became trapped.) The bottom part of Fig. 13 shows the concurrent changes in neutrality along these same gradient walks. First a structure descends to the minimum free energy configuration, and the neutrality drops to a minimum. Then the neutrality increases sharply to levels above 0.4. This value lies in the tail of the distribution of neutralities for all sequences with that minimum free energy structure (not shown). The mean of this distribution is 0.3 (Matt Bell, personal communication, August, '99). Note that the fraction of one-error neighbors compatible with the shown structure is approximately 0.55. (The exact compatibility fraction depends on the U content of a sequence.) One walk (up-triangles) approaches but does not sustain this upper bound for neutrality.

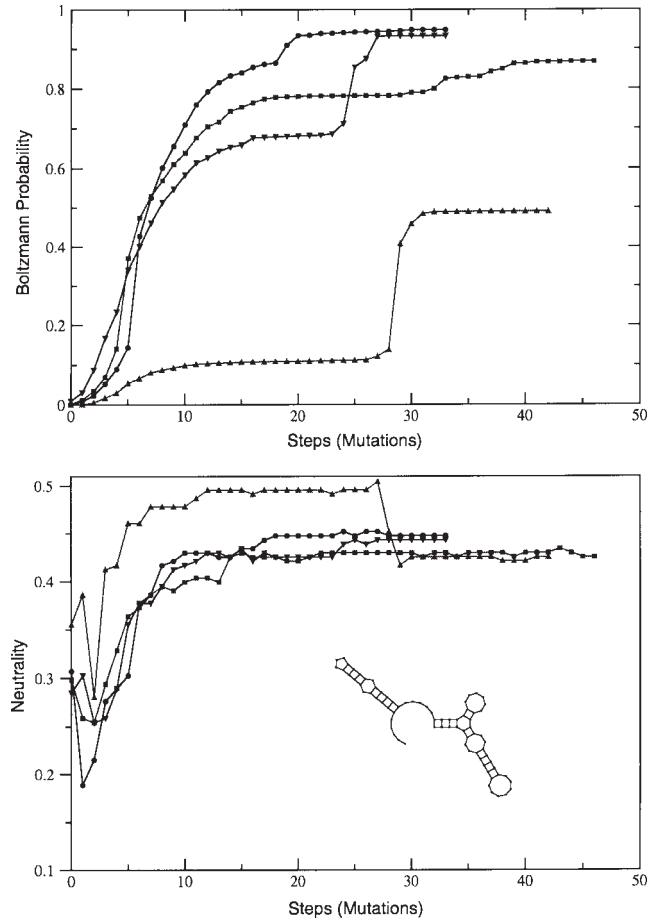


Fig. 13. Gradient walks in plastogenetic space. Gradient walks are generated by moving from the current sequence to its one-error mutant that most improves the Boltzmann probability of a prespecified structure. Walks begin with sequences that have that structure within their $5kT$ plastic repertoire. **Top:** Progress profiles along sample walks. **Bottom:** Concurrent development of neutrality (fraction of one-error mutants with the same minimum free energy structure) along these walks.

After the desired structure has moved into the minimum free energy position, point mutations that increase its Boltzmann probability typically also increase the neutrality of the sequence. These changes have an epistatic effect of the kind we define in section 1.5 and Fig. 4.

Recall that neutrality indicates robustness to genetic modification and the Boltzmann probability of the minimum free energy structure indicates robustness to thermodynamic perturbation. As such, we conclude that genetic canalization occurs in tandem with environmental canalization. In these walks and in our full-blown simulations, environmental canalization—the reduction of plasticity—is the directly selected response to selection

in a population. Genetic canalization, however, is not an adaptation. It is instead a byproduct of environmental canalization, to which it is linked by common genetic underpinnings. The possibility that environmental and genetic canalization share a genetic basis which might account for the evolution of canalization was recently hypothesized by Wagner et al. ('97). Our study of the evolution of plasticity in RNA provides a mechanistic realization for this hypothesis.

Figures 14 and 15 provide support for the simultaneous drops in plasticity and evolvability. The first of these figures illustrates the surprising extent of plasticity reduction in RNA. It compares the density of structural states of three sequences that have been obtained through different processes but that share the same minimum free energy structure. The common minimum free energy shape is the characteristic shape α found at the dead-end of the evolutionary process depicted in Fig. 6 (section 2).

The random sequence obtained by inverse folding α (Hofacker et al., '94) has 574 different structures within 3 kcal of α . It spends only 3% of the time in α , and the combined probability of the 574 alternative configurations accounts for only 58% of the partition function. The different configurations cover a wide range of structural diversity.

The second sequence is the product of evolution under the simple map. We inoculated a simulated flow reactor with the aforementioned inverse folded sequence, and specified its minimum free energy structure α as the target structure. This leaves no room for phenotypic improvement. Despite the absence of direct selection pressure on the well-definition of the ground state, we observe a reduction of plasticity by one order of magnitude. In van Nimwegen et al. ('99a), we learn that a population evolving on a neutral network under the simple map will concentrate on sequences with higher than average neutrality (see Figs. 15 and 18). Because mutations off the neutral network yield, on average, much lower fitness phenotypes, there is indirect selection against sequences in the neutral network that have a high fraction of one-error mutants off the network, i.e., low neutrality. This is a second order effect that depends on the probability of deleterious mutation. Similar observations were made with a model of genetic regulatory networks by A. Wagner ('96). The evolved sequences that have become mutationally more robust also exhibit a correspondingly lower plasticity. This is a manifestation of plastogenetic congruence, which is (weakly) effective even under the simple map.

The last sequence in Fig. 14 is evolved under the plastic map where reduced plasticity is the direct consequence of natural selection. The size of the plastic repertoire decreases by another order of magnitude to only four structures in addition to the ground state which is occupied 67% of the time. If the molecule equilibrates over all its states, it spends 94% of its time in these five configurations that are highly similar to each other.

Figure 15 shows the neutral positions of these sequences. We call a position neutral if at least one mutation at that position leaves the minimum free energy structure unchanged. Figure 15 additionally includes three sequences obtained similarly for the structural end point in the evolutionary process of Fig. 5, section 2. For the structures on the left, the neutrality of the sequences increases from 0.184 for the random sequence to 0.412 for the neutrally evolved sequence to 0.456 for the canalized sequences. Similarly on the right, the neutrality increases from 0.158 to 0.311 to 0.430 from top to bottom. The neutral coverage in the canalized case is almost perfect. That is, a position will tend to be neutral if it permits a mutation that does not destroy the compatibility of the sequence with its minimum free energy structure. The intermediate neutrality of the neutrally evolved sequences reflects the second-order selection for increased neutrality under the simple map as discussed above. The further increase in neutrality that occurs under the plastic map is a side effect of the sharp reduction in plasticity.

3.3. Neutral confinement

Plastogenetic congruence provides both the mechanism for an evolutionary reduction of plasticity and the link between this reduction and the ultimate evolutionary dead-end. Plasticity is costly because more structures in the plastic repertoire implies less time spent in any one. Plastogenetic congruence enables populations to reduce these fitness costs through genetic assimilation—the movement of an advantageous suboptimal structure to the minimum free energy structure, and subsequent reduction in plasticity. At the same time, plastogenetic congruence causes a loss in variability. In this section we closely examine the extent to which populations end up genetically isolated from phenotypic novelty.

Figure 16 compares the frequency distribution of neutralities at the end of the plastic simulations depicted in Figs. 5 and 6 (filled circles) and the corresponding simple runs (down-triangles). (The data represented by squares are discussed

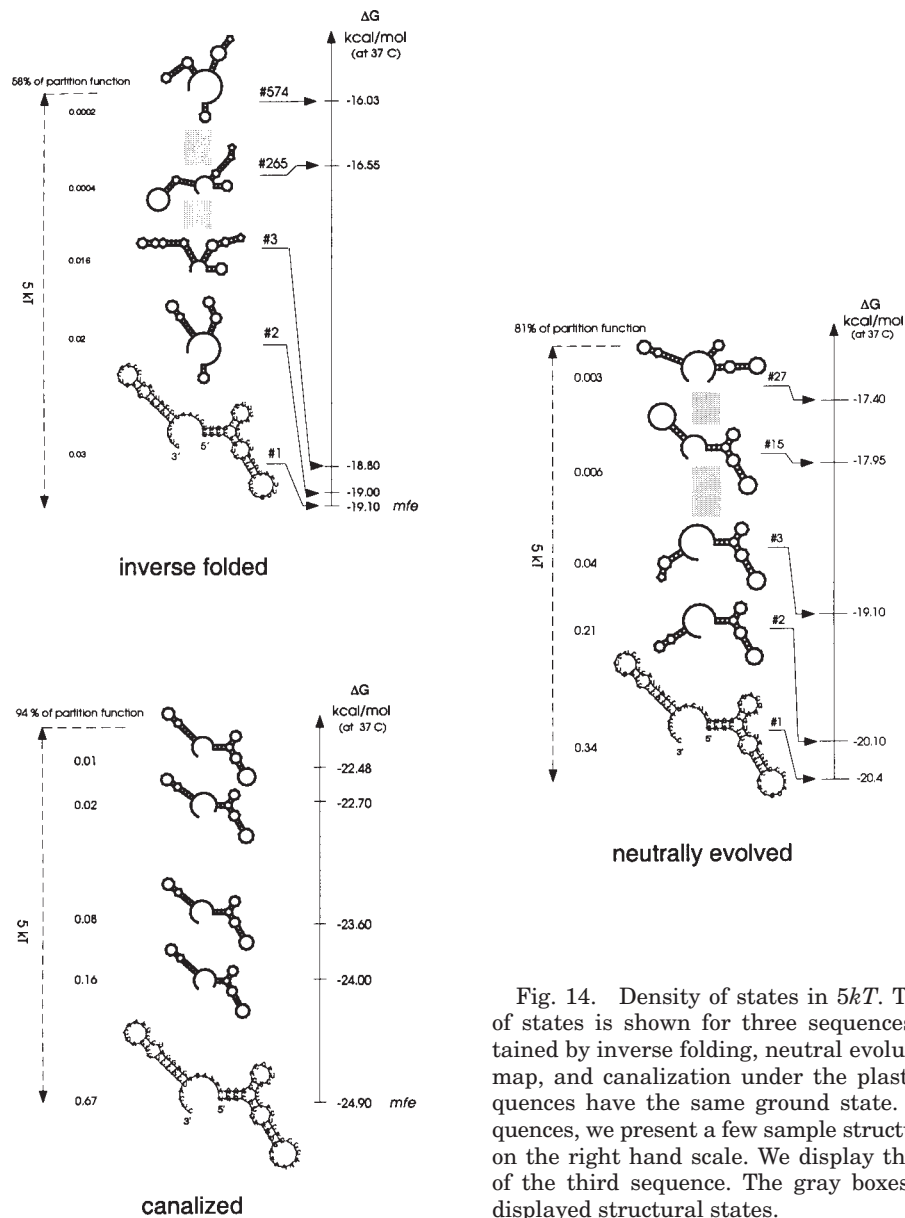


Fig. 14. Density of states in $5kT$. The structural density of states is shown for three sequences that have been obtained by inverse folding, neutral evolution under the simple map, and canalization under the plastic map. All three sequences have the same ground state. For the first two sequences, we present a few sample structures with their energy on the right hand scale. We display the complete repertoire of the third sequence. The gray boxes indicate many non-displayed structural states.

below.) Again, the neutrality of a sequence is the fraction of its one-error mutants that share its minimum free energy structure. It indicates the lack of phenotypic variability of the sequence. The plastic population evolves to a distribution of neutralities with a much higher average and lower kurtosis than that of the non-plastic population.

Sequences with high neutrality have only a small proportion of distinct phenotypes in their mutational vicinities, and are therefore unlikely to mutate towards a higher fitness phenotype. Furthermore, we argue that the minimum free energy structures of the few non-neutral neighbors offer little phenotypic diversity, making the discovery of novel advanta-

geous mutants even more improbable. High neutrality correlates with high structural similarity between the ground state structure and the other configurations in the plastic repertoire (Wuchty et al., '99). By plastogenetic congruence, this comes to mean that a one-error mutant of a low-plasticity sequence either folds into the same minimum free energy structure (neutrality) or into a structure that is very similar to it. Very similar here means that the structures differ slightly with respect to stack lengths or loop sizes. As a consequence, the discovery of new advantageous shapes is considerably slowed down and eventually halted. We say the population is "neutrally confined."




Fig. 15. Neutral positions. The emphasized positions are those with at least one neutral nucleotide substitution. We highlight these neutral positions for inverse folded, neutrally

evolved, and canalized sequences on two minimum free energy structures.

Highly neutral regions of a neutral network appear to be wide-spread and connected (Matt Bell, personal communication). Genetic diversification, therefore, still occurs within these regions. Yet the extent of genetic variation is insufficient to produce difficult shape transformations, such as a shift or the de novo creation of a stack (Fig. 2).

Neutral confinement in RNA seems independent of the mutation rate. We simulate evolution on a neutral network by starting a population with sequences that have the designated target structure as their minimum free energy structure. Figure 17 depicts such simulations that use the structure shown at the top of Fig. 16 as both the start-

ing minimum free energy shape and the target. Figure 17 monitors the fraction of sequences with that structure as a function of the replication accuracy per position. This is done for two values of the superiority (see caption for Fig. 17) and for both the plastic and the simple map. The phenotypic error threshold (Huynen et al., '96; Reidys et al., '98) is the replication accuracy at which that minimum free energy structure is lost from the population. At the same time we monitor the average neutrality of sequences with that structure. The independence of the average neutrality from error rate and superiority was predicted for the simple case by van Nimwegen et al. ('99a). In the

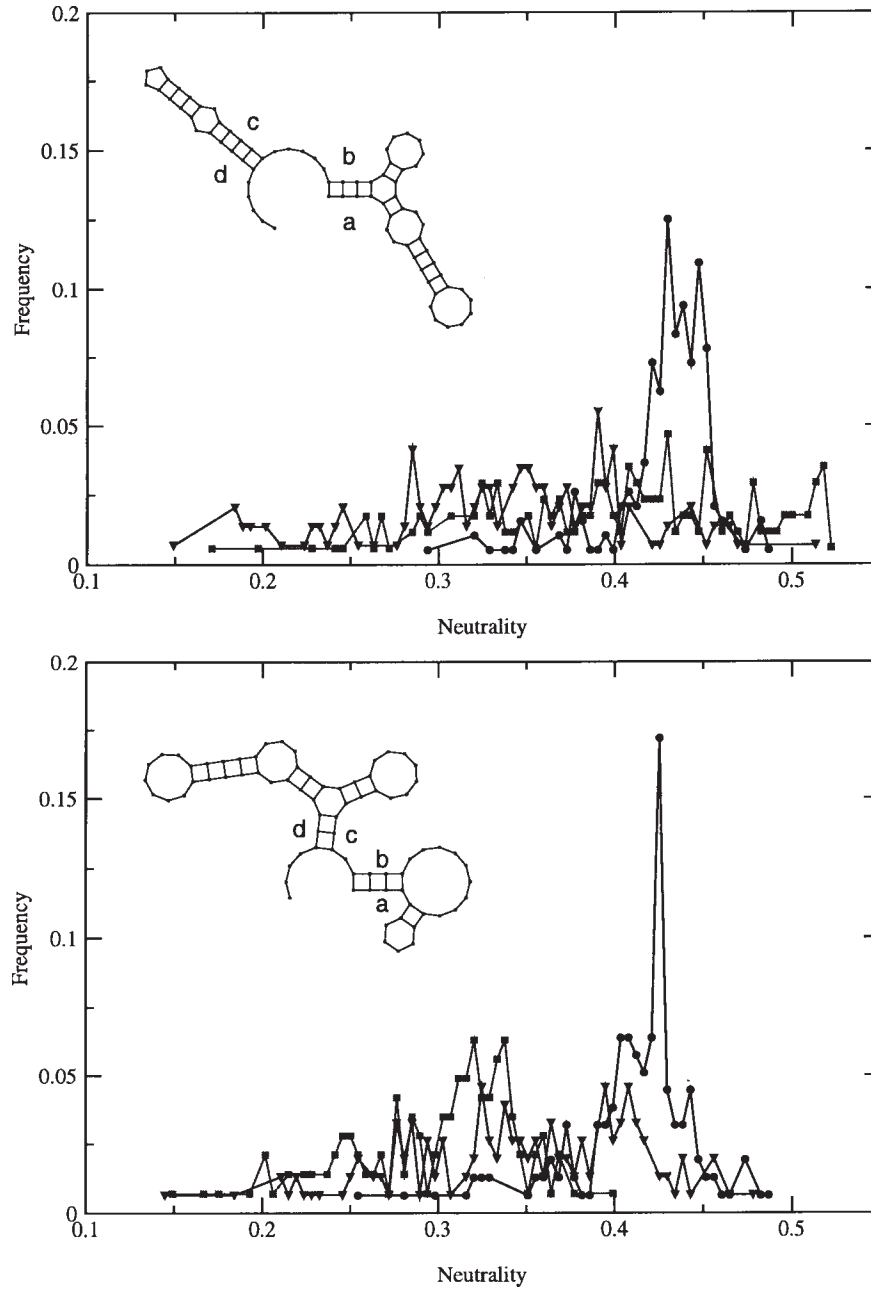


Fig. 16. Neutral confinement. We measure the neutrality of each sequence *species* in the population support and plot a frequency distribution. This just means that each sequence type is weighted the same, irrespective of its frequency in the population. Frequency-weighted plots look similar, with more dramatic high-neutrality peaks and a subdued low-neutrality spectrum. **Top:** Populations pertaining to the simulation of Fig. 6. Filled circles: population support neutrality after 4.2×10^6 replications under the plastic map (the structure shown in the inset dominates). Down-triangles: neutral-

ity at the end of the simple run (Fig. 6) with the simple map. Squares: neutrality of a population that has evolved for 5.6×10^6 replications from the one underlying the filled-circle data when plasticity was switched off. **Bottom:** Similar analysis for populations pertaining to the simulation shown in Fig. 5. Filled circles: with plasticity after 30×10^6 replications. Down triangles: simple run without plasticity after 32×10^6 replications. Squares: evolved from the filled-circle data for 4.4×10^6 replications with plasticity switched off.




Fig. 17. Neutrality as a function of replication accuracy. Populations were obtained by neutral evolution with different replication accuracies on the structure α shown at the top of Fig. 16. The graphs that decay exponentially with decreasing replication accuracy represent the stationary frequency of the target phenotype α in the population. The nearly constant graphs monitor the average neutrality (in the population support sense of Fig. 16; essentially the same figure obtains when sequences are weighted by their frequency) of the sequences on the “master network” (i.e., whose minimum free energy configuration is the target α). The replication accuracy at which the master network is lost is known as the phenotypic error threshold (Huynen et al., '96; Reidys et al., '98). For infinite populations the threshold accuracy q_{\min} depends on the average neutrality of the α -network, λ_{α} , and the superiority of α , σ_{α} , indicating how much better the replication rate of α , r_{α} , is with respect to the remaining phenotypes in the population (Reidys et al., '98; Schuster and Fontana, '99): $q_{\min} = ((1 - \lambda_{\alpha}\sigma_{\alpha}) / (1 - \lambda_{\alpha}))^{1/n}$ with $\sigma_{\alpha} = r_{\alpha} / \sum_{\beta \neq \alpha} r_{\beta} x_{\beta} / (1 - x_{\alpha})$, where x_{β} is the fraction of sequences with ground state β . In the plastic case, r_{α} is not constant, since different sequences with α as a ground state will have different sub-optimal configurations. We replace r_{α} by the average over such sequences in the population. The graphs for the plastic case always dominate those of the simple case (higher neutrality and hence lower threshold accuracy). Filled circles: neutral evolution at superiority 10. Down-triangles: neutral evolution at superiority 1.5. Solid curves are linear regressions and exponential fits. Note the near constancy of the population neutrality in the plastic case regardless of the superiority and error rate.

plastic case, however, the independence is rather unexpected. In section 4 we present formal models of evolution under the plastic map. They predict a shift in the equilibrium distribution from high neutrality (of neutral confinement) to lower neutrality at a mutation rate higher than the error threshold. This suggests that as the mutation rate increases, the population goes directly from neutral confinement into falling completely off the neutral network (error catastrophe).

To summarize, we have demonstrated that

plastogenetic congruence plays three critical roles. First it enables the genetic assimilation of sub-optimal shapes into ground states. Second it facilitates the reduction of plasticity. This occurs under stabilizing selection following a transition, in the early periods of apparent stasis when the genetically accessible phenotypes are either neutral or of lower fitness. Third, during this reduction, plastogenetic congruence yields the advance towards sequences with low phenotypic variability—sequences in an evolutionary dead-end.

3.4. Fitness landscape perspective

Our analysis repeatedly appeals to neutrality with respect to minimum free energy structures. Under the simple map from genotype to minimum free energy structure, such neutral networks are also invariant with respect to phenotype and fitness. Plasticity adds texture to these neutral networks in terms of both phenotype and fitness. Two sequences that share a minimum free energy structure will have divergent structures in their remaining plastic repertoires, and hence different fitnesses. As discussed in section 1.6, we nevertheless continue to refer to such networks as neutral.

When populations of plastic sequences arrive at an evolutionary dead-end, they are confined to the most neutral recesses of a neutral network. These regions correspond to local fitness maxima under the plastic map. One might therefore argue, that these dynamics can be *solely* conceived in terms of evolution toward a local fitness optimum. We counter that the ruggedness of the fitness function that results from the plastic map is remarkable by virtue of plastogenetic congruence.

The fitness peaks at which plastic populations come to rest correspond not just to good phenotypes (that is, repertoires with a stable minimum free energy structure close to the target), but also to regions that are mutationally isolated. Plastogenetic congruence means that the set of genotypes with highest fitness in a neutral network are mutationally highly interconnected with very small boundaries (if any) with other neutral networks.

Consider a fitness function in which a randomly chosen connected sub-network of a neutral network was assigned high fitness relative to the rest of the network. We assert that a population evolving under such a fitness function is much less likely to be trapped than a population evolving under our plastic fitness function. Arbitrarily chosen fitness peaks of this kind will not suffer the mutational buffering induced by plastogenetic congruence.

3.5. Relationship to work in protein folding

A correlation between the thermodynamic stability of a minimum free energy structure and its mutational robustness has been observed in lattice models of protein folding (Bussemaker et al., '97; Vendruscolo et al., '97; Govindarajan and Goldstein, '97; Bornberg-Bauer and Chan, '99), and in RNA models (Wuchty et al., '99). "Thermodynamic stability" refers to the free energy ΔG (Govindarajan and Goldstein, '97; Bornberg-Bauer and Chan, '99), the thermal stability (melting temperature) of the ground state structure (Vendruscolo et al., '97), or the energy gap between it and the first excited state (Bussemaker et al., '97). We relate our findings to this body of work in two respects.

Although plastogenetic congruence in RNA implies an analogous correlation between thermodynamic and mutational stability, it has broader implications. The plasticity of an RNA molecule refers to alternative *structures* in the vicinity of its ground state at constant temperature. [Bornberg-Bauer and Chan ('99) use the term "plasticity" unconventionally to refer to mutational (in)stability, rather than to alternative phenotypes of a single genotype in response to the environment.] Plastogenetic congruence states that the alternative structures available to a given sequence are frequently found as the minimum free energy structures of its mutational neighbors. This indicates *which* minimum free energy structures will be accessible through mutation, and therefore goes beyond a correlation between the thermodynamic and mutational stability of the ground state. Hence we introduce the new terminology "congruence."

Bornberg-Bauer and Chan ('99) report that neutral networks in protein models "center around a single prototype sequence of maximum mutational stability". This is not true in RNA. Our evolutionary simulations and numerical studies by Matt Bell (personal communication, August, '99) indicate that the RNA sequences with highest neutrality constitute extended and connected subnetworks of neutral networks. Furthermore, gradient walks to optimize the Boltzmann probability of the minimum free energy structure (Fig. 13) reach maximum neutrality after a dozen steps. Following this, the walks continue to improve the thermodynamic stability of the ground state structure while maintaining constant neutrality.

4. ANALYTIC MODELS OF NEUTRAL CONFINEMENT

4.1. Model assumptions

In this section, we discuss three simplified models of plastic RNA evolution. In each, we characterize the dynamics of a plastic population in the final stage of evolution, and derive equilibrium conditions in terms of neutrality and fitness. The two assumptions that make these approaches analytically tractable are (1) perfect plastogenetic congruence and (2) containment within a single neutral network in which all sequences share the same minimum free energy structure. By perfect plastogenetic congruence, we mean that structures accessible to a sequence through plasticity are exactly those found as minimum free energy structures of sequences accessible through single mutations. Although simulated populations typically consist of sequences from several equally fit neutral networks, our containment assumption simplifies the population to sequences that share a single minimum free energy structure, while their suboptimal structures may vary.

Recall that the plastic genotype–phenotype map gives rise to a much more rugged fitness landscape than the simple map. Neutral networks of sequences with the same minimum free energy structure no longer share the same fitness. Rather, the fitness of a sequence depends directly on other structures in its plastic repertoire. Since plastogenetic congruence ties these structures to the minimum free energy structures of mutational neighbors, our simplifying assumption (1) means that the fitness of a sequence becomes a function of the phenotypes of its genetic neighbors. By extracting this brand of ruggedness in a simpler framework, we demonstrate that the extreme genetic isolation in RNA should come as no surprise.

Furthermore, these models provide insight into the role of mutation. We demonstrate that the genetic isolation may be such that no amount of mutation can facilitate an escape. The minimum amount of mutation necessary to produce novel phenotypes is so great that the resulting mutants have regressed completely from the target.

4.2. Frequency distribution within a neutral network

Consider a population of RNA sequences with phenotypic plasticity. A sequence can assume a range of structures, all within an energetic neighborhood of its minimum free energy structure. The fitness of a sequence is determined by a weighted

average of the distances between its low energy configurations and a target structure. The assumption that plasticity gives a perfect picture of evolutionarily adjacent structures allows us to construe the fitness of a sequence entirely in terms of the fitness of its mutational neighbors:

$$w(S_i) = \sum_{S_j \in \mathcal{N}(S_i)} f(d(\alpha_0^j, \tau)) \cdot \frac{\delta_\mu(S_j, S_i)}{\sum_{S_j \in \mathcal{N}(S_i)} \delta_\mu(S_j, S_i)} \quad (7)$$

where S_i is a sequence, α_0^j is the minimum free energy structure of S_j , $\mathcal{N}(S_i)$ is the mutational neighborhood of S_i , $d()$ determines structural distance between two shapes, $f()$ is a monotonically decreasing function of structural distance $d()$, and $\delta_\mu()$ is a monotonically decreasing function of mutational distance. This definition is analogous to eq. (3) in section 2. Because of perfect plastogenetic congruence, the structures that determine the fitness of S_i are the minimum free energy structures in its mutational neighborhood. The second factor in eq. (7) is analogous to the Boltzmann probability in that it weighs a neighboring structure with the likelihood of reaching its sequence by mutation.

Suppose a population is concentrated on a neutral network G of relatively high fitness. That is, all sequences in G have the same minimum free energy structure, and most one-error mutants of these sequences that lie outside of G have relatively much lower fitness. Let $|G|$ be the number of sequences in G .

We use an approach proposed by van Nimwegen et al. ('99a) in which G is viewed as a graph. Each sequence corresponds to a node, and two nodes are connected by an edge when the sequences they represent differ by exactly one mutation. The degree d_i of a node is the number of edges that connect it to another node in G ; in other words, it is the number of single mutations that preserve the minimum free energy structure. A node S_i will therefore have d_i one-error mutants in G , and $3n - d_i$ one-error mutants that lie outside of G , where n is the length of sequences. Assuming perfect plastogenetic congruence, we can approximate the fitness of an individual node S_i by the average fitness of its one mutant neighbors. We assume fitness off of G is relatively negligible, and therefore ignore the fitness contributions of one-error mutants not in G . Therefore the fitness of S_i is $fd_i + 0(3n - d_i) = fd_i$ where f is the selective value of the minimum free energy structure shared by all sequences in G .

We can express the per generation change in the frequency distribution of sequences in G with a system of $|G|$ equations. For any $S \in G$,

$$P'_s = \frac{f}{\bar{w}} (P_s d_s (1 - \mu) + \frac{\mu}{3n} \sum_{t \in \mathcal{N}(s)} P_t d_t), \quad (8)$$

where P_s gives the frequency of sequence S , μ gives the per sequence per generation mutation rate, n is the length of the sequences, and \bar{w} is the average fitness of the population. Note that this formulation ignores the possibility of mutations onto G from sequences outside of G .

We now translate this system of equations into a transition matrix \mathbf{M} such that $\vec{P}' = \mathbf{M} \vec{P}$ where \vec{P} is the frequency distribution vector. Let \mathbf{I} denote the identity matrix; \mathbf{A} denote the adjacency matrix where $A_{ij} = 1$ if S_i and S_j are one-error mutants of each other and $A_{ij} = 0$ otherwise; and \mathbf{D} denote the diagonal matrix of degrees with $D_{ii} = d_i$ for all i and $D_{ij} = 0$ for all $i \neq j$. Then we derive

$$\mathbf{M} = ((1 - \mu)\mathbf{I} + \frac{\mu}{3n} \mathbf{A})\mathbf{D}. \quad (9)$$

Assume there exists a unique node c such that $d_c > d_i$ for all $i \neq c$. Without loss of generality, we let $c = 0$. Then S_0 is the most neutral sequence in G with degree d_0 . When $\mu = 0$, that is, there is no mutation, then $\mathbf{M} = \mathbf{D}$ which has a leading eigenvalue of d_0 and an associated eigenvector $[1, 0, 0, \dots, 0]$. In the absence of mutation then, the equilibrium population is made up entirely of S_0 .

In general, the equilibrium distribution is the solution \vec{P} to the eigenvalue equation $\mathbf{M}\vec{P} = \lambda\vec{P}$ where λ is the leading eigenvalue of \mathbf{M} . The resulting distribution of sequences will reflect a mutation-selection balance. Because we can break our transition matrix into a diagonalizable matrix and a small remainder term,

$$\mathbf{M} = \mathbf{D} + \mu \left(\frac{1}{3n} \mathbf{A} - \mathbf{D} \right) = [M_0] + \mu [M_1], \quad (10)$$

we can use perturbation theory to get a first-order approximation of the equilibrium distribution for small μ . We find that

$$\hat{P}_k = \begin{cases} \frac{1}{1 + \mu} \left(1 + \mu \left(1 - \frac{d_0}{3n} \sum_{k=1}^{|G|} \frac{A_{0k}}{d_0 - d_k} \right) \right) & \text{if } k = 0, \\ \frac{\mu}{1 + \mu} \left(\frac{d_0}{3n(d_0 - d_k)} \right) & \text{if } A_{0k} = 1, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$




Fig. 18. Network distributions. The equilibrium distributions for a non-plastic population (dotted) and a plastic population (solid). In network **a** (left ordinate), node 1 has much higher neutrality than all other nodes: $d_1 = 19$ and $d_n = 1$ for

$n = 2-20$. In network **b** (right ordinate), nodes 1–10 form a high-neutrality subnetwork: $d_n = 10$ for $n = 1-10$ and $d_n = 1$ for $n = 11-20$.

Recall that A_{ij} indicates whether sequences S_i and S_j differ by a single mutation. According to this approximation, the equilibrium distribution is made up of the most connected sequence S_0 and its one-error mutants. As μ decreases, the proportion of non- S_0 sequence shrinks. Because the denominator $3n(d_0 - d_k)$ depends on d_k , the frequency of a non- S_0 sequence is an increasing function of its degree. Even sequences with degree close to d_0 will appear at low frequency. For example, consider an S_k where $d_k = d_0 - 1$ and the maximal $\mu = 0.1$. Then $\hat{P}_k = 0.3 d_0/n$, which is very small for large n .

For two simple networks, we compare the equilibrium distribution of neutrality predicted by our plastic model to that of the original non-plastic model (van Nimwegen et al., '99a). Both networks contain 20 nodes (sequences). In the first, node 0 is connected to all 19 other nodes, and all other nodes are only connected to 0. In the second, nodes 0–10 are connected to all other nodes while nodes 11–20 are only connected to nodes 0–10. Figure 18 graphs the equilibrium distributions for non-plastic (dotted) and plastic (solid) populations. The plastic populations are much more concentrated on the nodes with highest degree.

4.3. Three-tiered model: exploration and error thresholds

Premised on the previous model, we turn to the role of mutation. Can mutation enable the population to avoid the evolutionary dead-end in the first place? Intuition suggests that an increased mutation rate might enable populations to search beyond their immediate genetic vicinity, into regions where novel phenotypes exist. We show here, however, that for certain parameter values there is no mutation rate that provides such an escape. Under low mutation rates, populations, like those in our simulations, are confined in an “exploration catastrophe,” yet under high mutation rates, populations are in an “error catastrophe” where they have slipped away completely from the target.

Again consider a population that has reached a relatively high fitness neutral network G . We break down the population into three distinct classes, two within the neutral network and one representing the rest of sequence space. For any sequence $S_i \in G$, the one-error mutants of S_i that are also in G are called its neutral neighbors, and the number of such neighbors is called d_i , the degree of S_i . Assume that the neutral network contains a single sequence

S_0 with maximal connectivity d_0 within G , and the degree of all other sequences in G is d_1 . That is, for any $S_i \in G$ where $S_i \neq S_0$, $d_i = d_1 < d_0$.

We define three classes of sequences: C_0 is just the sequence S_0 , C_1 consists of the sequences in $G - \{S_0\}$, and C_2 are all sequences not contained in G . Changes in the frequency distribution among these classes results from a combination of mutation and natural selection.

Let μ be the probability of mutation per sequence per generation. Sequences in C_0 mutate to sequences in C_1 at a rate $\mu d_0/3n$, where $3n$ is the number of possible one-error mutants of a sequence. Mutations take sequences in C_0 off the network, into C_2 at a rate $\mu(1 - \mu d_0/3n)$. Likewise the mutation rates from C_1 to itself and to C_2 are $\mu d_1/3n$ and $\mu(1 - d_1/3n)$, respectively. We ignore back mutation of sequences in C_1 to C_0 and of sequences in C_2 to either C_0 or C_1 . We discuss the implications of this assumption below.

The sequences in this model are again endowed with perfect plastogenetic congruence as described in the previous section. Every sequence has a selective value which is a measure of the similarity between the sequence's minimum free energy structure and a pre-determined target structure. Sequences within G have a selective value of 1 while sequences off G have a relative selective value of $f < 1$. The relative fitness of a sequence in class C_i is w_i , an average of the selective values of its one-error mutants:

$$w_0 = \frac{d_0}{3n} + \left(1 - \frac{d_0}{3n}\right) f, \quad (12)$$

$$w_1 = \frac{d_1}{3n} + \left(1 - \frac{d_1}{3n}\right) f, \quad (13)$$

$$w_2 = f. \quad (14)$$

Recall that $d_i/3n$ ($i = 0, 1$) is, for any sequence in C_i , the fraction of its one-error mutants that also lie within G . We construct a transition matrix \mathbf{T} that incorporates the effects of mutation and selection. If $\vec{P} = (P_1, P_2, P_3)$ describes the occupancy of the three classes at some time, then $\vec{P}' = \mathbf{T}\vec{P}$ gives the frequency distribution in the next generation where \mathbf{T} is as follows:

$$\mathbf{T} = \begin{pmatrix} w_0(1-\mu) & 0 & 0 \\ w_0\mu\frac{d_0}{3n} & w_1(1-\mu) + w_1\mu\frac{d_1}{3n} & 0 \\ w_0\mu(1-\frac{d_0}{3n}) & w_1\mu(1-\frac{d_1}{3n}) & w_2 \end{pmatrix}. \quad (15)$$

The leading eigenvector of \mathbf{T} provides the population equilibrium distribution. The eigensystem of \mathbf{T} is given by

$$\begin{pmatrix} \lambda_0 \\ \lambda_1 \\ \lambda_2 \end{pmatrix} = \begin{pmatrix} w_0(1-\mu) \\ w_1(1-\mu(1-\frac{d_1}{3n})) \\ w_2 \end{pmatrix} \quad (16)$$

and

$$\vec{P}_0 = \left(\frac{w_0(1-\mu) - w_1(1-\mu(1-\frac{d_1}{3n}))}{w_0\mu\frac{d_0}{3n}}, \mathbf{1}, \frac{w_1(1-\frac{d_0}{3n} - \mu(1-\frac{d_1}{3n})) - w_0(1-\frac{d_0}{3n} - \mu(1-\frac{d_0}{3n}))}{-w_0\frac{d_0}{3n}(1-\mu) + w_2} \right), \quad (17)$$

$$\vec{P}_1 = \left(0, \frac{w_1(1-\mu(1-\frac{d_1}{3n})) - w_2}{w_1\mu(1-\frac{d_1}{3n})}, \mathbf{1} \right) \quad (18)$$

$$\vec{P}_2 = (0, 0, \mathbf{1}). \quad (19)$$

The subscripts of the eigenvalues and eigenvectors refer to the class which dominates the distribution, and not to their magnitudes. As parameter values vary, so does the leading eigenvalue $\lambda = \max(\lambda_0, \lambda_1, \lambda_2)$.

We seek parameter ranges that allow the population to explore phenotype space. A population concentrated in C_0 will have mostly neutral mutants, and therefore will be unlikely to find a higher fitness phenotype through mutation. We call this an exploration catastrophe. A population lost in C_2 has regressed from the higher fitness network. This is the error catastrophe. C_1 on the other hand is a high fitness platform from which a population can explore phenotype space for better options. For these reasons, we say that a population with a high concentration of C_1 is below the error threshold which moves the population off of G , and above the exploration threshold which contains the population in a highly inward-looking subset of G .

In Fig. 19, we display equilibrium distributions for various parameter ranges. In every case, we set the length of the sequences to $n = 100$. Each right-hand graph shows the concentration of C_0 (labeled "0"), C_1 (labeled "1"), and C_2 (labeled "2"). Populations dominated by C_2 are above the error threshold while populations dominated by C_0 are below the exploration threshold. On the left we graph the eigenvalues. The topmost plane at each point is the leading eigenvalue. Its associated eigenvector is that which determined the frequen-

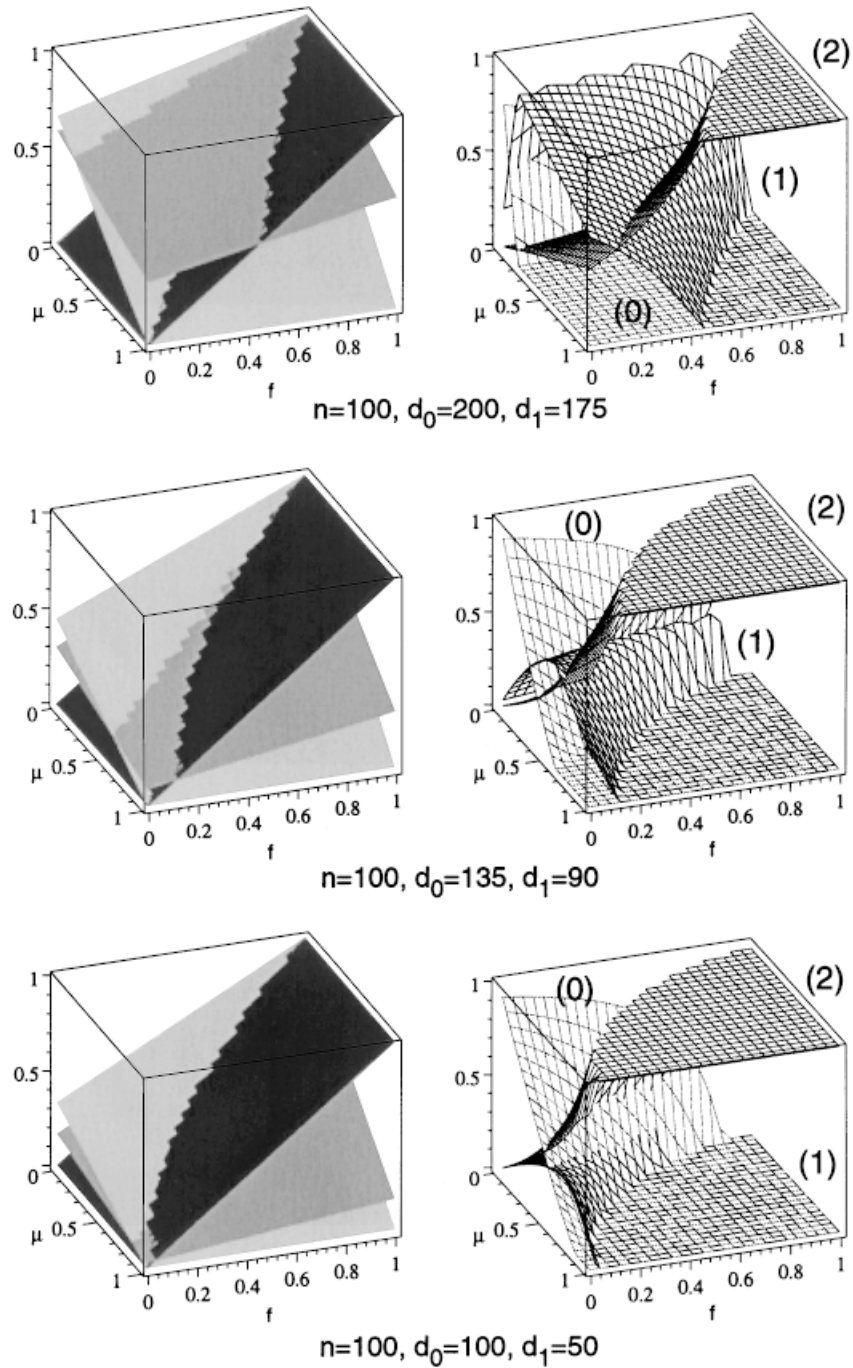


Fig. 19. Eigenvalues (left graphs) and equilibrium distributions (right graphs). Given values for μ and f , we graph λ_0 , λ_1 , and λ_2 in light, medium, and dark gray, respectively, and

equilibrium concentrations for C_0 , C_1 , and C_2 labeled as (0), (1), and (2), respectively.

cies in the opposing graph. Light gray, medium gray, and dark gray represent λ_0 , λ_1 , and λ_2 , respectively.

The center distribution assumes neutralities of $d_0/3n = 0.45$ for C_0 and $d_1/3n = 0.3$ for C_1 . These

are the values we obtain from actual simulation. The equilibrium distribution suggests that for the low values of μ we use in simulations, the population should be trapped on C_0 in an exploration catastrophe. This is consistent with the neutral

confinement we find for low-plasticity sequences. We discuss these distributions further in the next section.

4.4. Neutrality in the subnetwork C_0

In simulation, RNA populations evolve through neutral networks that are much more complex than the idealized C_0 and C_1 of the three-tiered model. Populations move among highly interconnected subnetworks of multiple structurally similar neutral networks. In this extension of the three-tiered model, we still assume the population converges on a single neutral network G , but attempt a more realistic conception of the structure of G .

Suppose now that C_0 is enlarged to a subnetwork of the neutral network G in which all $|C_0|$ sequences have degree d_0 . As above, the remaining sequences in G occupy C_1 and share degree d_1 . We define two new parameters χ_0 and χ_1 , which are the fractions of neutral mutations that remain within C_0 and C_1 , respectively. Thus $1 - \chi_0$ is the fraction of neutral mutations from C_0 to C_1 , and $1 - \chi_1$ is the fraction from C_1 to C_0 .

Mutation operates according to the following table of mutation rates (Table 2). The entry in the C_i row and C_j column is the probability that an individual in C_i mutates into C_j in a given generation. We will use these rates in the transition matrix.

In this case, we also consider mutations from C_1 into C_0 , but again we ignore mutations from C_2 into G .

As in the original formulation, fitness is the weighted average of the selective values of the minimum free energy structures of all one-error mutants of a sequence. For any sequence in C_0 or C_1 the total number of one-error mutants that are also in G is still d_0 or d_1 , respectively. Therefore the extension of C_0 to a sub-network of G does not alter the fitnesses of the three classes.

Again we construct a transition matrix that de-

scribes flow between classes. In the following matrix μ_{ij} , $0 \leq i, j \leq 2$, is the mutation rate from C_i to C_j as given in Table 2:

$$\mathbf{T}' = \begin{pmatrix} w_0(1-\mu) + w_0\mu_{00} & w_1\mu_{10} & 0 \\ w_0\mu_{01} & w_1(1-\mu) + w_1\mu_{11} & 0 \\ w_0\mu_{02} & w_1\mu_{12} & w_2 \end{pmatrix}. \quad (20)$$

Though it still only involves a quadratic and a linear term, the eigensystem of \mathbf{T}' looks more complicated than that of \mathbf{T} because \mathbf{T}' is not triangular, and we have added the new parameters χ_0 and χ_1 . In the following analysis, we assume $\chi_0 = \chi_1 = \chi$, and explore the effects of χ and μ on the predicted equilibrium distributions.

Figure 20 depicts equilibrium distributions for $f = 0.9, 0.5, 0.1$ from left to right. As χ increases, the equilibrium concentration of C_0 increases slightly and the frequency of C_1 decreases, making exploration even more difficult. For very larger χ , the frequency of C_1 remains negligible for all μ . See for example Fig. 21.

The lower eigenvalue graphs reveal that under the extended model, populations are always in either the λ_0 phase or the λ_2 phase. In the previous model, we found three distinct regimes, each dominated by a unique eigenvector. This discrepancy arises from the role of mutation. In the first model recall that there is no back mutation from C_2 into C_0 and C_1 , and no back mutation from C_1 to C_0 . In the second model, we add mutation from C_1 to C_0 . A model that includes mutation in all directions among three classes would yield a single maximum eigenvector over all of parameter space. In the first model, we use the transitions between leading eigenvalues to mark the phase transitions between regimes dominated by different subsets of the population. In the current model and the hypothetical all-mutations model, however, transitions between such phases in the frequency distribution are difficult to identify. One innovative approach to this problem considers finite-population dynamics in a simplex (van Nimwegen et al., '99b). Here a qualitative perspective suggests that there are parameter values for which populations go directly from an exploration catastrophe, confined to C_0 , to an exploration catastrophe, diffusing in C_2 .

4.5. Exploration and error catastrophes

In both versions of the model we find parameter ranges for which the population is confined

TABLE 2. Mutation rates

Origin	Destination		
	C_0	C_1	C_2
C_0	$\chi_0\mu\frac{d_0}{3n}$	$(1-\chi_0)\mu\frac{d_0}{3n}$	$\mu(1-\frac{d_0}{3n})$
C_1	$(1-\chi_1)\mu\frac{d_1}{3n}$	$\chi_1\mu\frac{d_1}{3n}$	$\mu(1-\frac{d_1}{3n})$
C_2	0	0	μ

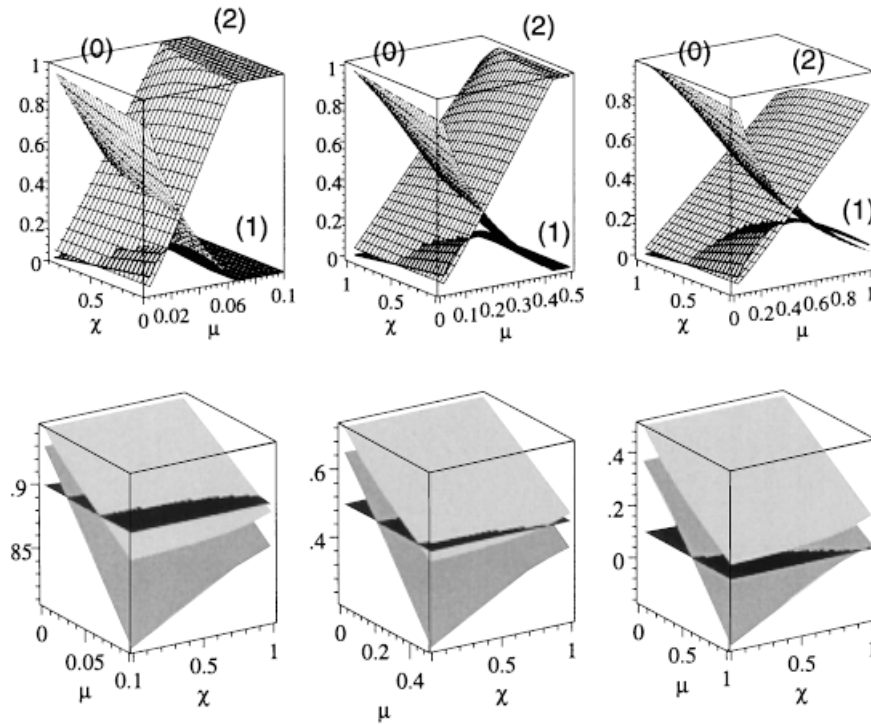


Fig. 20. Extended model equilibrium distributions and eigenvalues. Given values for μ and χ , these display equilibrium concentrations for C_0 , C_1 , and C_2 labeled as (0), (1), and (2), respectively. The bottom row gives eigenvalues λ_0 in light

gray, λ_1 in medium gray, and λ_2 in dark gray. All graphs assume $n = 100$, $d_0 = 135$, and $d_1 = 95$. From left to right, $f = 0.9, 0.5, 0.1$. Note that the range for μ varies across graphs.

to C_0 and therefore is unlikely to find phenotypic novelty through mutation (exploration catastrophe). For given values of f , n , and χ we can identify a threshold μ below which the population is in an exploration catastrophe. Generally, the exploration threshold decreases as f increases and is




Fig. 21. Equilibrium distribution for $\chi = 0.9$. For low μ , the population is confined to C_0 . Around $\mu = 0.25$, the population moves off C_0 into C_2 , with virtually no occupancy of C_1 .

a slightly monotonically increasing function of χ . This implies that as the size and interconnectivity of C_0 grows, so does the likelihood of reaching an evolutionary dead-end.

There also exist parameter ranges for which the population is lost in C_2 . For specified f , n , r , and χ we can find a threshold μ above which the population reaches such an error catastrophe. Intuitively, as f —the relative fitness of C_2 —increases, the error threshold decreases.

These transitions are formally identical to the well-known genotypic error threshold, denoting the mutation rate at which a genotype with optimal phenotype (“master genotype”) is lost (Eigen, '71), and the phenotypic error threshold, denoting the mutation rate at which the optimal phenotype is no longer maintained in the population (Huynen et al., '96; Reidys et al., '98).

In the simple version of our model, the exploration threshold corresponds to the genotypic error threshold in which there is a single master sequence with superiority approximately $w_0/[d_0w_1/3n + (1 - d_0/3n)w_2] = [d_0(1 - f) + 3nf]/[d_1d_0(1 - f)/3n + 3nf]$. The genotypic error threshold divides the low mutation rates at which the master sequence is preserved at equilibrium from the

high mutation rates at which the master class is catastrophically lost. The exploration threshold similarly divides the low mutation rates at which the population is confined to C_0 from the high mutation rates at which the population reaches C_1 .

What we call the error catastrophe in our simple model is equivalent to a phenotypic error threshold where the mean fraction of neutral neighbors is approximately $v \sim d_1/3n$. In our extended model both the exploration catastrophe and the error catastrophe correspond to a phenotypic error threshold because C_0 becomes a neutral network rather than a single sequence. Increasing the size of C_0 , i.e., increasing the neutrality of C_0 , postpones the exploration threshold. A larger C_0 entails extended ranges of mutation rates at which the population will be in an exploration catastrophe. Consequently the regions of parameter space in which the population goes directly from an exploration catastrophe into an error catastrophe also increase.

5. MODULARITY

Modularity is a hallmark of biological organization and an important source of evolutionary innovation (Bonner, '88; Wagner and Altenberg, '96; Hartwell et al., '99). Once it exists, modularity constitutes an obvious advantage by enabling the recombination of stable subunits into novel phenotypes. Yet, the origin of modules remains a problem for evolutionary biology, even in the case of the most basic protein or RNA domains (Westhof et al., '96). Here we offer a possible origin for such organization.

In section 2, we demonstrated that plasticity is rapidly reduced by natural selection under a biophysically motivated fitness function that weighs the selective value of each shape in the plastic repertoire of a sequence by its Boltzmann probability. The reduction of plasticity has, in addition to genetic canalization, a further side effect which is best characterized as modularity.

The point we are making in this section is, in essence, yet another characterization of plasto-genetic congruence in RNA: structural units that appear autonomous from an environmental and developmental perspective appear at the same time autonomous from a genetic perspective. Here the environment refers to temperature, and development refers to the kinetic process by which an RNA sequence folds from an open chain into its minimum free energy structure. We show that the evolutionary reduction of plasticity crystallizes RNA structures into environmentally, developmentally, and genetically autonomous units.

5.1. Norms of reaction: melting behavior

We view plasticity as a stochastic choice among alternative structural states of a biopolymer in contact with a heat bath at constant temperature. Biologists often characterize the plasticity of a genotype using a *norm of reaction*. This is typically a map from some parameter in the environment to a phenotype. The RNA plastic repertoires deviate from this standard framework in that the relevant environmental input—Brownian motion in a heat bath—is not easily scaled on an axis. In this section we use a more conventional norm of reaction: a map from temperature to minimum free energy structure. The suite of minimum free energy structures as they change with temperature is known in biophysics as “melting” profile. Figures 22 and 23 compare the (computed) melting profiles of the sequences from Fig. 15 at three levels.

First we compute the melting series, that is, the suite of minimum free energy structures in the temperature range from 0 to 100°C. Second we calculate the heat capacity (at constant pressure), labeled H in Figs. 22 and 23, from the Gibbs free energy, G , of the ensemble of structural states by means of the partition function Z (McCaskill, '90):

$$H = -T \frac{\partial^2 G}{\partial T^2} \text{ with } G = -RT \log Z, \quad (21)$$

where R is the universal gas constant. H can be measured empirically by differential scanning calorimetry (DSC). DSC is widely used to determine the thermophysical properties of materials. As a sample is heated over a range of temperature, the material starts to undergo a phase change that releases or absorbs heat. The calorimeter measures the heat flow (enthalpy change) into or out of the sample undergoing the phase change, thereby providing data from which the heat capacity can be quantitatively recovered. Our third perspective on structural transitions are the Boltzmann probabilities as a function of temperature for each minimum free energy structure appearing in the melting series. These show that the peaks in H correspond to major structural phase transitions occurring at the temperature at which the Boltzmann probabilities of the outgoing and incoming structures intersect.

In Fig. 22, we compare the melting profile for the three sequences on the left of Fig. 15. Recall that these sequences share a common minimum free energy structure at 37°C. The melting behav-

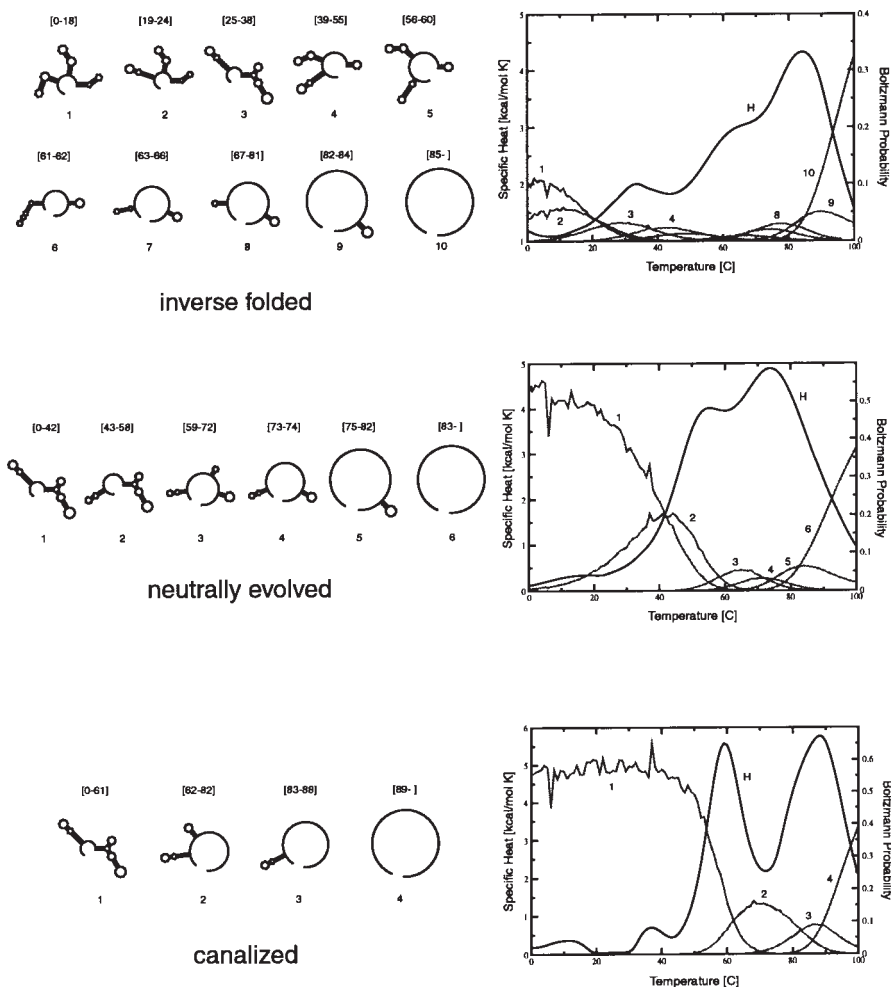


Fig. 22. Melting behavior I. Left of each graph are the minimum free energy structures in the temperature range 0–100°C for three sequences obtained by inverse folding, neutral evolution, and canalization on the same structure (Fig.

15, left). The graphs trace the temperature dependence of the specific heat (H) and the Boltzmann probabilities of the individual structures in the melting series.

ior of the inverse folded sequence is markedly more disorderly than that of the neutrally evolved and canalized sequences. Its melting series consists of as many as 10 shapes. The relevant observation, however, is the absence of structural features that rearrange *locally*, particularly at low temperatures (when there is more structure). In other words, temperature variations induce global refoldings of the shape. Furthermore, each structure in the melting series remains stable only over a small temperature range. As a consequence, the heat capacity consists of several small and closely occurring humps.

The canalized sequence, in contrast, has a highly localized melting behavior. We easily identify structural features that melt individually at distinct temperatures *without affecting the integ-*

ity of other parts of the structure. Often, but not always, such structural units coincide with components in the sense of Fig. 1. The structure of the canalized sequence is stable up to 61°C, when the large T-like feature at the 5' end disappears almost entirely in a single step. Remarkably, the other feature at the 3' end is not affected, despite a new open sequence segment that is now available for interaction. The 3' feature melts at its own transition temperature of 89°C in a single step. The major transitions from #1 to #2 and from #3 to #4 are well separated and marked by two sharp peaks in the heat capacity.

The neutrally evolved sequence occupies a middle ground. It has fewer transitions and more highly preserved structural similarity across transitions than the inverse folded sequence, but to a

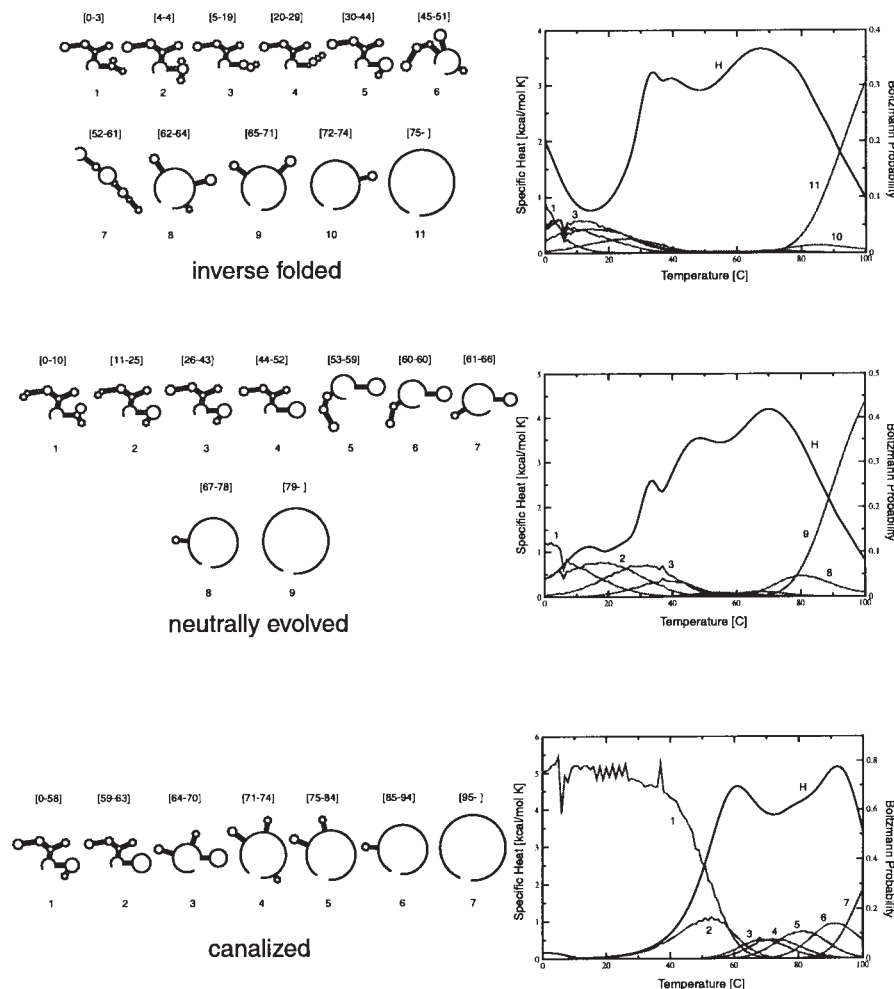


Fig. 23. Melting behavior II. See caption to Fig. 22. The 37°C ground state structure for these sequences is the one pictured on the right of Fig. 15.

lesser extent than the canalized sequence. It departs significantly from the canalized sequence in its lack of independently melting features. For example, the transition from #1 to #2, involving the hairpin structure at the 3' end, neither melts this feature completely nor does it preserve some of its subfeatures. The same holds for the structure portion at the 5' end.

This comparative analysis of melting behaviors suggests one facet of “modularity”: the thermophysical independence of a structural trait from other traits over a wide temperature range. Note that this is a quite different notion of “unit” than what is obtained from parsing a shape into structural units based on morphology alone. In fact, all three sequences considered have the same morphology at 37°C, but only the canalized sequence is modular in the sense of thermophysically autonomous units.

Figure 23 illustrates the same points. Note, however, that the modules evidenced in Fig. 22 are larger structural assemblies than in the case of Fig. 23.

5.2 Kinetics: modularity and funnels

A recent stochastic model of kinetic RNA folding at elementary step resolution permits the study of folding pathways (Flamm et al., '99). The folding kinetics of our three sequences provides a further perspective on modularity. The folding pathways of RNA play a role analogous to developmental pathways of organisms.

The tree graphs in Fig. 24 represent different organizations of the energy landscape on which the folding process occurs. For each of our three sequences on the left of Fig. 15 the portion of the energy landscape shown comprises the 50 lowest local free energy minima (courtesy Christoph

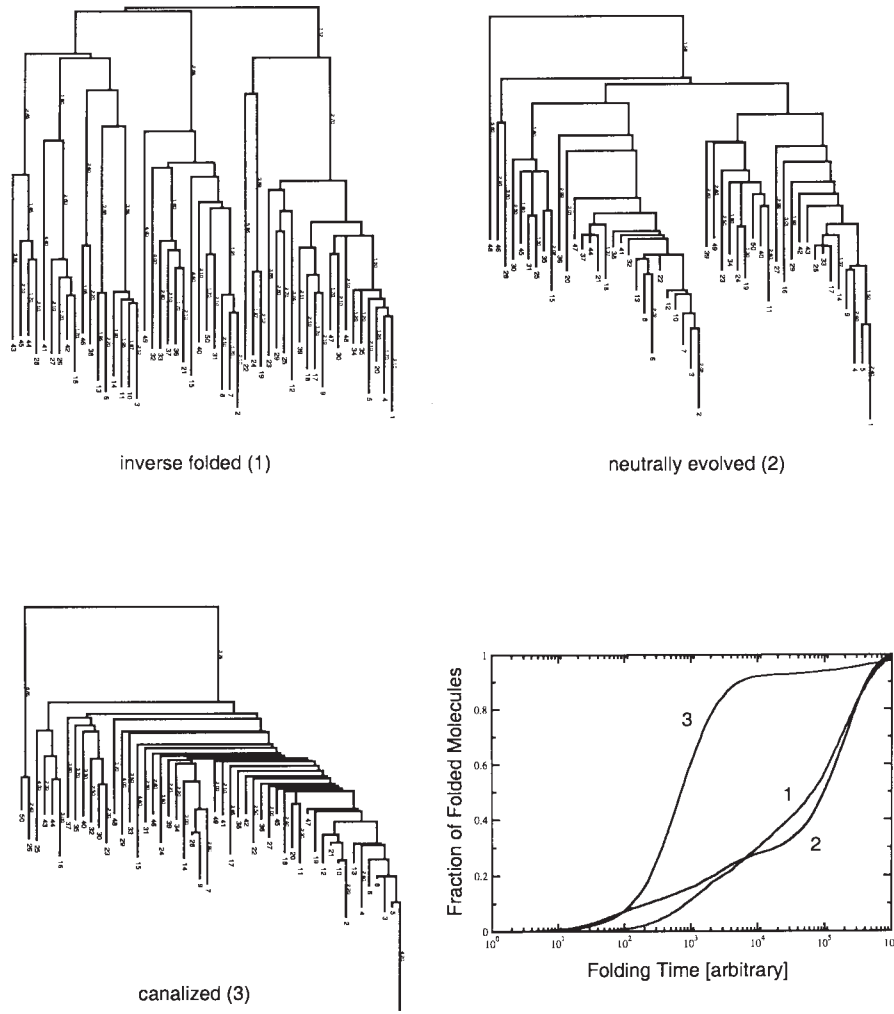


Fig. 24. Folding kinetics and energy landscapes. The graph on the lower right shows the folding time distributions for the three classes of sequences shown on the left of Fig. 15, based on a stochastic model of kinetic folding whose elementary moves consist in the making, breaking and shifting of a single base pair (Flamm et al., '99). Images (1), (2),

and (3) are the inverse folded, neutrally evolved, and canalized sequences, respectively. The trees depict the energy landscape associated with each sequence in terms of the hierarchical organization of barriers separating individual states and their basins. The distance from the root (top) of the tree represents the free energy.

Flamm, University of Vienna). A local minimum corresponds to a leaf, and leaves are grouped into basins which are further linked to one another in a hierarchical fashion. The heights of internal nodes represent energy barriers connecting two local minima or their basins (Flamm et al., '99). In other words, the height of the lowest internal node connecting two leaves represents the energetic requirement of refolding from one structure into the other. The lower right graph of Fig. 24 shows the distribution of folding times (first passage times from the unfolded sequence to the native structure) for these sequences.

The differences are striking. Not only does the

canalized sequence fold much more rapidly than the others, but it folds predominantly along one well-defined pathway and has, therefore, one dominant time scale. Its energy surface resembles a funnel (Bryngelson and Wolynes, '87; Dill and Chan, '97), suggesting that individual structural units fold independently from one another. Sequence segments of one unit are unlikely to crossfold with segments of other units, which would cause traps delaying the formation of the native structure. (There is a folding trap [not shown], visited with low probability, that accounts for the tail of the distribution. It is due to early diffusion among several high-energy shapes until

the molecule drops into the funnel and folds). In contrast, the inverse folded sequence has an energy landscape without much structure, as well as high barriers separating individual states. The sequence can misfold in many ways, giving rise to a broad distribution of folding times. The neutrally evolved sequence exhibits two funnels, one of which leads to a misfold.

5.3. Context insensitivity

We used the melting behavior and the organization of the energy landscape to identify modularity with respect to temperature change and folding dynamics. We next turn to modularity as identified through genetic perturbation.

Point mutations cause only local disruption (if any) to the minimum free energy structure of a low-plasticity sequence. A thermodynamically well-defined stacking region is unlikely to unfold completely when a point mutation knocks out a single base pair, while a marginally stable stacking region is likely to unwind, leading to a global rearrangement of the structure. Modules buffer against extensive rearrangements. This aspect of

modularity is similar to low pleiotropy in an organismal context (Wagner and Altenberg, '96).

Perhaps the most defining property of modules is the maintenance of their structural integrity across changing genetic contexts. We slice from each sequence the segment s_τ that folds into a particular structural unit τ . We then paste to the left and right of s_τ random segments half its size and ask whether s_τ still folds into τ despite the new genetic embedding. The chart in Fig. 25 shows the fraction of 1,000 such foldings that maintain τ . Indeed, canalized sequences have structural components that are much more context independent than inverse folded sequences. In a few cases, neutrally evolved sequences withstand contextual modification as well as the canalized ones.

To summarize, plastogenetic congruence states that plasticity (the environmental variance of phenotype) mirrors variability (the mutational sensitivity of phenotype). This section on modularity, in essence, further elaborates this theme. Thermodynamic and kinetic autonomy of units, as manifest in the norms of reaction to temperature and the organization of the energy landscape, correlates to the autonomy of those same units with respect to changing genetic contexts. A computational analysis of naturally occurring sequences suggests that functionally important structures have heightened context insensitivity (Wagner and Stadler, '99). Recall, however, that in our simulations sequences were never selected for modularity, only for reduced plasticity. Direct selection pressures for modularity may exist, but this analysis demonstrates that the emergence of modularity does not require them. Modularity arises, like genetic canalization, as a byproduct of environmental canalization.

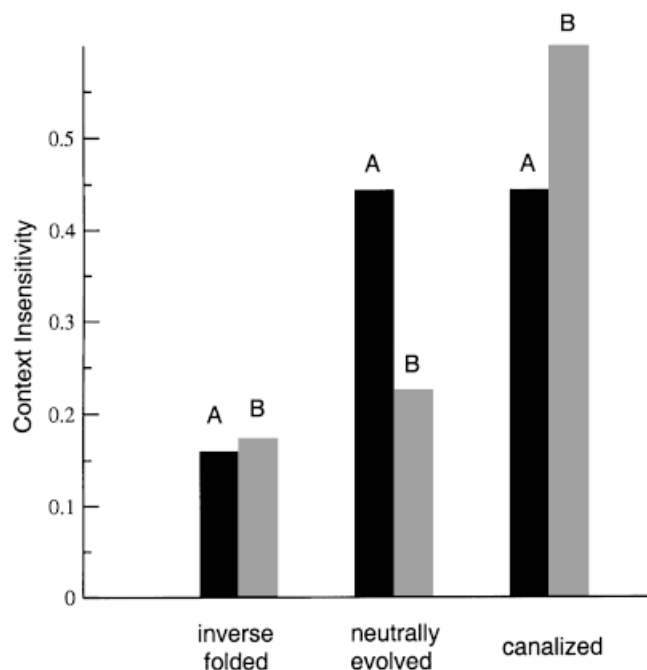


Fig. 25. Context insensitivity of modules. The sequence segments underlying the 5' structural component (A) and the 3' component (B) of the shapes on the left and right of Fig. 15, respectively, are embedded in random sequence contexts. The chart shows the frequency with which the segment retained its original structure if it originated in an inverse folded, neutrally evolved, and canalized sequence.

6. DISCUSSION

Biological evolution changes not only the frequencies of extant phenotypes, but the phenotypes themselves. A population genetic analysis of the fate of innovations under natural selection provides only a partial story that must be integrated with a theory of phenotypic innovations (Buss, '87).

We turn our attention to a simple but nontrivial evolvable object: an RNA molecule. RNA provides both a theoretically and empirically well-characterized high-dimensional relation between genotype (sequence) and phenotype (structure). The main algorithms for RNA folding were developed twenty years ago (Nussinov et al., '78; Waterman and Smith, '78; Zuker and Stiegler, '81), not with the present questions in mind but rather as a tool

to assist experimentalists. This divergence of applications suggests that inscription errors (that is, the construction of models which predetermine the desired output) are less likely here than for genotype–phenotype models constructed with particular definitions of epistasis, plasticity, or modularity in mind.

Recent advances in RNA folding (McCaskill, '90; Flamm et al., '99; Wuchty et al., '99) enable us to introduce and analyze a form of environment–gene interaction which we call plasticity. The result is a powerful model system in which concepts like plasticity, evolvability, epistasis, and modularity not only can be precisely defined and statistically measured, but reveal simultaneous and profoundly non-independent effects of natural selection. Although these concepts were introduced in the study of organismal evolution, we demonstrate that they apply to the molecular domain as well and are optimistic that lessons from RNA may, in turn, provide robust insight.

The secondary structure into which an RNA sequence folds is determined not by the primary sequence alone but also by environmental inputs such as temperature and the presence of other potentially interacting molecules. As a surrogate for such environmental factors, we map a sequence to a repertoire of its thermodynamically most stable structures. We assume that the Boltzmann coefficient—a variable reflecting the thermodynamic stability of a structure relative to all other structures within the configuration space of a sequence—is proportional to the time that sequence would spend in the given structure under a heterogeneous environment.

The most striking outcome of our simulations is the dramatic loss of variability that accompanies the evolutionary reduction of diversity in the plastic repertoires. Recall that variability is the potential of a population of sequences to innovate phenotypically. We gain a deeper understanding of the loss of variability through two lines of inquiry. First we construct the causal bridge from the assumptions of our model—point mutation, a plastic genotype–phenotype map, and fitness based on the average structural distance to target over all structures in a plastic repertoire—to the observed evolutionary dead-end. Second, we characterize in as many dimensions as possible the typical genotype and phenotype distribution for a steady state population evolved under the plastic map. These two objectives are highly interrelated. The link between plasticity and variability is shown by a close look at the evolving

distribution of sequences and their shape repertoires.

The loss of variability stems from two simple observations:

1. The more variation in the plastic repertoire, the less time a sequence spends in its best structure. In this way, plasticity is costly and is ultimately reduced by natural selection in constant environments.
2. There is a significant overlap between the shapes in the plastic repertoire of a sequence and the set of minimum free energy structures of genetically proximate sequences, i.e., of sequences that differ from it by one mutation. We call this property plastogenetic congruence.

The first observation is a straightforward consequence of our (biophysically motivated) plastic fitness function. We verify the second, which rests on the intuition that a point mutation can tip the folding landscape of a sequence in favor of any low-energy structure, through several statistical assays. Phenocopies, environmentally triggered traits that correspond to mutant phenotypes, provide evidence for the generality of plastogenetic congruence in nature. The effects of high temperature on moth antenna morphology (Goldschmidt, '41), of ether on *Drosophila melanogaster* thoracic development (Waddington, '42; Gibson and Hogness, '96) and of gold foil on the fibular crest in birds (Müller, '90) among many other environmental perturbations have been shown to mimic known mutants or ancestral morphologies (Stearns, '93; Schlichting and Pigliucci, '98). Because the alignment between environmental variability (plasticity) and genetic variability has met the skepticism of evolutionary geneticists, it is not well-integrated into mainstream evolutionary thinking. Our example of the plastic RNA map illustrates its relevance to evolutionary theory and instantiates the claim that plastogenetic congruence is a fairly ubiquitous property of genotype–phenotype relationships.

Under the conditions studied here, natural selection reduces plasticity. Plastogenetic congruence then implies that a drop in diversity in the plastic repertoire entails a drop in the diversity of minimum free energy structures in the one-error neighborhood (the set of one-error mutants of a given sequence). Natural selection on plasticity indirectly curtails phenotypic novelty accessible by mutation, and hence the potential to evolve.

Through idealized models of a plastic popula-

tion, we formalize this explanation for the decline in evolvability. These models become analytically tractable under the assumption that the plastogenetic congruence is perfect. In other words, the plastic repertoire for any given sequence is exactly the set of minimum free energy structures in the one-error neighborhood. As the fraction of one-error mutants with identical minimum free energy structure (neutrality) increases, the ability to evolve decreases.

The models formally connect mutation rate, the topography of a phenotype space and evolvability. Assuming that the population has reached a neutral network relatively close to the target shape, we identify three phases of equilibrium distributions: the exploration catastrophe, when the population is concentrated in a highly neutral region and so cannot access phenotypic novelty; the error catastrophe, when the population falls off the neutral network into the rest of genotype space with on average much lower fitness; and the ideal phase, when the population remains in regions of the neutral network that have mutational access to the rest of phenotype space. The fate of the population as mutation increases depends on the structure of the neutral network. For some, increasing mutation rate takes the population from the exploration catastrophe through the ideal phase to the error catastrophe. For many neutral networks, however, the exploration threshold exceeds the error threshold. Upon increasing mutation rate, a population goes immediately from an exploration catastrophe to an error catastrophe. Simulation suggests that this is the predicament of our steady state RNA populations under the plastic map.

Plastogenetic congruence is a robust statistical feature of the RNA folding map from sequences to secondary structures with broad population genetic implications. In particular, it instantiates the hypothesis put forward by Wagner et al. ('97) that genetic canalization—buffering against phenotypic effects of mutation—occurs as a byproduct of environmental canalization—the evolution of resilience to environmental perturbation.

In meeting the second objective, a characterization of the phenotypic consequences of natural selection, we discovered a second, equally remarkable byproduct of environmental canalization: modularity. By modules we do not mean a syntactical property of structures (which is trivial in RNA) but rather autonomous components that maintain their structural integrity across a broad range of environmental and genetic contexts and

that lose integrity through sudden and discrete steps without affecting the remaining structure. Modularity manifests itself as a resistance to sustained environmental or genetic perturbation, and the dissolution of modules translates into sharp and well separated phase transitions. In section 5, we compare thermophysical, genetic and kinetic aspects of modularity across sequences that were generated by different processes but share the same minimum free energy structure. Modularity appears thermophysically as distinct melting temperatures of structural components that vanish upon melting, leaving an open chain segment, rather than a different structural arrangement. Modularity appears kinetically as a single primary folding funnel over the configurational landscape of a sequence. There is little probability of misfolding because the folding of a modular component does not interfere with the folding of other components. Modularity appears genetically as cassette-like behavior, by which modular structural components have a markedly increased probability of maintaining their integrity if transplanted into different sequence contexts. This agrees with the principles of RNA architecture discovered through recent crystallizations of catalytic RNAs [e.g., Cech et al. ('94); for an overview see Westhof et al. ('96)].

Modularity is both a manifestation of evolutionary lock-in, and provides the basic tool for escaping it. The evolutionary stability of modules makes them, in conjunction with their context-insensitivity (transposability), natural building blocks for constructing novelty at a higher combinatorial level (Wagner and Altenberg, '96). The shift toward a combinatorics of modular elements to escape evolutionary lock-in appears to be less of a convenient ad hoc innovation in evolutionary process than the only route left, once modules arise as a byproduct of environmental canalization in a constant environment.

Throughout this work, we hold the temperature and target shape constant, and find an intuitive evolutionary loss of plasticity and some surprising corollaries. What happens, though, when the environment is not constant? In particular, what conditions favor the maintenance of plasticity? We hope to analyze several modes of environmental variability: temperature heterogeneity, target fluctuations, and joint folds with co-occurring molecules.

ACKNOWLEDGMENTS

Thanks to our colleagues Ivo Hofacker, Peter Stadler, and Stefan Wuchty for their contributions

to the Vienna RNA package. We are grateful to Leo Buss, Marc Feldman, Christoph Flamm, Peter Godfrey-Smith, Ellen Goldberg, James Griesemer, Martijn Huynen, Erica Jen, Michael Lachmann, Laura Landweber, Mark Newman, Erik van Nimwegen, Peter Schuster, and Andreas Wagner for valuable discussions and comments on the manuscript. Günter Wagner spotted several mistakes in a previous draft. Junhyong Kim's thoughtful comments led to a much improved revision. Christoph Flamm kindly provided the folding kinetics of section 5. Matt Bell, 1999 REU summer intern at the Santa Fe Institute, shared valuable insight and back-up computations. The present work was supported by a Steinmetz Fellowship and a U.S. National Defense Science and Engineering Fellowship to L.A., by NIH Grant GM-28016 to M.W. Feldman, by the Keck Foundation, and by core grants to the Santa Fe Institute from the John D. and Catherine T. MacArthur Foundation, the National Science Foundation, and the U.S. Department of Energy. W.F. and his research program are supported by Michael A. Grantham.

REFERENCES

- Altenberg L. The schema theorem and Price's theorem. In: Whitley D, Vose MD, editors, *Foundations of Genetic Algorithms*, pages 23–49. MIT Press, Cambridge, MA, 1995.
- Ancel LW. 1999a. A quantitative model of the Simpson-Baldwin effect. *J Theor Biol* 196:197–209.
- Ancel LW. 1999b. Undermining the Baldwin expediting effect: does phenotypic plasticity accelerate evolution? *Theoretical Population Biol* (in press).
- Baldwin JM. 1896. A new factor in evolution. *Am Nat* 30:441–451.
- Bartel DP, Szostak JW. 1993. Isolation of new ribozymes from a large pool of random sequences. *Science* 261:1411–1418.
- Beaudry AA, Joyce GFR. 1992. Directed evolution of an RNA enzyme. *Science* 257:635–641.
- Bonner JT. 1988. *The evolution of complexity*. Princeton, NJ: Princeton University Press.
- Bornberg-Bauer E, Chan HS. 1999. Modeling evolutionary landscapes: mutational stability, topology, and superfunnels in sequence space. *Proc Natl Acad Sci USA* 96:10689–10694.
- Bryngelson JD, Wolynes PG. 1987. Spin glasses and the statistical mechanics of protein folding. *Proc Natl Acad Sci USA* 84:7524–7528.
- Buss LW. 1987. *The evolution of individuality*. Princeton, NJ: Princeton University Press.
- Bussemaker HJ, Thirumalai D, Bhattacharjee JK. 1997. Thermodynamic stability of folded proteins against mutations. *Phys Rev Lett* 79:3530–3533.
- Cech TR, Damberger SH, Gutell RR. 1994. Representation of the secondary and tertiary structure of group I introns. *Nat Struct Biol* 1:273–280.
- Dill KA, Chan HS. 1997. From Levinthal to pathways to funnels. *Nat Struct Biol* 4:10–19.
- Eigen M. 1971. Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften* 58:465–523.
- Eigen M, McCaskill JS, Schuster P. 1989. The molecular quasispecies. *Adv Chem Phys* 75:149–263.
- Eklund EH, Szostak JW, Bartel DP. 1995. Structurally complex and highly active RNA ligases derived from random RNA sequences. *Science* 269:364–370.
- Ellington AD. 1994. RNA selection. Aptamers achieve the desired recognition. *Curr Biol* 4:427–429.
- Ellington AD, Szostak JW. 1990. In vitro selection of RNA molecules that bind specific ligands. *Nature* 346:818–822.
- Flamm C, Fontana W, Hofacker IL, Schuster P. 2000. RNA folding at elementary step resolution. *RNA* 6:325–338.
- Fontana W, Schuster P. 1987. A computer model of evolutionary optimization. *Biophys Chem* 26:123–147.
- Fontana W, Schuster P. 1998a. Continuity in evolution: on the nature of transitions. *Science* 280:1451–1455.
- Fontana W, Schuster P. 1998b. Shaping space: the possible and the attainable in RNA genotype-phenotype mapping. *J Theor Biol* 194:491–515.
- Fontana W, Konings DAM, Stadler PF, Schuster P. 1993a. Statistics of RNA secondary structures. *Biopolymers* 33:1389–1404.
- Fontana W, Stadler PF, Bornberg-Bauer EG, Griesmacher T, Hofacker IL, Tacker M, Tarazona P, Weinberger ED, Schuster P. 1993b. RNA folding and combinatorial landscapes. *Phys Rev E* 47:2083–2099.
- Freier SM, Kierzek R, Jaeger JA, Sugimoto N, Caruthers MH, Neilson T, Turner DH. 1986. Improved free-energy parameters for prediction of RNA duplex stability. *Proc Natl Acad Sci USA* 83:9373–9377.
- Gavrilets S, Hastings A. 1994. A quantitative-genetic model for selection on developmental noise. *Evolution* 48:1478–1486.
- Gibson G, Hogness DS. 1996. Effect of polymorphism in the *Drosophila* regulatory gene ultrabithorax on homeotic stability. *Science* 271:200–203.
- Gillespie DT. 1976. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J Comp Phys* 22:403–434.
- Gillespie DT. 1977. Exact stochastic simulation of coupled chemical reactions. *J Phys Chem* 81:2340–2361.
- Goldschmidt RB. 1940. *The material basis of evolution*. New Haven, CT: Yale University Press.
- Govindarajan S, Goldstein RA. 1997. The foldability landscape of model proteins. *Biopolymers* 42:427–438.
- Grüner W, Giegerich R, Strothmann D, Reidys C, Weber J, Hofacker IL, Stadler PF, Schuster P. 1996a. Analysis of RNA sequence structure maps by exhaustive enumeration. I. Neutral networks. *Monatsh Chem* 127:355–374.
- Grüner W, Giegerich R, Strothmann D, Reidys C, Weber J, Hofacker IL, Stadler PF, Schuster P. 1996b. Analysis of RNA sequence structure maps by exhaustive enumeration. II. Structure of neutral networks and shape space covering. *Monatsh Chem* 127:375–389.
- Hartwell LH, Hopfield JJ, Leibler S, Murray AW. 1999. From molecular to modular cell biology. *Nature* 402:C47–C52.
- He L, Kierzek R, SantaLucia J, Walter AE, Turner DH. 1991. Nearest-neighbour parameters for G-U mismatches. *Biochemistry* 30:11124.
- Hinton GE, Nowlan SJ. 1987. How learning can guide evolution. *Complex Systems* 1:495–502.
- Hofacker IL, Fontana W, Stadler PF, Bonhoeffer S, Tacker M, Schuster P. 1994. Fast folding and comparison of RNA secondary structures. *Monatsh Chem* 125(2):167–188.
- Hofacker IL, Fontana W, Stadler PF, Schuster P. 1994–1998. Vienna RNA Package. <http://www.tbi.univie.ac.at/~ivo/RNA/> (free software).

- Hofacker IL, Schuster P, Stadler PF. 1999. Combinatorics of RNA secondary structures. *Discrete Appl Math* 89:177–207.
- Huynen MA, Stadler PF, Fontana W. 1996. Smoothness within ruggedness: The role of neutrality in adaptation. *Proc Natl Acad Sci USA* 93:397–401.
- Huynen MA, Gutell R, Kongings D. 1997. Assessing the reliability of RNA folding using statistical mechanics. *J Mol Biol* 267:1104–1112.
- Jaeger JA, Turner DH, Zuker M. 1989. Improved predictions of secondary structures for RNA. *Proc Natl Acad Sci USA* 86:7706–7710.
- Joyce GF. 1989. Amplification, mutation and selection of catalytic RNA. *Gene* 82:83–87.
- Landweber LF. 1999. Experimental RNA evolution. *Trends Ecol Evolution (TREE)* 14:353–358.
- Landweber LF, Pokrovskaya ID. 1999. Emergence of a dual-catalytic RNA with metal-specific cleavage and ligase activities: the spandrels of RNA evolution. *Proc Natl Acad Sci USA* 96:173–178.
- Maynard-Smith J. 1987. Natural selection: when learning guides evolution. *Nature* 329:761–762.
- McCaskill JS. 1990. The equilibrium partition function and base pair binding probabilities for RNA secondary structure. *Biopolymers* 29:1105–1119.
- Mills DR, Peterson RL, Spiegelman S. 1967. An extracellular Darwinian experiment with a self-duplicating nucleic acid molecule. *Proc Natl Acad Sci USA* 58:217–224.
- Müller GB. 1990. Developmental mechanisms at the origin of morphological novelty: a side-effect hypothesis. In: Nitecki MH, editor. *Evolutionary Innovations*. Chicago: University of Chicago Press. p 99–130.
- Nussinov R, Jacobson AB. 1980. Fast algorithm for predicting the secondary structure of single-stranded RNA. *Proc Natl Acad Sci USA* 77(11):6309–6313.
- Nussinov R, Piecznik G, Griggs JR, Kleitman DJ. 1978. Algorithms for loop matching. *SIAM J Appl Math* 35(1):68–82.
- Reidys C, Stadler PF. 1996. Bio-molecular shapes and algebraic structures. *Comput Chem* 20:85–94.
- Reidys C, Stadler PF, Schuster P. 1997. Generic properties of combinatorial maps: neutral networks of RNA secondary structures. *Bull Math Biol* 59:339–397.
- Reidys C, Forst CV, Schuster P. 1998. Replication and mutation on neutral networks. Preprint SFI 98-04-036. Available electronically from www.santafe.edu. Santa Fe: Santa Fe Institute.
- Scheiner SM. 1993. Genetics and evolution of phenotypic plasticity. *Annu Rev Ecol Syst* 24:35–368.
- Schlichting CD, Pigliucci M. 1998. Phenotypic evolution: a reaction norm perspective. Sunderland, MA: Sinauer Associates, Inc.
- Schmalhausen II. 1949. In: Dobzhansky T, editor. *Factors of evolution*. Philadelphia: Blakiston.
- Schuster P. 1997. Landscapes and molecular evolution. *Physica D* 107:351–365.
- Schuster P, Fontana W. 1999. Chance and necessity in evolution: Lessons from RNA. *Physica D: Nonlinear Phenomena* 133:427–452.
- Schuster P, Fontana W, Stadler PF, Hofacker I. 1994. From sequences to shapes and back: a case study in RNA secondary structures. *Proc R Soc (London) B* 255:279–284.
- Simpson GG. 1953. The Baldwin effect. *Evolution* 7:110–117.
- Spiegelman S. 1971. An approach to the experimental analysis of precellular evolution. *Q Rev Biophys* 4:213–253.
- Stearns SC. 1993. The evolutionary links between fixed and variable traits. *Acta Paleontol Pol* 38:1–17.
- Stein PR, Waterman MS. 1978. On some new sequences generalizing the Catalan and Motzkin numbers. *Discrete Math* 26:261–272.
- Tuerk C, Gold L. 1990. Systematic evolution of ligands by exponential enrichment (SELEX): RNA ligands to bacteriophage T4 DNA polymerase. *Science* 249:505–510.
- Turner DH, Sugimoto N, Freier S. 1988. RNA structure prediction. *Annu Rev Biophys Chem* 17:167–192.
- van Nimwegen E, Crutchfield JP, Huynen M. 1999a. Neutral evolution of mutational robustness. Preprint SFI 99-03-021. Santa Fe: Santa Fe Institute. (Available electronically from www.santafe.edu.)
- van Nimwegen E, Crutchfield JP, Mitchell M. 1999b. Statistical dynamics of the royal road genetic algorithm. *Theor Comp Sci* 229:41–102.
- Vendruscolo M, Maritan A, Banavar JR. 1997. Stability threshold as a selection principle for protein design. *Phys Rev Lett* 78:3967–3970.
- Waddington CH. 1942. Canalization of development and the inheritance of acquired characters. *Nature* 3811:563–565.
- Waddington CH. 1957. *The strategy of the genes*. New York: MacMillan Co.
- Wagner A. 1996. Does evolutionary plasticity evolve? *Evolution* 50(3):1008–1023.
- Wagner A, Stadler PF. 1999. Viral RNA and evolved mutational robustness. *J Exp Zool/MDE* 285:119–127. (SFI working paper 99-02-010, available from <http://www.santafe.edu>.)
- Wagner GP, Altenberg L. 1996. Complex adaptations and the evolution of evolvability. *Evolution* 50:967–976.
- Wagner GP, Booth G, Bagheri-Chaichian H. 1997. A population genetic theory of canalization. *Evolution* 51:329–347.
- Wagner GP, Laubichler MD, Bagheri-Chaichian H. 1998. Genetic measurement theory of epistatic effects. *Genetica* 102/103:569–580.
- Waterman MS. 1995. *Introduction to computational biology: sequences, maps and genomes*. London: Chapman & Hall.
- Waterman MS, Smith TF. 1978. RNA secondary structure: a complete mathematical analysis. *Math Biosci* 42:257–266.
- Weber J. 1997. Dynamics of neutral evolution. Ph.D. dissertation, Friedrich-Schiller-Universität, Jena, Germany, 1997 (available electronically from <http://www.tbi.univie.ac.at>).
- Westhof E, Masquida B, Jaeger L. 1996. RNA tectonics: towards RNA design. *Folding Design* 1:R78–R88.
- Wuchty S, Fontana W, Hofacker IL, Schuster P. 1999. Complete suboptimal folding of RNA and the stability of secondary structures. *Biopolymers* 49:145–165.
- Zuker M, Stiegler P. 1981. Optimal computer folding of larger RNA sequences using thermodynamics and auxiliary information. *Nucleic Acids Res* 9:133–148.
- Zuker M, Sankoff D. 1984. RNA secondary structures and their prediction. *Bull Math Biol* 46(4):591–621.