The Role of Randomness in Darwinian Evolution*

Andreas Wagner†‡

Historically, one of the most controversial aspects of Darwinian evolution has been the prominent role that randomness and random change play in it. Most biologists agree that mutations in DNA have random effects on fitness. However, fitness is a highly simplified scalar representation of an enormously complex phenotype. Challenges to Darwinian thinking have focused on such complex phenotypes. Whether mutations affect such complex phenotypes randomly is ill understood. Here I discuss three very different classes of well-studied molecular phenotypes in which mutations cause nonrandom changes, based on our current knowledge. What is more, this nonrandomness facilitates evolutionary adaptation. Thus, living beings may translate DNA change into nonrandom phenotypic change that facilitates Darwinian evolution.

1. Introduction. "In ordinary English, a random event is one without order, predictability or pattern. The word connotes disaggregation, falling apart, formless anarchy, and fear." This quote from the late Stephen J. Gould (1993) illustrates one reason why many nonbiologists—even highly educated ones—may feel uncomfortable with Darwinian evolution: Darwinian evolution centrally involves chance or randomness. Specifically, it involves a combination of natural selection and *random* or *chance* variation on which selection feeds. Most such variation is caused by mutations, genetic changes in DNA.

The discomfort with chance and randomness can partly account for statements like this: "To speak of chance for a universe which presents

*Received February 2011; revised June 2011.

†To contact the author, please write to: University of Zurich, Institute of Evolutionary Biology and Environmental Studies, Y27-J-54, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland; e-mail: andreas.wagner@ieu.uzh.ch.

‡I acknowledge support through Swiss National Science Foundation grants 315200-116814, 315200-119697, and 315230-129708, as well as through the YeastX project of SystemsX.ch.

Philosophy of Science, 79 (January 2012) pp. 95–119. 0031-8248/2012/7901-0006$10.00

such a complex organization in its elements and such marvelous finality in its life would be equivalent to giving up the search for an explanation of the world as it appears to us" (Schonborn 2005). For a refutation of this view, see Laubichler et al. (2005).

The concept of randomness has given rise to three main currents of literature in biology and in the philosophy of biology. By far the most prominent current regards the question whether mutations affect fitness randomly (Simpson 1953; Mayr 1961; Sober 1984, 2000; Dawkins 1996; Futuyma 1998; Eble 1999). In the words of the philosopher Elliott Sober, "Mutations are said to be random in that they do not arise because they would be beneficial to the organisms in which they occur" (2000, 37). I will come back to this notion of randomness in section 3 below. A second current regards the question how to distinguish between natural selection and genetic drift. On the one hand, natural selection has random aspects. For example, the viability of an organism—an important aspect of fitness—is often expressed as a *probability* that the organism survives from a zygote to reproductive age. On the other hand, genetic drift arises from the random sampling of alleles or genotypes from one to the next generation. Because selection and drift both involve chance, the question arises how to properly distinguish between them (Beatty 1984; Millstein 2000, 2002). A third current regards deterministic chaos, apparently random behavior that may arise from deterministic interactions of a system's components. Deterministic chaos may occur in ecological and neural systems and raises broad questions about determinism (May 1976; Mackey and Glass 1977; Wimsatt 1980; Earman 1986; Hastings et al. 1993; Elbert et al. 1994; Glass 2001).

This article does not fall plainly within any of these three currents. It is closest to the first one because it revolves around the effects of mutations. However, there is one key difference: it does not focus just on fitness but on complex phenotypes and how mutations affect them. Supported by very recent evidence (Lipman and Wilbur 1991; Schuster et al. 1994; Ciliberti, Martin, and Wagner 2007a, 2007b; Rodrigues and Wagner 2009, 2011; Ferrada and Wagner 2010; Samal et al. 2010), it argues that the phenotypic variation on which natural selection feeds can be viewed as nonrandom and highly structured. Not only that, it is structured in ways that facilitate evolutionary adaptation and innovation.

The role of chance and randomness in evolution can be examined for three different and variable aspects of a living system. The first of them is an organism's genotype. The second is its phenotype, which has many different facets that range from an organism's form, to its physiology, down to the spatial fold of the proteins inside its cells. The third aspect—I just mentioned it—is fitness, which collapses the immense complexity of a phenotype onto a single scalar quantity that indicates how well an

organism is adapted to its environment. Here I will discuss random change on genotypes and fitness only briefly (Hartl and Clark 2007). My main focus is random change in complex phenotypes because variation in these phenotypes is the substrate of natural selection and because recent work sheds light on how mutations change these phenotypes.

In section 2, I define a notion of randomness and random change that is suitable for my purpose. In sections 3 and 4, I discuss random change in fitness and in genotypes, mainly for completeness of the exposition, and because they illustrate applications for the notion of randomness I use. The remainder of the article focuses on complex phenotypes. Section 5 discusses why visible macroscopic phenotypes are currently ill suited for my purpose. Section 6 introduces three classes of systems whose phenotypes are better suited, partly because they are involved in many evolutionary processes. Specifically, these systems are large-scale metabolic networks, regulatory gene circuits, and protein or RNA macromolecules. Section 7 discusses the genotypes and phenotypes of these systems, and section 8 discusses recent insights into how these phenotypes are organized in the space of all possible genotypes. The three system classes are very different, but phenotypic changes in them are structured and nonrandom, in a sense that section 9 makes clear. Not only that, they are structured in ways that facilitate evolutionary adaptation and innovation.

2. The Notion of Randomness. Colloquial uses of words such as "random" or "chance" face the imprecision and ambiguities of everyday language. If we want to avoid such ambiguities, we can turn to the language of mathematics. Albeit itself not without limitations (Chaitin 1975, 2001), mathematics may be our best chance of lending some precision to the word "random." The relevant branch of mathematics is probability theory. Fundamentally, probability theory rests on conceptual experiments, such as the tossing of coins, the rolling of dice, the dealing of cards, or the change of letters in a string of text such as DNA. Each such experiment must have a set of well-defined outcomes: heads or tails, the numbers one through six, all possible DNA strings, and so on. In the lingo of probability theory, these outcomes constitute a *sample space*. Each outcome is called an *event*. And each event has a *probability*, such as the probability one-half of tossing heads with a fair coin. Events, sample spaces, and probabilities are primitive and undefined notions of probability theory, much as points and straight lines are in Euclidean geometry (Feller 1968, chap. 1).

Imagine you tossed a coin many times and it showed heads in 80% of these tosses. Colloquially, we would say that there is something nonrandom about how this coin falls. But this may not be so from the perspective of probability theory. The coin may simply not be a "fair" coin. For

example, one of its sides may be heavier than the other to the extent that it is much more likely to show a head than a tail. We can take this example to an extreme and imagine a coin that always shows heads, that is, with probability one. In common usage, such a coin toss has a decidedly nonrandom outcome. From a probabilistic standpoint, it is just an extreme example of a coin that is not fair, with a probability of showing heads equal to one.[1]

Probability theory is a way of viewing the world. Through its glasses, every process in the world becomes a random process. This is unhelpful if we want to ask in which sense Darwinian evolution might involve random change. To simply state that everything about the world is random leaves us unsatisfied. Most of us feel that there is a difference between a coin showing heads half of the time and showing heads 80% of the time.

The difference is that we have an unspoken expectation about the outcome of a random coin toss. It should produce heads about 50% of the time. In other words, both heads and tails should be *equiprobable*. This expectation is based on our prior experience with coin tosses and games of chance. It is also based on tacit assumptions about how a coin is—or should be—manufactured, namely, with equal mass on both sides.

This observation characterizes an important colloquial use of randomness: we use the notion of randomness to characterize an expected outcome of events in the world; deviations from this outcome constitute nonrandomness. In the absence of any other information, we often expect that possible events occur with equal probability. If our observations are consistent with this expectation, we say that they occur randomly. If they violate our expectation, we call them nonrandom. Equiprobability is not the only possible expectation, but some expectation must exist for this colloquial notion of randomness to apply. Unfortunately, this expectation is often tacit, unacknowledged, and imprecise.

This notion of randomness—an event's expected outcome by chance alone—is not only colloquially important. It has been made precise in the statistical notion of hypothesis testing, where the expectations I empha-

1. This is much less absurd than it may seem, especially if one studies continuous sample spaces with infinitely many members. For example, imagine you could choose one number at random among the real numbers between zero and 10 such that every real number has an equal probability of being chosen. In this case, the problem of not choosing any particular number, say 3/2, is equal to one because there are infinitely many (yes, uncountably many) such numbers. At the same time, some number will be the chosen one, even though the probability of choosing it was infinitesimally small. In mathematical language, the subset of the interval (0, 10) that corresponds to this number is a subset of measure zero. Even from a colloquial standpoint, we would feel comfortable referring to such a choice as a random choice, even though a specific outcome may have zero probability.

sized are called *null hypotheses* (Sokal and Rohlf 1981). In hypothesis testing, one asks whether a series of events could have occurred by chance alone, meaning that they are consistent with a prior expectation or, in statistical language, with a null hypothesis. If not, the null hypothesis is rejected. For example, for our earlier coin toss, a relevant null hypothesis is that we will observe heads half of the time. Statistical tests tell us how much deviation from this expectation is tolerable if we toss a coin N times. If the frequency of heads lies outside the range of tolerance, we would say that it is not expected by chance alone. It is nonrandom in this sense.[2]

I will use this notion of randomness here: consistency with an explicit expected outcome, that is, with a null hypothesis or a (statistical) model of a process. It is a more precise version of the colloquial use of randomness and is useful to discuss several notions of randomness in Darwinian evolution. I will also find it useful to use the notions of sample space and events because they help us make our expectations more precise.

3. Randomness and Fitness. The most widely discussed notion of randomness in evolutionary biology regards the effects of mutations on an organism's fitness. Here, one can distinguish at least two kinds of events: mutations that are good (beneficial) and those that are bad (deleterious) for fitness. If you did not know much about our world, you might expect that mutations are equally likely to be good or bad. If you knew more (and especially if you had children), you would be aware that a haphazard change of any one object—toy, machine, and so forth—is more likely to break it than to improve its function. You might extrapolate this insight to living things and thus argue that most mutations might be deleterious rather than beneficial. These observations can form the basis of an expectation defining random effects on fitness: random mutation would typically be deleterious, not beneficial, to an organism. They would typically reduce its fitness. In contrast, if mutations were nonrandom, they might be mostly beneficial or even always beneficial, in which case natural selection might become unnecessary. In other words, the more strongly nonrandom mutations are, the less important natural selection would be in organic evolution.

All evidence available thus far suggests that mutations are random in

2. Strictly speaking, an observation's deviation from expectations implies either that it is nonrandom or that it is very unusual. In practice, replication of observations can be used to distinguish between these two possibilities. One should be aware, though, that statistical tests by their very nature do not provide absolute certainty about nonrandomness. They are merely the best practical means to distinguish random from nonrandom observations. However, many successes of modern science and engineering are built on them. This power of statistical testing provides another motivation to use a statistical definition of randomness here.

this sense. They do not preferentially cause an organism's fitness to increase, much as a broken part in a machine does not usually cause the machine to work better. Occasionally this notion is challenged by experimental data. The last serious challenge occurred decades ago (Cairns, Overbaugh, and Miller 1988), but it has been readily deflected by better data (Hendrickson et al. 2002). And with the analogy of man-made devices in mind, we can see that this randomness of mutations may be with us for a good long time. Organisms are highly ordered and extremely complex systems, much more so than machines, and perform complex activities (feeding, self-defense, reproduction, etc.) to persist. The chances of improving such a system through haphazard change are small.

In sum, evolution is random in the sense that mutations do not usually improve fitness. I need not say more about this subject because many others have (Simpson 1953; Mayr 1961; Sober 1984, 2000; Dawkins 1996; Futuyma 1998; Eble 1999). However, it is worth saying that fitness—a simple scalar quantity—is not even a caricature of an organism's phenotypic complexity. As I will show below, if we focus on complex phenotypes, we can learn more interesting things about the role of nonrandom variation in evolution. Before focusing on such phenotypes, however, I need to discuss how mutations may affect genotypes, because any effort to discuss randomness without them would be incomplete.

4. Randomness and Genotype. It is often stated that mutations are random changes in DNA, meaning that they affect a genotype randomly. The relevant sample space is the space of all possible DNA molecules. This space is also sometimes called a *sequence space* or *genotype space*. Each sequence is a single member or point in this space and constitutes one of the possible *events* or outcomes of mutation. That is, the *events* are all possible sequences in this space. If we ask how mutation changes a string of DNA, say, the coding region of a given gene in an organism's genome, then we can distinguish multiple different kinds of mutations. Some mutations—point mutations—change one or more single nucleotides into some other nucleotide; others—inversions—change the orientation of a DNA molecule; yet others duplicate part of the DNA string; and so on.

What are the hidden expectations behind the statement that mutations are random changes in DNA? The simplest possible expectation would be that all events in our sample space are equally likely. This would mean that any one mutation could create all possible sequences and would do so with equal probability. To anybody who knows the first thing about biochemistry and about the molecular mechanisms behind mutation, this expectation is ludicrous. For example, point mutations usually change one base pair at a time, and inversions can produce only the reverse

complement of a (double-stranded) DNA text before mutation. The biochemical mechanisms behind genotypic change severely constrain how mutations alter DNA. If we adopted the very naive expectation I just mentioned, mutations would be clearly nonrandom in their effects on DNA.

A less naive expectation is that random point mutations should change any nucleotide (A, C, G, or T) into any other nucleotide with equal probability. Any deviation from this pattern would constitute nonrandomness. Molecular evolutionists have developed sophisticated methods to determine whether mutation is random in this sense. These methods rely on a combination of comparative analysis and experimental work and are greatly aided by our ability to determine the DNA sequences of entire genomes (Li 1997; Drake et al. 1998; Omilian et al. 2006; Denver et al. 2009; Ossowski et al. 2009). They show that the above expectation is wrong. We know, for example, that transition mutations (A $\leftrightarrow$ G, C $\leftrightarrow$ T) are typically twice as frequent as transversion mutations (A, G $\leftrightarrow$ C, T), even though there are only two possible transition mutations but four possible transversion mutations (Li 1997). The reason lies in the biochemical mechanisms of mutation. For example, the chemical structure of the base adenine (A) is much more similar to that of guanine (G) than it is to the other two bases, which makes it much easier to convert A into G or vice versa. Thus, our expectation is violated, and we might call mutation nonrandom from this perspective.

Once we know of this so-called transition-transversion bias, we can form new, better-informed expectations. For example, we might call point mutations random if the likelihood that a mutation transforms a base into another base depends only on whether the mutation would be a transition or transversion but would otherwise be independent of the base considered. This expectation might seem quite reasonable, but it also turns out to be violated. For example, mutations are often context dependent. That is, whether, say, a C mutates into a T may depend on whether there is a G next to it (Morton 2003; Niu, Lin, and Zhang 2003; Jia and Higgs 2008; Touchon and Rocha 2008). Thus, mutation is nonrandom with respect to the expectation above.

These are just two, increasingly better-informed, expectations of what constitutes random mutation. I could propose many more on the basis of what we know about mutation. For example, some bases are methylated, which influences their propensity to mutate; so does the active replication of DNA, which favors certain kinds of point mutations over others; and the DNA strand ("top" or "bottom") of the double-stranded DNA helix in which a base occurs also influences the kind of changes that this base can undergo (Morton 2003; Niu et al. 2003; Jia and Higgs 2008; Touchon and Rocha 2008). The list could go on and on.

This example illustrates an important principle. Nonrandomness in the sense I defined it—as a deviation from a prior expectation—is a property not just of a natural phenomenon but also of our knowledge about this phenomenon. As our knowledge increases, this expectation may change. Nonrandomness is thus a moving target. Whether we call mutations in genotypes random may depend on our knowledge about genotypes and the mechanisms behind their change.

Even at our present state of knowledge, after decades of studying DNA sequences, and with billions of DNA sequences in public databases available for analysis, we do not have an expectation sophisticated enough to explain the observed patterns of genotypic change we see in nature. And this holds despite considerable effort to look for statistical models of DNA (Karlin and Brendel 1993; Karlin and Cardon 1994). Thus, genotypic change is nonrandom on the basis of what we presently know.[3]

5. Randomness and Complex Phenotypes. In section 4, I discussed genotypic change, which illustrates how our knowledge can shape our expectations about randomness. Such expectations are key to the statistical definition of randomness that I use here (sec. 2). Before that, in section 3, I discussed change in fitness and the consensus view that mutations do affect fitness randomly. Both sections are included here for the sake of completeness, but they are a mere lead-in to my main focus, complex phenotypes.

There are at least three reasons to focus a discussion of randomness on complex phenotypes. First, natural selection does not directly act on genotypes but on phenotypes. Second, fitness is a highly aggregate property, a compressed scalar representation of immensely complex phenotypes, traits of living systems that extend through space and time. One consequence of reducing complex phenotypes to simple fitness is this: the expectation behind the assertion that mutations affect fitness randomly, namely, that mutations are more likely to damage a system than improve it, is very crude. As I mentioned in section 2, prior knowledge about mechanisms of change can be an important basis for expectations like this. However, because fitness reflects so many different aspects of phe-

3. At some point in the distant future, after having studied changing DNA for many more years, we may discover an appropriate null hypothesis, a statistical model of changing DNA. With such a model in hand, any mutational change in DNA would become a random change under the definition of randomness I use here. The chances of finding such a general null hypothesis, however, may be slim, for it is well known that rates and patterns of mutations depend in subtle ways on an organism's genotype and on the environment that it finds itself in. Any such null hypothesis may be specific to one organism and one environment rather than general enough for a comprehensive understanding of random genotypic change in evolution.

notype, there may not be one mechanism of change that could help us refine this crude expectation. Only a less compressed notion of phenotype could circumvent this problem. A third reason to focus on complex phenotypes is that they encompass the astounding complexity of life that Darwinian evolution aims to explain. For these reasons, I will now turn to complex phenotypes.

Until recently, we did not have sufficient knowledge about such phenotypes to ascertain whether mutations affect phenotypes randomly. Those studying *phenotypic constraints* might cringe at this statement and argue that, yes, we have known for a while that phenotypes change nonrandomly. A phenotypic constraint is a bias or limitation in phenotypic variation that a biological system produces (Maynard-Smith et al. 1985). Phenotypic constraints are everywhere. Dramatic examples include the nonexistence of birds with horns or the general lack of teeth in the lower jaw of frogs (Maynard-Smith et al. 1985; Futuyma 1998). Many more subtle examples exist (Maynard-Smith et al. 1985). Together, they show that not all conceivable phenotypic variation exists in nature.

The problem with phenotypic constraints is that many candidate constrained characters are macroscopic characters. For such characters, it is not clear what the relevant sample space—the space of all possible characters—and possible events following a mutation should be. Efforts to define and explore a "morphospace" of macroscopic traits (Gould 1991; Newman and Bhat 2009) have thus far found limited resonance, perhaps because many macroscopic traits are too complex to be easily embedded in such a space. In addition, to characterize the relevant sample space, it is useful to understand how change in genotypes affects phenotypes, that is, how genotypes map onto phenotypes. The reason is that most phenotypic change that is heritable and thus relevant for Darwinian evolution is caused by genotypic change. We do not understand the relationship between genotypic and phenotypic change well for macroscopic traits.

Thus, for most macroscopic traits, we cannot currently make clear what would constitute an expectation for random phenotypic variation and what would constitute constrained variation. This is why the rich existing literature on phenotypic constraints contains many ambiguities (Odell et al. 1981; Cheverud 1984; Maynard-Smith et al. 1985; Shubin and Alberch 1986; Oster et al. 1988; Hodin 2000; Brakefield 2006; Newman and Bhat 2009). For the same reason, phenotypic constraints are not exempt from the assertion that we are ignorant about the randomness of phenotypic variation.

To understand the role of randomness in phenotypic change, we need to study phenotypes that are complex but not too complex to apply concepts such as that of a sample space to them. I will next discuss three very different classes of systems that meet this requirement. For these

systems, we are beginning to understand the relationship between genotype and phenotype. They are the basis of many evolutionary adaptations and innovations and thus are central to Darwinian evolution. These features make them attractive study objects. As we shall see, phenotypic variation in these systems can be viewed as nonrandom. And here is the most important feature of this nonrandomness: it can facilitate evolutionary adaptations and innovation. As opposed to constraints, which connote a restricting and confining role in Darwinian evolution, this kind of nonrandomness can play a positive role in Darwinian evolution.

6. Three Classes of Systems Important to Evolutionary Adaptation and Innovation. All macroscopic traits are composed of microscopic, submicroscopic, and molecular traits, down to the level of DNA. And DNA change can percolate all the way up to macroscopic traits. To understand the full complexity of this hierarchical organization is beyond our current means. But we can study important classes of systems that we know are involved in producing macroscopic traits and changes therein.

One class of such systems is genome scale metabolic networks. These are networks of hundreds to thousands of chemical reactions—catalyzed by enzymes that are encoded by genes—that synthesize all small molecules in biomass from environmental nutrients. In addition, they produce energy and many important secondary metabolites. The metabolic phenotypes of such networks are involved in many new phenotypes that occur anywhere from microbes to higher organisms. Examples include the ability of microbes to grow on synthetic antibiotics or other toxic xenobiotic compounds, such as polychlorinated biphenyls, chlorobenzenes, or pentachlorophenol (Cline et al. 1989; van der Meer 1995; van der Meer et al. 1998; Copley 2000; Rehmann and Daugulis 2008). They also include the urea cycle, a metabolic innovation of land-living animals that allows them to convert toxic ammonia into urea for excretion.

Novel metabolic traits often involve new combinations of chemical reactions (enzymes) that already exist elsewhere. For example, a novel metabolic pathway to degrade pentachlorophenol involves four steps that its host organism assembled—probably through horizontal gene transfer—from enzymes processing naturally occurring chlorinated chemicals as well as from an enzyme involved in tyrosine metabolism (Copley 2000). The urea cycle arose when four widespread enzymatic reactions involved in arginine biosynthesis combined with arginase, a reaction involved in arginine degradation (Takiguchi et al. 1989)

The second system class is regulatory circuits. They are systems of interacting gene products that influence each other's biological activity. Their phenotypes are gene expression phenotypes or, more generally, molecular activities of gene products with important biological functions.

Such circuits are involved whenever cells and tissues communicate and whenever gene expression is regulated (Gilbert 1997; Carroll, Grenier, and Weatherbee 2001). Both processes are indispensable for the development of any multicellular organism and thus for the formation of all macroscopic phenotypes. The most important kinds of circuits are transcriptional regulation circuits because transcriptional regulation provides a regulatory backbone to most organisms and because such circuits drive many pattern formation processes in embryonic development. The genes in such circuits encode transcriptional regulators that bind DNA near other genes and regulate their transcription. Among the best-known examples are Hox genes, which pattern limbs and many other body structures in animals, and MADS box genes that are involved in patterning flowers (Hughes and Kaufman 2002; Irish 2003; Wagner, Amemiya, and Ruddle 2003; Causier et al. 2005; Lemons and McGinnis 2006; Hueber and Lohmann 2008).

Regulatory change is also involved in forming many new macroscopic phenotypes. For example, the KNOX (KNOTTED1-like homeobox) transcription factors and their overexpression may stand behind the origin of dissected leaves, an innovation of some plants that may aid thermoregulation (Bharathan et al. 2002). The predator-deterring eyespots of butterflies form where the transcription factor *Distal-less* is overexpressed. In these and many other examples (Carroll et al. 2001), changes in the regulation of existing molecules that also serve other purposes are associated with the formation of novel phenotypes.

The third and final system class is protein and RNA macromolecules. Their phenotypes are three-dimensional molecular structures and their varied functions. They catalyze thousands of biochemical reactions, exchange chemicals between cells and their environment, give structural support to cells, are central to locomotion, and serve many other useful functions. Not surprisingly then, important adaptations in evolution are traceable to changes in macromolecules. One example regards antifreeze proteins, which allow organisms in cold environments to survive where the body temperature of others would freeze. They evolved independently, rapidly, and from different ancestors in arctic and antarctic fish (Chen, DeVries, and Cheng 1997; Cheng 1998). Another example involves the hemoglobin molecules important for oxygen transport. In the hemoglobin subunits of the bar-headed goose (*Anser indicus*), a single proline to alanine substitution in one of hemoglobin's subunits increases the protein's affinity to oxygen. The change helps this bird migrate over the Himalayas at altitudes exceeding 10 kilometers (Golding and Dean 1998; Liang et al. 2001).

7. Genotypes and Phenotypes in the Three Study Systems. In all three

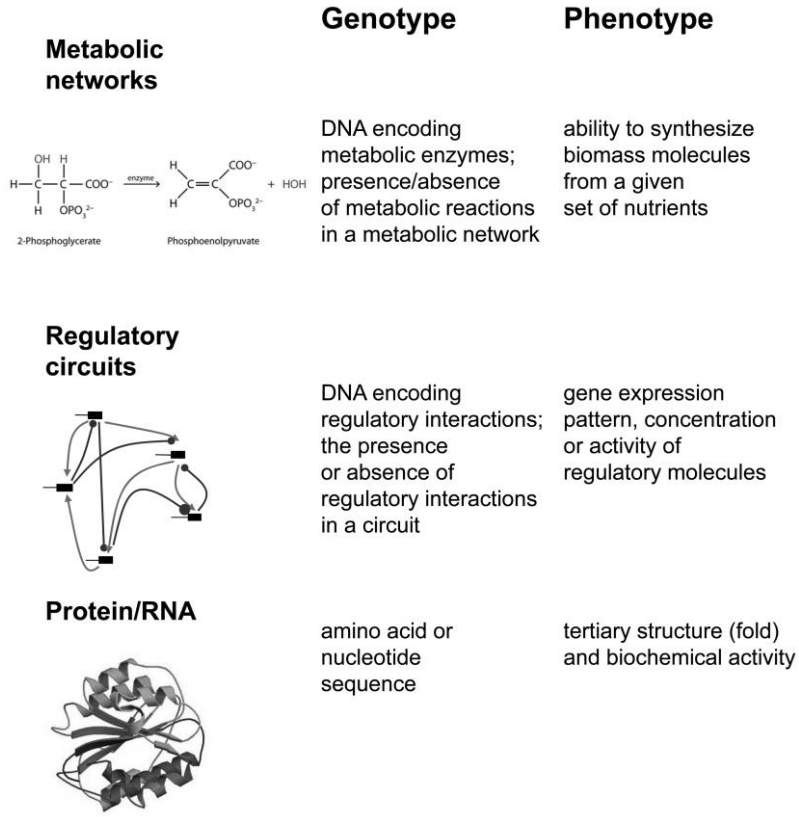

Figure 1. Overview over the concepts of genotypes and phenotypes in the three system classes discussed in this article.

system classes that I just discussed, we have recently gained a better understanding of the relationship between genotype and phenotype. I will next summarize some insights from recent work on these systems (Lipman and Wilbur 1991; Schuster et al. 1994; Fontana and Schuster 1998; MacCarthy, Seymour, and Pomiankowski 2003; Ciliberti et al. 2007a, 2007b; Rodrigues and Wagner 2009, 2011; Ferrada and Wagner 2010; Samal et al. 2010).

Genotypes in all three systems (fig. 1) exist in a vast genotype space. Although such genotypes are ultimately DNA sequences, it is often useful to represent them in more compact ways. For example, a metabolic genotype is the part of an organism's genome that encodes metabolic enzymes, but rather than representing this genotype as a DNA string, it can

be represented more compactly through information about the presence or absence of individual enzymes or of individual enzyme-catalyzed reactions in a metabolic network. In other words, one can represent an organism's metabolic genotype through the set of enzymes or enzyme-catalyzed reactions that occur in its metabolic network.

The genotype space of metabolic networks is the space of all possible metabolic networks. The current known "universe" of metabolic reactions comprises more than 5,000 such reactions, each of which can be present or absent in any one network. Thus, metabolic genotype space is a vast, hyperastronomical space comprising more than $2^{5,000}$ metabolic genotypes, each one corresponding to a metabolism with a different set of enzyme-catalyzed reactions (Rodrigues and Wagner 2009; Samal et al. 2010).

Metabolic phenotypes emerge from the joint action of all chemical reactions in a metabolic network, from how these reactions cooperate to transform energy and chemical elements contained in environmental nutrients into biomass. There are different ways of classifying metabolic phenotypes. One useful approach is to classify phenotypes according to their ability to sustain life—to synthesize biomass—in different chemical environments (Rodrigues and Wagner 2009). For example, if one focuses on carbon metabolism, one can ask which molecules can serve as sole carbon and energy sources for a metabolic network. To represent such phenotypes systematically, one can use some number of common carbon sources, say 100 different molecules, and write these down as a list. A metabolic phenotype can then be represented as a binary string, where one writes a 1 next to a carbon source in our list if the network can sustain life on it and a 0 if it cannot. Recent methodological advances in computational biochemistry allow us to compute metabolic phenotypes from metabolic genotypes (Price et al. 2003). Note that for 100 carbon sources, there is already an astronomical number of 2^{100} possible metabolic phenotypes, each of them encapsulating viability in a different spectrum of chemical environments. Analogous classifications are possible for sources of other chemical elements (Rodrigues and Wagner 2011).

The genotypes of regulatory circuits (fig. 1) are DNA sequences that encode transcriptional regulators as well as the regulatory regions on DNA that these regulators bind to and that mediate gene regulation. Instead of representing these genotypes as DNA strings, one can also represent them more compactly as patterns of regulatory interactions. In a transcriptional regulation circuit, for example, any one gene X can have an activating, repressing, or no effect on the expression of another gene Y, as determined partly by regulatory DNA sequences near gene Y. The space of all circuits comprises all possible patterns of interactions between a given set of genes. These interactions bring forth a circuit's gene expression phenotype (MacCarthy et al. 2003; Wagner 2005; Ciliberti et al.

2007a, 2007b). Even if one considers only whether a gene is expressed or not, then a circuit of N genes has 2^N possible phenotypes. Finer distinctions, among different levels of expression, would give rise to even more phenotypes. Thus, also in this system class, many possible phenotypes exist. One can classify them and compare them systematically (Ciliberti et al. 2007a, 2007b).

The genotypes of molecules, finally, exist in a space of amino acid and nucleotide sequences. Exploration of their genotype spaces has the longest history (Maynard-Smith 1970; Lipman and Wilbur 1991; Schuster et al. 1994). The relevant phenotypes include secondary and tertiary structures of protein and RNA molecules, as well as enzymatic and other functions of these molecules. One can classify and compare such phenotypes systematically. Examples include protein structures and enzymatic functions (Todd, Orengo, and Thornton 2001; Levitt 2009). The number of such phenotypes is again large. For example, even if one considers only single-domain polypeptides—proteins that consist of one autonomously folding amino acid string—there are more than 10^4 known protein tertiary structure phenotypes (Levitt 2009). If one considers proteins with multiple domains and takes biochemical functions into account, this number would become astronomical. Mapping genotypes onto phenotypes for macromolecules is becoming increasingly feasible. It involves computational predictions of structure phenotypes but also the analysis of existing data on the sequence, structure, and function of tens of thousands of molecules (Todd et al. 2001; Ferrada and Wagner 2008, 2010; Levitt 2009).

In sum, metabolic networks, regulatory circuits, and macromolecules are important for the formation of most phenotypes: they allow us to classify and compare both genotypes and phenotypes, and they permit us to infer or predict phenotype from genotype to some extent. These three features make them ideal to ask whether phenotypic change is random or not.

8. The Organization of Phenotypes in Genotype Space. The genotype spaces of the above system classes share several important similarities (Lipman and Wilbur 1991; Schuster et al. 1994; Fontana and Schuster 1998; Ciliberti et al. 2007a, 2007b; Rodrigues and Wagner 2009, 2011; Ferrada and Wagner 2010; Samal et al. 2010). First, typically an astronomical or hyperastronomical number of genotypes have the same phenotype. For example, for part of the λ repressor, a protein from the bacteriophage λ, 5×10^{56} sequences may yield a functional protein (Reidhaar-Olson and Sauer 1990). Note that even large numbers like this may correspond to a tiny fraction of genotype space. In the case of the λ repressor, this fraction may amount to only one 10^{-63}rd of all amino acid sequences (Reidhaar-Olson and Sauer 1990). Analogous observations

hold for the phenotypes of metabolic networks, regulatory circuits, and RNA molecules (Gruner et al. 1996; Ciliberti et al. 2007b; Jörg, Martin, and Wagner 2008; Samal et al. 2010).

Before discussing the second commonality of genotype space organization, I need to introduce the concepts of a *neighbor* and of a *neighborhood* in genotype space. Two metabolic network genotypes are neighbors if they differ in exactly one chemical reaction, two regulatory circuits are neighbors if they differ in exactly one regulatory interaction, and two proteins or RNA molecules are neighbors if they differ in exactly one amino acid or nucleotide. A genotype's neighborhood comprises all its neighbors. I note that a neighborhood typically comprises many neighbors. For example, each protein of 100 amino acids has $100 \times 19 = 1{,}900$ neighbors because each of the protein's 20 amino acids can change into 19 other amino acids. The notions of neighbor and neighborhood can be extended to genotypes that differ in some number $k > 1$ of metabolic reactions, regulatory interactions, or amino acids.

These concepts allow me to introduce the second commonality of the three system classes I discuss: genotypes with the same phenotype form vast connected genotype networks that extend far through genotype space. This means that one can step from one genotype to its neighbor, to the neighbor's neighbor, and so on without ever changing a phenotype. Thus, not only can very different genotypes have the same phenotype, these genotypes can also be connected in genotype space. The many small open circles in figure 2*a* are a caricature of such a network in a much larger genotype space (*large rectangle*). Circles that are neighbors are connected by a straight line. The network that these circles form illustrates that a genotype network can extend far through genotype space.[4]

A third commonality is that the neighborhood of any two genotypes on the same genotype network (i.e., with the same phenotype) contains very different novel phenotypes. This means that even if two genotypes differ only modestly, their neighborhoods do not contain the same phenotypes. The two dashed circles in figure 2*a* illustrate this principle. Each circle circumscribes the neighborhood of the genotype in the circle's center. Symbols of different shape and shading indicate genotypes whose phenotypes are different from those on the genotype network shown. For the genotype in the center of the left circle, three neighbors are shown, two of which have different phenotypes (*filled star* and *open hexagon*). For the genotype in the center of the right circle, four neighbors are shown,

4. The observation that the same phenotype can be formed by many genotypes has implications for the problem of reductionism in biology (Hull 1972, 1979) as well as on the question of what genes are and what exactly it is they encode (Griffiths and Neumann-Held 1999; Griffiths and Stotz 2007).

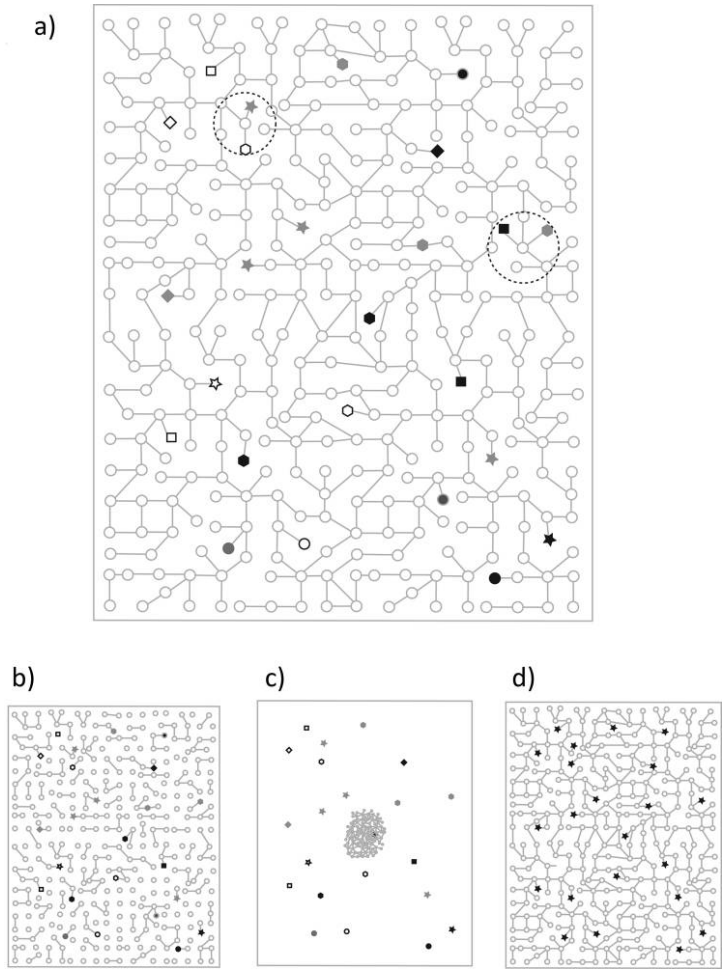

Figure 2. Connected genotype networks facilitate accessibility of diverse new phenotypes. Panel *a* schematically represents a set of hypothetical genotypes (*open circles*) in a genotype space (*rectangle*) that all share the same phenotype and form a single genotype network; neighboring genotypes are connected by lines. Symbols of different shapes and shading indicate genotypes whose phenotypes differ from the phenotype shared by all genotypes on this genotype network. The two dashed circles denote the neighborhood of two different genotypes on the genotype network. The figure illustrates that many different novel phenotypes can be accessed from a connected genotype network that spreads far through genotype space. Panels *b*, *c*, and *d* indicate three counterfactual scenarios for genotype network organization. Source: Wagner (2011). See the text for details.

two of which have different phenotypes (*filled square* and *filled hexagon*). The phenotypes (symbols) in the two neighborhoods are very different.

I note that genotype spaces are high-dimensional spaces whose features cannot be captured adequately in a two-dimensional figure. For example, a genotype may have hundreds to thousands of neighbors, only a few of which are shown for each genotype in figure 2*a*. Also, each genotype whose phenotype is different from that of the genotype network in figure 2*a* would itself be part of a vast genotype network that is not shown.

Together, the properties I just discussed facilitate the exploration of novel phenotypes through blind evolutionary searches in genotype space. They solve a major problem that the discovery of novel and useful phenotypes poses to living systems: they need to preserve existing, well-adapted phenotypes while exploring innumerable new phenotypes, only a few of which may be improvements over the status quo. Here is how they accomplish this: envision a population of organisms that preserves its existing phenotype (through stabilizing selection) while being exposed to mutational change. The existence of genotype networks means that the population can gradually change its genotype while preserving its phenotype. In doing so, it can explore different regions of genotype space. The immediate neighborhood of the population will contain very different novel phenotypes, depending on where its members are located in genotype space. The existence of genotype networks combined with the diversity of their neighborhoods thus allows exploration of many novel phenotypes.[5]

Both the existence of genotype networks and the diversity of their neighborhoods are necessary to explore many novel phenotypes. To see this, consider several counterfactual scenarios shown in figure 2*b*–2*d*. In the first scenario (fig. 2*b*), the number of genotypes that form the same phenotype is just as large as in figure 2*a*, and these genotypes are just as widely distributed through sequence space. However, these genotypes ei-

5. The evolutionary dynamics I describe here resembles that on adaptive landscapes used to study the evolution of reproductive isolation, such as in Sewall Wright's shifting balance theory (Futuyma 1998, 408) and Gavrilets's (1997) holey adaptive landscapes, in that evolutionary change can be neutral and occurs preferably along adaptive ridges with high fitness. However, two differences from this previous work are more important than these similarities. The first is that the present approach considers not just scalar fitness but multivariate and complex phenotypes. To study speciation and reproductive isolation, it may be adequate to consider just fitness, but doing so is no longer sufficient if one wants to study how novel phenotypes originate. The second major difference is that in the older models, scalar fitness is often computed from simple population genetic models or it is simply randomly assigned to genotypes. In contrast, the complex phenotypes I discuss here are based on a deeper, mechanistic understanding of how phenotypes emerge from information in genotypes.

ther are isolated from one another or form only small groups of connected genotypes. Their disconnectedness hinders access of new phenotypic variants because evolving genotypes remain confined to small regions of this space. They can no longer explore large regions of this space through mutations that leave the phenotype unchanged.

The second scenario (fig. 2*c*) shows a genotype network that is connected but does not extend far through genotype space. Instead, it is confined to a small region of this space. Therefore, many novel phenotypes occurring elsewhere in genotype space remain inaccessible to it. Figure 2*d* shows a final counterfactual scenario, where a sprawling and connected genotype network exists but the phenotypes in its neighborhoods are all the same. In this case, the network is irrelevant to explore novel phenotypes because regardless of where a genotype occurs on this network, it has access to the same novel phenotypes.

I emphasize again that the imagery of figure 2 has to be taken with a grain of salt. Genotype spaces are very high-dimensional spaces akin to hypercubes, where any one genotype may have thousands of neighbors. A two-dimensional caricature like that of figure 2 can thus mislead in several ways. Nonetheless, it serves to provide a modicum of intuition about genotype space organization.

9. Back to Random Phenotypic Change. To ask whether a mutation has random effects on any of the complex phenotypes I discuss here, we need to first identify an expectation defining such randomness. Recall that such expectations are at the core of the statistical notion of randomness that I use here (sec. 2). For simplicity, I will here discuss a very common class of mutations, point mutations that change a genotype into its neighbor.[6]

One naive expectation is that a mutation of any one genotype could produce all phenotypes, and do so with equal probability for all phenotypes. However, in all three systems classes, this is not so. And once we have understood the structure of genotype space, it becomes clear that it could not be so. The reason is that the number of possible phenotypes vastly exceeds the number of neighbors any one genotype has. A genotype may have thousands of neighbors, but their numbers pale compared to

6. I tacitly make the simplifying assumption that the genotypic change that underlies this phenotypic change is random under a simple statistical model, namely, that each system part, such as an amino acid in a protein, is equally likely to be replaced with any other system part, e.g., any other amino acid. Analogous assumptions are possible for regulatory circuits and metabolic networks. Even this simplifying assumption will lead me below to the conclusion that phenotypic change is nonrandom. Because genotypic change is actually not random, as I argued in sec. 4, phenotypic change deviates even more from randomness than this simple model suggests.

the astronomical numbers of possible phenotypes. There is simply not enough room in any one neighborhood to accommodate all phenotypes.

Thus, we need to form a more sophisticated expectation. If we want to understand effects of mutations that go beyond any one genotype and its phenotype, it must be an expectation about the organization of phenotypes in genotype space. The following expectation fits the bill. Let us go through all possible genotypes one by one. For the first genotype, pick one of the possible phenotypes at random (with equal probability to choose each phenotype) and assign it to this genotype. Repeat this procedure for the next genotype, and so on, until you have assigned a phenotype to each genotype. The result is a random and equiprobable assignment of a phenotype to each genotype. For brevity, I will call this organization or distribution of phenotypes in genotype space *random.* (My earlier caveat that any distribution could be viewed as random applies.) Are the known effects of mutations on phenotypes consistent with this random distribution? If so, we could say that mutations affect phenotypes randomly.

The consequences of this random organization are no longer so intuitive, but they can be identified mathematically. The first consequence is this: under a random organization of phenotypes, the neighborhood of two different genotypes (whether they have the same phenotype or not) will generally contain very different phenotypes (Ciliberti et al. 2007b; Wagner 2011). This holds as long as there are many more phenotypes than a genotype has neighbors. This prediction looks reassuring because it resembles the observation of diverse phenotypic neighborhoods that I discussed above. Unfortunately, it is wrong in many details. For example, many neighbors of any one genotype G with phenotype P have a phenotype similar to P, and not just any arbitrary phenotype (Sumedha, Martin, and Wagner 2007). We can, however, leave this issue aside because it pales in comparison to the following fatal problem.

For randomly organized phenotypes, do genotypes with the same phenotype form extended genotype networks? That is, are many or all of them connected? The answer is a resounding no. The vast majority of genotypes with any given phenotype will be isolated. That is, a genotype will typically have no neighbors with the same phenotype. Random phenotype organization does not support the existence of genotype networks. A mathematical argument is needed to see this (Erdös and Renyi 1960; Bollobas, Kohayakawa, and Luczak 1992; Reidys and Stadler 2002; Ciliberti et al. 2007a), but fundamentally, the reason is that the number of genotypes with any one phenotype—albeit large—constitutes a tiny fraction of genotype space.

Genotype networks are thus a nonrandom feature of genotype space organization, under the expectation I just stated. This also means that

the effects of mutations on phenotypes are nonrandom. They are more likely to preserve a phenotype than expected by chance. This feature, combined with the phenotypic diversity of different phenotypic neighborhoods, facilitates the exploration of novel phenotypes and thus phenotypic innovation.[7]

Genotype space has many additional features for which we currently have no statistical mode, no null hypothesis that would describe these features as random. They include the fact that the phenotypes of many neighbors of a genotype *G* are similar to that of *G*, whereas a few are very different phenotypes, that the number of mutations necessary to transform a phenotype into a completely dissimilar phenotype can be large (and system dependent), that the distribution of the number of genotypes with a given phenotype can be quite different for different phenotypes, and so on. Moreover, each of these features can depend on the system class—metabolism, regulatory circuits, or macromolecules (Lipman and Wilbur 1991; Schuster et al. 1994; Reidys, Stadler, and Schuster 1997; Fontana and Schuster 1998; Ciliberti et al. 2007a, 2007b; Rodrigues and Wagner 2009, 2011; Ferrada and Wagner 2010; Samal et al. 2010).

In sum, the random expectation of genotype space organization I mentioned has not been replaced by an unequivocally better model of this space. Until that time, one can say that mutations affect complex phenotypes nonrandomly. And more importantly, they do so in ways that facilitate evolutionary adaptation and innovation. The reason is that genotype networks and their diverse neighborhoods allow the exploration of many different neighborhoods of genotype space and the myriad new phenotypes that occur therein. These features help provide phenotypic raw material for natural selection, raw material that contains not only inferior phenotypes but also phenotypes that may be superior to a current, well-adapted phenotype. Because the vast majority of new phenotypes in the neighborhood of a well-adapted genotype *G* are worse than the phenotype of *G* itself, this ability to explore many new phenotypes can be

7. As in the section on DNA change, one can look beyond this null hypothesis toward more sophisticated models of how phenotypes are organized in genotype space. One model is worth mentioning in this regard. It revolves around the fraction of a genotype's neighbors that have the same phenotype as itself and can also be referred to as a genotype's robustness to point mutations. Such robustness can be shown to be both necessary and sufficient for the existence of genotype networks. However, this model of genotype space organization fails to capture other important aspects of genotype space organization, some of which are discussed in the next paragraph. At present, we do not know how to build an adequate such model, which would have to explain all known features of phenotype organization in genotype space and do so for systems as different as metabolism, regulatory circuits, and macromolecules.

very important to find superior phenotypes. Put more starkly, if genotype networks did not exist, then phenotypic improvement through natural selection might be impossible, because it might be prohibitive to find rare superior phenotypes in a vast genotype space without destroying present, well-adapted phenotypes.

In section 3, I stated that highly nonrandom—always beneficial—effects of mutations on fitness could imply that natural selection would become unnecessary. In contrast, the kind of nonrandomness I describe here for complex phenotypes does not obviate the need for natural selection. Instead, it makes evolutionary adaptation through natural selection easier or feasible. The relationship between natural selection and the nonrandom phenotypic change I describe here can be briefly described as follows. On the one hand, natural selection without genotype networks might not allow adaptive evolution, let alone the gradual emergence of complex phenotypes through successive adaptations. On the other hand, genotype networks without natural selection would be useless for adaptation, because natural selection is necessary to preserve existing well-adapted genotypes and to preserve new and superior phenotypes once they have been "discovered." Both natural selection and the nonrandom organization of complex phenotypes in genotype space are essential for organic evolution. Together they can help us understand how immensely complex phenotypes can arise through multiple cumulative adaptations.

10. Summary. To make the colloquial notion of randomness in Darwinian evolution more precise, I used the concept of sample spaces, events, and probabilities from probability theory. In probability theory, every event or process can be viewed as random; more colloquially, randomness means a conformity of observations to a prior expectation, a null hypothesis, or a statistical model. This second notion is useful because it allows us to distinguish between random and nonrandom aspects of the world. I discussed three very different system classes: metabolic networks, regulatory circuits, and macromolecules. These systems have complex phenotypes involved in many evolutionary innovations. In all three of them, we can study the relationship between genotype and phenotype systematically. All three systems show highly intertwined and connected genotype networks with diverse phenotypic neighborhoods. Such networks allow the exploration of novel phenotypes while preserving existing phenotypes and thus facilitate evolutionary adaptation and innovation. No statistical model or null hypothesis that could reproduce all or most aspects of genotype space organization currently exists. On the basis of what we know today, mutations affect both DNA and also complex phenotypes nonrandomly. Especially remarkable about their effects on complex phenotypes is that they facilitate Darwinian evolution.

REFERENCES

Beatty, J. 1984. "Chance and Natural Selection." *Philosophy of Science* 51:183–211.

Bharathan, G., T. E. Goliber, C. Moore, S. Kessler, T. Pham, and N. R. Sinha. 2002. "Homologies in Leaf Form Inferred from Knoxi Gene Expression during Development." *Science* 296:1858–60.

Bollobas, B., Y. Kohayakawa, and T. Luczak. 1992. "The Evolution of Random Subgraphs of the Cube." *Random Structures and Algorithms* 3:55–90.

Brakefield, P. M. 2006. "Evo-Devo and Constraints on Selection." *Trends in Ecology and Evolution* 21:362–68.

Cairns, J., J. Overbaugh, and S. Miller. 1988. "The Origin of Mutants." *Nature* 335:142–45.

Carroll, S. B., J. K. Grenier, and S. D. Weatherbee. 2001. *From DNA to Diversity: Molecular Genetics and the Evolution of Animal Design.* Malden, MA: Blackwell.

Causier, B., R. Castillo, J. L. Zhou, R. Ingram, Y. B. Xue, Z. Schwarz-Sommer, and B. Davies. 2005. "Evolution in Action: Following Function in Duplicated Floral Homeotic Genes." *Current Biology* 15:1508–12.

Chaitin, G. J. 1975. "Randomness and Mathematical Proof." *Scientific American* 5:47–52.

———. 2001. *Exploring Randomness*. London: Springer.

Chen, L. B., A. L. DeVries, and C. H. C. Cheng. 1997. "Convergent Evolution of Antifreeze Glycoproteins in Antarctic Notothenioid Fish and Arctic Cod." *Proceedings of the National Academy of Sciences of the United States of America* 94:3817–22.

Cheng, C. C.-H. 1998. "Evolution of the Diverse Antifreeze Proteins." *Current Opinion in Genetics and Development* 8:715–20.

Cheverud, J. M. 1984. "Quantitative Genetics and Developmental Constraints on Evolution by Selection." *Journal of Theoretical Biology* 110:155–71.

Ciliberti, S., O. C. Martin, and A. Wagner. 2007a. "Circuit Topology and the Evolution of Robustness in Complex Regulatory Gene Networks." *PLoS Computational Biology* 3 (2): e15.

———. 2007b. "Innovation and Robustness in Complex Regulatory Gene Networks." *Proceedings of the National Academy of Sciences of the United States of America* 104:13591–96.

Cline, R. E., R. H. Hill, D. L. Phillips, and L. L. Needham. 1989. "Pentachlorophenol Measurements in Body-Fluids of People in Log Homes and Workplaces." *Archives of Environmental Contamination and Toxicology* 18:475–81.

Copley, S. D. 2000. "Evolution of a Metabolic Pathway for Degradation of a Toxic Xenobiotic: The Patchwork Approach." *Trends in Biochemical Sciences* 25:261–65.

Dawkins, R. 1996. *Climbing Mount Improbable*. New York: Norton.

Denver, D. R., et al. 2009. "A Genome-Wide View of Caenorhabditis Elegans Base-Substitution Mutation Processes." *Proceedings of the National Academy of Sciences of the United States of America* 106:16310–14.

Drake, J. W., B. Charlesworth, D. Charlesworth, and J. F. Crow. 1998. "Rates of Spontaneous Mutation." *Genetics* 148:1667–86.

Earman, J. 1986. *A Primer on Determinism*. Dordrecht: Reidel.

Eble, G. J. 1999. "On the Dual Nature of Chance in Evolutionary Biology and Paleobiology." *Paleobiology* 25:75–87.

Elbert, T., W. J. Ray, Z. J. Kowalik, J. E. Skinner, K. E. Graf, and N. Birbaumer. 1994. "Chaos and Physiology—Deterministic Chaos in Excitable Cell Assemblies." *Physiological Reviews* 74:1–47.

Erdös, P., and A. Renyi. 1960. "On the Evolution of Random Graphs." *Evolution* 5:17–61.

Feller, W. 1968. *An Introduction to Probability Theory and Its Applications*. New York: Wiley.

Ferrada, E., and A. Wagner. 2008. "Protein Robustness Promotes Evolutionary Innovations on Large Evolutionary Time Scales." *Proceedings of the Royal Society of London* B 275:1595–1602.

———. 2010. "Evolutionary Innovation and the Organization of Protein Functions in Sequence Space." *PLoS ONE* 5 (11): e14172.

Fontana, W., and P. Schuster. 1998. "Continuity in Evolution: On the Nature of Transitions." *Science* 280:1451–55.

Futuyma, D. J. 1998. *Evolutionary Biology*. 3rd ed. Sunderland, MA: Sinauer.

Gavrilets, S. 1997. "Evolution and Speciation on Holey Adaptive Landscapes." *Trends in Ecology and Evolution* 12:307–12.

Gilbert, S. F. 1997. *Developmental Biology*. 5th ed. Sunderland, MA: Sinauer.

Glass, L. 2001. "Synchronization and Rhythmic Processes in Physiology." *Nature* 410:277–84.

Golding, G. B., and A. M. Dean. 1998. "The Structural Basis of Molecular Adaptation." *Molecular Biology and Evolution* 15:355–69.

Gould, S. J. 1991. "The Disparity of the Burgess Shale Arthropod Fauna and the Limits of Cladistic-Analysis—Why We Must Strive to Quantify Morphospace." *Paleobiology* 17:411–23.

———. 1993. *Eight Little Piggies: Reflections in Natural History*. New York: Norton.

Griffiths, P. E., and E. M. Neumann-Held. 1999. "The Many Faces of the Gene." *BioScience* 49:656–62.

Griffiths, P. E., and K. Stotz. 2007. "Gene." In *Cambridge Companion to the Philosophy of Biology*, ed. D. Hull and M. Ruse, 85–102. New York: Cambridge University Press.

Gruner, W., R. Giegerich, D. Strothmann, C. Reidys, J. Weber, I. L. Hofacker, P. F. Stadler, and P. Schuster. 1996. "Analysis of RNA Sequence Structure Maps by Exhaustive Enumeration." Pt. 2, "Structures of Neutral Networks and Shape Space Covering." *Monatshefte fur Chemie* 127:375–89.

Hartl, D. L., and A. G. Clark. 2007. *Principles of Population Genetics*. 4th ed. Sunderland, MA: Sinauer.

Hastings, A., C. L. Hom, S. Ellner, P. Turchin, and H. C. J. Godfray. 1993. "Chaos in Ecology—Is Mother Nature a Strange Attractor?" *Annual Review of Ecology and Systematics* 24:1–33.

Hendrickson, H., E. S. Slechta, U. Bergthorsson, D. I. Andersson, and J. R. Roth. 2002. "Amplification-Mutagenesis: Evidence That 'Directed' Adaptive Mutation and General Hypermutability Result from Growth with a Selected Gene Amplification." *Proceedings of the National Academy of Sciences of the United States of America* 99:2164–69.

Hodin, J. 2000. "Plasticity and Constraints in Development and Evolution." Paper presented at the Modularity of Animal Form Workshop, Friday Harbor, WA.

Hueber, S. D., and I. Lohmann. 2008. "Shaping Segments: Hox Gene Function in the Genomic Age." *Bioessays* 30:965–79.

Hughes, C. L., and T. C. Kaufman. 2002. "Hox Genes and the Evolution of the Arthropod Body Plan." *Evolution and Development* 4:459–99.

Hull, D. L. 1972. "Reduction in Genetics—Biology or Philosophy." *Philosophy of Science* 39:491–99.

———. 1979. "Reduction in Genetics." *Philosophy of Science* 46:316–20.

Irish, V. F. 2003. "The Evolution of Floral Homeotic Gene Function." *Bioessays* 25:637–46.

Jia, W. L., and P. G. Higgs. 2008. "Codon Usage in Mitochondrial Genomes: Distinguishing Context-Dependent Mutation from Translational Selection." *Molecular Biology and Evolution* 25:339–51.

Jörg, T., O. C. Martin, and A. Wagner. 2008. "Neutral Network Sizes of Biological RNA Molecules Can Be Computed and Are Not Atypically Small." *BMC Bioinformatics* 9: 464.

Karlin, S., and V. Brendel. 1993. "Patchiness and Correlations in DNA Sequences." *Science* 259:677–80.

Karlin, S., and L. R. Cardon. 1994. "Computational DNA Sequence Analysis." *Annual Review of Microbiology* 48:619–54.

Laubichler, M. D., G. B. Muller, W. Fontana, and G. P. Wagner. 2005. "Sacrificing Dialogue for Politics." *Science* 309:324.

Lemons, D., and W. McGinnis. 2006. "Genomic Evolution of Hox Gene Clusters." *Science* 313:1918–22.

Levitt, M. 2009. "Nature of the Protein Universe." *Proceedings of the National Academy of Sciences of the United States of America* 106:11079–84.
Li, W.-H. 1997. *Molecular Evolution*. Sunderland, MA: Sinauer.
Liang, Y. H., Z. Q. Hua, X. Liang, Q. Xu, and G. Y. Lu. 2001. "The Crystal Structure of Bar-Headed Goose Hemoglobin in Deoxy Form: The Allosteric Mechanism of a Hemoglobin Species with High Oxygen Affinity." *Journal of Molecular Biology* 313:123–37.
Lipman, D. J., and W. J. Wilbur. 1991. "Modeling Neutral and Selective Evolution of Protein Folding." *Proceedings of the Royal Society of London* B 245:7–11.
MacCarthy, T., R. Seymour, and A. Pomiankowski. 2003. "The Evolutionary Potential of the Drosophila Sex Determination Gene Network." *Journal of Theoretical Biology* 225: 461–68.
Mackey, M. C., and L. Glass. 1977. "Oscillation and Chaos in Physiological Control Systems." *Science* 197:287–88.
May, R. M. 1976. "Simple Mathematical Models with Very Complicated Dynamics." *Nature* 261:459–67.
Maynard-Smith, J. 1970. "Natural Selection and the Concept of a Protein Space." *Nature* 255:563–64.
Maynard-Smith, J., R. Burian, S. Kauffman, P. Alberch, J. Campbell, B. Goodwin, R. Lande, D. Raup, and L. Wolpert. 1985. "Developmental Constraints and Evolution." *Quarterly Review of Biology* 60:265–87.
Mayr, E. 1961. "Cause and Effect in Biology." *Science* 134:1501–6.
Millstein, R. L. 2000. "Chance and Macroevolution." *Philosophy of Science* 67:603–24.
———. 2002. "Are Random Drift and Natural Selection Conceptually Distinct?" *Biology and Philosophy* 17:33–53.
Morton, B. R. 2003. "The Role of Context-Dependent Mutations in Generating Compositional and Codon Usage Bias in Grass Chloroplast DNA." *Journal of Molecular Evolution* 56:616–29.
Newman, S. A., and R. Bhat. 2009. "Dynamical Patterning Modules: A 'Pattern Language' for Development and Evolution of Multicellular Form." *International Journal of Developmental Biology* 53:693–705.
Niu, D. K., K. Lin, and D. Y. Zhang. 2003. "Strand Compositional Asymmetries of Nuclear DNA in Eukaryotes." *Journal of Molecular Evolution* 57:325–34.
Odell, G. M., G. Oster, P. Alberch, and B. Burnside. 1981. "The Mechanical Basis of Morphogenesis." Pt. 1, "Epithelial Folding and Invagination." *Developmental Biology* 85:446–62.
Omilian, A. R., M. E. A. Cristescu, J. L. Dudycha, and M. Lynch. 2006. "Ameiotic Recombination in Asexual Lineages of Daphnia." *Proceedings of the National Academy of Sciences of the United States of America* 103:18638–43.
Ossowski, S., K. Schneeberger, J. I. Lucas-Lledo, N. Warthmann, R. M. Clark, R. G. Shaw, D. Weigel, and M. Lynch. 2009. "The Rate and Molecular Spectrum of Spontaneous Mutations in Arabidopsis Thaliana." *Science* 327:92–94.
Oster, G. F., N. Shubin, J. D. Murray, and P. Alberch. 1988. "Evolution and Morphogenetic Rules—the Shape of the Vertebrate Limb in Ontogeny and Phylogeny." *Evolution* 42: 862–84.
Price, N. D., J. A. Papin, C. H. Schilling, and B. O. Palsson. 2003. "Genome-Scale Microbial in Silico Models: The Constraints-Based Approach." *Trends in Biotechnology* 21:162–69.
Rehmann, L., and A. J. Daugulis. 2008. "Enhancement of PCB Degradation by *Burkholderia xenovorans* LB400 in Biphasic Systems by Manipulating Culture Conditions." *Biotechnology and Bioengineering* 99:521–28.
Reidhaar-Olson, J. F., and R. T. Sauer. 1990. "Functionally Acceptable Substitutions in 2 Alpha-Helical Regions of Lambda Repressor." *Proteins: Structure, Function, and Genetics* 7:306–16.
Reidys, C. M., and P. F. Stadler. 2002. "Combinatorial Landscapes." *SIAM Review* 44:3–54.
Reidys, C. M., P. F. Stadler, and P. Schuster. 1997. "Generic Properties of Combinatory

Maps: Neutral Networks of RNA Secondary Structures." *Bulletin of Mathematical Biology* 59:339–97.

Rodrigues, J. F., and A. Wagner. 2009. "Evolutionary Plasticity and Innovations in Complex Metabolic Reaction Networks." *PLoS Computational Biology* 5:e1000613.

———. 2011. "Genotype Networks in Sulfur Metabolism." *BMC Systems Biology* 5:39.

Samal, A., J. F. M. Rodrigues, J. Jost, O. C. Martin, and A. Wagner. 2010. "Genotype Networks in Metabolic Reaction Spaces." *BMC Systems Biology* 4:30.

Schonborn, C. 2005. "Finding Design in Nature." *New York Times*, July 7.

Schuster, P., W. Fontana, P. Stadler, and I. Hofacker. 1994. "From Sequences to Shapes and Back—a Case-Study in RNA Secondary Structures." *Proceedings of the Royal Society of London* B 255:279–84.

Shubin, N. H., and P. Alberch. 1986. "A Morphogenetic Approach to the Origin and Basic Organization of the Tetrapod Limb." *Evolutionary Biology* 20:319–87.

Simpson, G. G. 1953. *The Major Features of Evolution*. New York: Columbia University Press.

Sober, E. 1984. *The Nature of Selection.* Cambridge, MA: MIT Press.

———. 2000. *Philosophy of Biology*. Boulder, CO: Westview.

Sokal, R. R., and F. J. Rohlf. 1981. *Biometry*. New York: Freeman.

Sumedha, O., C. Martin, and A. Wagner. 2007. "New Structural Variation in Evolutionary Searches of RNA Neutral Networks." *Biosystems* 90:475–85.

Takiguchi, M., T. Matsubasa, Y. Amaya, and M. Mori. 1989. "Evolutionary Aspects of Urea Cycle Enzyme Genes." *Bioessays* 10:163–66.

Todd, A. E., C. A. Orengo, and J. M. Thornton. 2001. "Evolution of Function in Protein Superfamilies, from a Structural Perspective." *Journal of Molecular Biology* 307:1113–43.

Touchon, M., and E. P. C. Rocha. 2008. "From Gc Skews to Wavelets: A Gentle Guide to the Analysis of Compositional Asymmetries in Genomic Data." *Biochimie* 90:648–59.

van der Meer, J. R. 1995. "Evolution of Novel Metabolic Pathways for the Degradation of Chloroaromatic Compounds." Paper presented at the Beijerinck Centennial Symposium on Microbial Physiology and Gene Regulation—Emerging Principles and Applications, The Hague, December.

van der Meer, J. R., C. Werlen, S. F. Nishino, and J. C. Spain. 1998. "Evolution of a Pathway for Chlorobenzene Metabolism Leads to Natural Attenuation in Contaminated Groundwater." *Applied and Environmental Microbiology* 64:4185–93.

Wagner, A. 2005. "Circuit Topology and the Evolution of Robustness in Two-Gene Circadian Oscillators." *Proceedings of the National Academy of Sciences of the United States of America* 102:11775–80.

———. 2011. *The Origins of Evolutionary Innovations: A Theory of Transformative Change in Living Systems.* Oxford: Oxford University Press.

Wagner, G. P., C. Amemiya, and F. Ruddle. 2003. "Hox Cluster Duplications and the Opportunity for Evolutionary Novelties." *Proceedings of the National Academy of Sciences of the United States of America* 100:14603–6.

Wimsatt, W. C. 1980. "Randomness and Perceived Randomness in Evolutionary Biology." *Synthese* 43:287–329.