

Generative models versus underlying symmetries to explain biological pattern

S. A. FRANK

Department of Ecology and Evolutionary Biology, University of California, Irvine, CA, USA

Keywords:

evolutionary genetics;
extreme value theory;
limiting distributions;
mathematical models;
systems biology;
theoretical biology.

Abstract

Mathematical models play an increasingly important role in the interpretation of biological experiments. Studies often present a model that generates the observations, connecting hypothesized process to an observed pattern. Such generative models confirm the plausibility of an explanation and make testable hypotheses for further experiments. However, studies rarely consider the broad family of alternative models that match the same observed pattern. The symmetries that define the broad class of matching models are in fact the only aspects of information truly revealed by observed pattern. Commonly observed patterns derive from simple underlying symmetries. This article illustrates the problem by showing the symmetry associated with the observed rate of increase in fitness in a constant environment. That underlying symmetry reveals how each particular generative model defines a single example within the broad class of matching models. Further progress on the relation between pattern and process requires deeper consideration of the underlying symmetries.

Introduction

In a recent experiment, bacterial fitness increased steadily over time (Wiser *et al.*, 2013). The logarithm of the rate of increase in fitness declined approximately linearly with respect to the logarithm of time. Such a simple pattern contains information about the underlying processes that determine how fitness increases. But exactly what sort of information?

To evaluate the match between an observed pattern and a hypothesized process, mathematical models have become the standard in biology. Typically, one puts together a set of plausible assumptions about process and then studies the resulting model for how well it generates the target outcome. A successful match implies a plausible generative model of process. Ideally, the generative model will make additional testable hypotheses, which can be studied in further experiments.

But does a successful generative model, by itself, really provide much information about underlying

process? Probably not. The more commonly a pattern is observed, the more important it is to understand the underlying process. At the same time, it is almost always true that the more common a pattern, the greater the number of underlying generative models that match the pattern. The simple law of nature is that the commonness of a pattern associates with the number of distinctive underlying processes that lead to that pattern (Jaynes, 2003). Put another way, it is overwhelmingly easy to make a generative model that matches a simple, common pattern, but the match provides little information about the true underlying process (Frank, 2009).

What is the real information in a pattern about underlying process? The real information is the constraints that define the underlying class of matching generative models (Jaynes, 2003). Such constraints express the fundamental symmetries that determine pattern (Feynman, 1967; Anderson, 1972). In this case, symmetry means the group of alternative models that produce the same pattern. Symmetry defines alternative states that are the same with regard to some measure (Weyl, 1983). For example, a square is the same after it is rotated by ninety degrees, so a square is symmetric with respect to ninety degree rotations. Similarly, one may change many assumptions of an

Correspondence: Steven A. Frank, Department of Ecology and Evolutionary Biology, University of California, Irvine, CA 92697–2525, USA. Tel.: +1 949 824 2244; fax: +1 949 824 2181; e-mail: safrank@uci.edu

evolutionary model and still generate a linear decline in the logarithm of fitness increase with respect to the logarithm of time. The rate of increase in fitness is symmetric with respect to such changes in underlying generative assumptions.

If we knew all of the symmetries with respect to the rate of increase in fitness, then we would know the full class of underlying generative models consistent with that pattern. We would know the essence of process required to generate the pattern, and those aspects of particular generative models that do not matter. For simple patterns, the associated symmetries tend to be simple, and most details of particular generative models do not matter. However, we can only distinguish those aspects that matter from those that do not if we know the defining symmetries.

Why are most studies limited to expressing the existence of a matching generative model? Because finding a matching model is easy, whereas discovering the underlying symmetries that truly define the relations between pattern and process is often difficult. In consequence, it has become common to ignore the inherent structure of the problem, and to be satisfied with a generative match. Indeed, the very idea that one should look for underlying symmetries rather than a generative match is rarely acknowledged and perhaps not widely understood.

I do not have a solution to the difficulty of discovering the underlying symmetries. However, to have a chance, one must first recognize the problem. With the proper goal in mind, certain steps often help in learning about the underlying symmetries of process that lead to observed pattern. In this article, I use the example of increasing bacterial fitness to illustrate some of these issues.

Overview

The argument in this article is a bit more abstract than usual. A brief overview may help before starting. A common sequence of science is hypothesis, test by observation and updated hypothesis in light of observed pattern. How well one does with that sequence depends on how good one is at finding useful updated hypotheses. An updated hypothesis must, of course, be constrained by the observed pattern. But consistency with observed pattern may often be, by itself, a rather weak way of generating new hypotheses.

For example, if one observed a Gaussian (Normal) distribution of measurements, then a detailed hypothesis about why the specific natural history or biochemistry of certain processes led to that Gaussian pattern could easily be constructed to generate the pattern. But such a detailed explanation of process would almost certainly be wrong, because the central limit theorem tells us that Gaussian patterns arise inevitably when underlying processes tend to add up to make final

pattern, irrespective of almost all of the details concerning the underlying processes.

The phrase 'irrespective of almost all of the details concerning the underlying processes' means that pattern in that case is symmetric or invariant to many changes in details. One must know that, otherwise the hypothetico-deductive process is almost certain to lead one to false paths, because one may mistakenly put too much weight on hypothesized details of process that in fact do not matter. The same problem arises for many patterns that are not Gaussian in shape. The difficulty is that the underlying symmetries are not always immediately obvious and so require some thought in order to avoid false paths. To illustrate those points, I have structured this article as follows.

The first section introduces a particular observed pattern for the increase in fitness over time in bacterial populations. I express that observed pattern in a variety of alternative ways. Those alternative expressions allow one to see the pattern from a variety of perspectives. Those different perspectives provide a deeper sense of the pattern and its 'shape'. A sense of shape helps to see what may matter and may not with regard to underlying process, that is, the underlying symmetries.

The second and third sections introduce the particular symmetries that may explain the form of the observed pattern for the increase in fitness. In this case, the symmetries arise from extreme value theory. That theory has the same structure as the central limit theorem. In particular, a wide variety of seemingly different processes turn out to lead to the same pattern, when the pattern depends primarily on rare or extreme events. At first glance, it may seem that rare events would be particularly difficult to predict and so be difficult to analyse with regard to the consequences for pattern. However, although each rare event is hard to predict, in the aggregate over several rare events, the outcomes tend to converge to a very regular form. One must recognize that symmetry to find meaningful hypotheses about the generation of certain types of pattern.

The fourth section applies extreme value theory to the observed pattern for the increase in fitness in bacterial populations. The final sections interpret the particular results with respect to the broader problem of understanding symmetries and their role in generating useful hypotheses for the interpretation of commonly observed patterns.

Observed power law scaling for the increase in fitness

Wiser *et al.* (2013) studied the change in fitness over time in experimental populations of bacteria maintained in a constant environment. They showed that the observed fitnesses over time follow a power law relation

$$w = (1 + bt)^a, \quad (1)$$

where w is mean fitness at time t , estimated from 12 replicate populations, and a and b are fitted parameters. For my purposes, it will be useful to express Wisser *et al.*'s relation in a variety of alternative ways to highlight the related forms of simplicity that define the essential pattern. Most of the simple forms begin with logarithmic scaling

$$\log w = a \log(1 + bt), \quad (2)$$

which, for $bt \gg 1$, approaches

$$\log w = a \log b + a \log t. \quad (3)$$

Rate of change in mean fitness

From eqn (3)

$$\frac{d \log w}{d \log t} = a, \quad (4)$$

which shows a very simple relation between the change in log mean fitness and the change in the logarithm of time.

We can also write the change in log fitness with respect to time as

$$\frac{d \log w}{dt} = \gamma \quad (5)$$

where $\gamma \equiv \gamma(t)$ is a function of time that describes the rate of increase in log mean fitness at any time. Note that if selection is the only force, then $\gamma \equiv J$ from Frank (2012), and γ is the increase in information accumulated by a population from the action of natural selection.

It follows from eqn (2) that the increase in information at any point in time is

$$\gamma = \frac{ab}{1 + bt} \rightarrow \frac{a}{t}, \quad (6)$$

where the limiting value on the right arises when $bt \gg 1$.

Decay in the rate of change in log fitness

From eqn (6), it follows that

$$\frac{d \log \gamma}{d \log t} = -1, \quad (7)$$

which shows that the logarithm of the rate of increase in fitness declines directly in proportion to $\log t$, the logarithm of the amount of time that has passed.

Time required for a fixed change in log fitness

We can write eqn (4) as

$$\frac{d \log w}{d \log t} = \gamma t = a, \quad (8)$$

because $dt = t d \log t$, and from eqn (6), $\gamma t = a$.

How much time must pass for a fixed change in log fitness, $d \log w = k$? From eqn (8),

$$d \log t = k/a. \quad (9)$$

Because $d \log t = dt/t$, we have

$$dt = kt/a \propto t, \quad (10)$$

where the symbol \propto means 'is proportional to'. Thus, the time increment, dt , needed to obtain a fixed change in log fitness is proportional to the amount of time, t , that has passed since the first measurement of fitness.

Interpretation of simple scaling relations

Observation of such simple and elegant scaling relations implies a powerful underlying force. That underlying force must erase all the details of selection and evolution that are particular to each population, exposing the minimal symmetry that constrains pattern. No particular generative or dynamical model can make a primary claim to explaining such simplicity. Rather, one must search for the way in which aggregation and the loss of the particular information in each population causes the simple underlying symmetry to dominate (Frank, 2009).

Extreme values and the increase in fitness

Several prior studies have analysed the accumulation of beneficial mutations in a constant environment. Fitness increases over time as each beneficial mutation gets added to a population by selection. The most interesting studies with regard to underlying symmetries emphasize extreme value theory (Gillespie, 1983, 1984, 1994; Orr, 2003, 2010; Beisel *et al.*, 2007; Joyce *et al.*, 2008). That theory describes the probability distribution for the maximum observed value in a sample.

For the case of increasing fitness, one can think of a sample as the mutations that occur during a fixed time period. The most beneficial mutation that spreads through a population during a particular period would often be the extreme value, which would determine the advance in fitness over that time increment. The overall rate of increase in fitness depends primarily on the sequence of extreme values over sequential time periods.

The prior studies used extreme value theory to show that many details of underlying models do not matter with regard to the distribution of the effects of beneficial mutations. Because the distribution of beneficial mutations influences the rate of increase in fitness, many details of underlying models do not matter with regard to the increase in fitness. Here, I build on the insight of those prior applications of extreme value theory. In particular, I use extreme value theory to expose

the underlying symmetries that may explain why the observed pattern of increase in fitness follows the simple scaling law observed by Wisner *et al.* (2013).

Extreme value theory

This section summarizes extreme value theory. I limit my presentation here to essential aspects that illustrate the concepts. Frank (2009) provides a full introduction and citations to comprehensive treatises on extreme value theory.

I begin with a few definitions. Write the probability that a random process, Y , takes on a value above the given constant y , as

$$P(Y > y) = \hat{F}(y).$$

This expression describes the probability of observing a value in the upper tail of the distribution, that is, a value greater than y . However, our primary interest concerns the maximum value, M_n , in a sample of size n , which we may also think of as the extreme value of a sample. The probability that M_n is less than y is the probability that none of the n samples is greater than y , which is

$$P(M_n < y) = [1 - \hat{F}(y)]^n. \quad (11)$$

This expression is the right idea, because it tells us about the probability that the extreme value in a sample is bounded by a particular upper value, y . However, there is a mathematical problem. As the sample size, n , becomes large, the probability that the maximum value is bounded by y becomes zero, as long as there is some chance that an observation could be above y . In a big enough sample, eventually some observation will be above the bound y . So this expression by itself is not very helpful with regard to general principles, because the expression depends on the sample size, n , which will vary in each particular study.

To obtain a useful general expression for extreme values, we need to find a form that does not depend on sample size. Put another way, we want an expression for the probability of observing a particular extreme value that is symmetric, or invariant, with respect to changes in sample size. If we can find that underlying symmetry, we can greatly expand our understanding of the general underlying principles that determine the distribution of extreme values.

The form of eqn (11) suggests that the following mathematical identity will help in finding an expression that is independent of sample size

$$\left[1 - \frac{\hat{F}(y)}{n}\right]^n \longrightarrow e^{-\hat{F}(y)}, \quad (12)$$

in which the right hand side does not depend on the sample size, n , and is an increasingly good approximation

for the left hand side as $\hat{F}(y)/n$ decreases. The problem now becomes how to express the probability for the maximum in a sample, given in eqn (11), in the form given by eqn (12). In particular, the problem is that the maximum of the sample, M_n , depends on the sample size, n . How can we standardize M_n to remove the dependence on sample size? One possibility is to account for the increase in the expected maximum value with n by using the value for the upper bound $\alpha_n y + \beta_n$, in which the coefficients α_n and β_n depend on n in a way that accounts for the expected increase in the upper bound. In particular, if we can find values of α_n and β_n such that

$$\hat{F}(\alpha_n y + \beta_n) = \frac{\hat{F}(y)}{n},$$

then we can rewrite the form in eqn (11) as

$$P(M_n < \alpha_n y + \beta_n) = \left[1 - \frac{\hat{F}(y)}{n}\right]^n \longrightarrow e^{-\hat{F}(y)}. \quad (13)$$

The probability of obtaining a particular extreme value depends on how $\hat{F}(y)$, the upper tail probability of observing a value above y , declines with an increase in y . One commonly observed pattern is a power law decline in the upper tail probability with increasing y , such that the expected probability in the upper tail is approximately proportional to $\hat{F}(y) = y^{-1/a}$.

Given a power law scaling in the upper tail, we must choose $\alpha_n = n^a$ and $\beta_n = 0$ to satisfy eqn (13), because $\hat{F}(n^a y) = \hat{F}(y)/n$. Those constants lead to the extreme value distribution of the Fréchet form

$$P(M_n < n^a y) = e^{-y^{-1/a}}.$$

For our purposes, we only need to use the fact that, for a sufficiently small $\hat{F}(y)/n$, the expected value of the maximum, M_n , increases with the sample size n in proportion to n^a . Using angle brackets to denote expected values, we can express this key fact in symbols as

$$\langle M_n \rangle = n^a \langle y \rangle, \quad (14)$$

under the assumption that $a < 1$, so that M_n increases at a diminishing rate with n .

The increase in fitness

Our goal is to understand the observed power law scaling between relative fitness and time. That scaling was given in eqn (1), repeated here for convenience

$$w = (1 + bt)^a. \quad (15)$$

My claim is that eqn (14) captures the essence of the observed power law scaling for fitness. To support that claim, we must fill in some gaps to show how the extreme value result in eqn (15) expresses the simple underlying symmetries that lead to eqn (15).

To start, we must show the essential equivalence of the extreme value theory expression in eqn (14) and the observed scaling of fitness with time in eqn (15). Then, we must express the underlying symmetries in the extreme value theory and how those symmetries may clarify the types of generative processes that could lead to the observed pattern.

To establish the equivalence, we start by letting $\langle M_n \rangle \equiv w$, where w in this context is the expected relative fitness. The idea is that the best potentially selected mutant, M_n , has spread through the population and is the maximum observation, or extreme value. Other mutants may have had a higher value but, for whatever reason, were constrained from spreading and do not enter into the set of potentially selected mutants.

The expression $\langle M_n \rangle$ is the expected maximum value as a function of the sample size, n . We will interpret the sample size as an increasing function of the time that has passed. The relative fitness, w , at any time, is equivalent to the best mutant that has spread through the population up to that time.

The equivalence between fitness and the extreme value means that

$$w = \langle M_n \rangle = n^a \langle y \rangle.$$

Because w is relative fitness, which we may scale by any constant, we may choose $\langle y \rangle \equiv 1$ without loss of generality. We thus obtain the expression for the expected value of relative fitness

$$w = n^a. \quad (16)$$

This very simple result requires only that the probability distribution for potentially selected mutants has an upper tail that decays as a power law. A power law decay in the upper tail for fitness, w , is equivalent to an exponentially decaying upper tail when values for the mutants are expressed on a logarithmic scale of fitness, $\log w$. The actual shape of an upper tail is a difficult empirical problem (Clauset *et al.*, 2009). A power law may be a good approximation, because fitness has a natural tendency to scale logarithmically (Wagner, 2010). Logarithmic scaling typically corresponds to power laws (Frank, 2009). For specific discussion of alternative tail shapes of mutational effects in relation to extreme value theory, see the excellent analysis in Joyce *et al.* (2008).

The result in eqn (16) describes fitness in terms of sample size, n . However, the observed pattern of fitness in eqn (15) expresses the change in fitness with respect to time, t , rather than with respect to sample size, n . Thus, we must define a reasonable relation between sample size, n , and time, t . The relation that transforms eqn (16) into eqn (15) is

$$n = 1 + bt. \quad (17)$$

This equivalence works mathematically, but can we justify it biologically? In fact, the equivalence is

inevitable if we take the simplest interpretation of two necessary relations. First, at $t = 0$, the biologically observed expression in eqn (15) uses the arbitrary assumption that relative fitness is one. Using that same assumption, we must have $t = 0$ corresponding to $n = 1$, for which eqn (17) uses the simplest expression of that assumption. Second, if we follow the inevitable fact that sample size increases with time, then the simplest assumption is that the number of samples increases linearly with time. In eqn (17), the assumption is that the sample size, n , increases linearly with time as bt . For these two reasons, eqn (17) seems the simplest, fully justified way to express the relation between sample size and time.

Using the relation between sample size and time from eqn (17) in the universal extreme value scaling law in eqn (16), we obtain the observed pattern for the change in fitness with time in eqn (15). The only assumptions are upper tail events are relatively rare; upper tail events that increase fitness decay as a power law; and sample size increases linearly with time. The observed power law pattern contains only the information in those three assumptions. For any generative process that matches the observed pattern, no additional information can be added beyond those key assumptions.

Symmetry interpretation versus generative models

For the observed power law increase in fitness with time in Wisser *et al.* (2013), all matching generative models are symmetric with regard to details beyond the two key aspects of information contained in the observed pattern. To repeat the above conclusion, the key aspects of information are upper tail events that increase fitness decay as a power law and upper tail events are relatively rare (Gillespie, 1994; Joyce *et al.*, 2008; Orr, 2010). Here, symmetry means invariance, in the sense that the pattern generated by a matching model is invariant with regard to any details beyond the two components of information contained in the observed pattern.

The central limit theorem provides the best known example of symmetry and invariance (Jaynes, 2003). When summing up a series of random processes, one often observes a normal, or Gaussian, distribution. The reason is that the particular details in each component process tend to average out, exposing only the underlying information about the mean and variance of the aggregate process. Thus, the mean and variance are sufficient statistics to define a Gaussian distribution. Put another way, given the mean and the variance, all other details of a particular generative process get washed away in the aggregate. The Gaussian pattern is symmetric, or invariant, to the other particular details of the great many generative processes that lead to that

outcome. Indeed, the Gaussian is so very common because there are so many different generative processes that satisfy the key underlying symmetry.

The central limit theorem is well known. Thus, an observed Gaussian pattern rarely motivates anyone to produce a detailed generative model to match that pattern. The fact that so many different generative models lead to that pattern is understood. By contrast, in biology, almost any other observed pattern often motivates the production of a detailed generative model. But many other patterns besides the Gaussian share the essential feature of attracting a wide class of underlying generative processes to the same pattern. The extreme value pattern is one of the most powerful attractors, on the same level as the Gaussian pattern (Embrechts *et al.*, 1997; Kotz & Nadarajah, 2000; Frank, 2009; Frank & Smith, 2011). If the set of attracting patterns were better known, the widespread tendency of explaining simple general patterns by overly particular and often misleading generative models would not be so common (Gutenkunst *et al.*, 2007; Frank, 2011, 2013; Machta *et al.*, 2013).

The philosophical literature contains a distinct analysis of similar problems. As Stanford (2013) puts it: 'At the heart of the underdetermination of scientific theory by evidence is the simple idea that the evidence available to us at a given time may be insufficient to determine what beliefs we should hold in response to it'. Focus on symmetries provides a way forward, by telling us something about the particular nature of underdetermination.

Discussion

In this section, I first provide additional context about symmetries in relation to explanation. I then discuss how my approach to the particular example of fitness relates to past work.

Underlying symmetries do not explain all of the details of observed pattern. Rather, the symmetries express the main force that sets the broad features of pattern. In the same way, the average value over a range of outcomes does not express all of the variability, but rather the expected outcome or the general location of pattern. Thus, my argument is not that the details of particular models are necessarily uninteresting or unhelpful. Instead, one must first locate the primary cause of pattern before one can understand the causes of variation around that primary location. The first step in explanation is to figure what does not matter in setting the main shape of pattern. One may then fill in the details of how, in particular situations, additional processes set the variations around the primary shape.

Often, a complicated and detailed generative model, which appears to capture many aspects of realism, attracts strongly to the very simple pattern set by the basic underlying symmetries. The many details act like

random perturbations relative to the core pattern. The greater the number of random perturbations, the more likely they tend to cancel in the aggregate, leaving only the core pattern.

Wiser *et al.* (2013) present a detailed generative model to explain their observed power law scaling of fitness with time. They assume that beneficial mutations follow an exponential distribution when measured with respect to $\log w$, which is equivalent to a distribution with a power law upper tail for w . They also make many particular assumptions, leading to a relatively complicated model and analysis. Their primary conclusion is that their model matches the simple power law pattern for the increase in fitness with respect to time. They also include analysis and discussion of many potentially realistic and informative details of evolutionary process. Those details could provide much insight. The main limitation is that, by not expressing the underlying symmetry and the simplicity of the relation between pattern and process, it is hard to see what matters and what does not matter with respect to the primary pattern. Therefore, it is hard to know how the particular details do or do not influence variation around the primary trend set by the underlying symmetry.

Much past work has promoted the power of extreme value theory for analysing the distribution of beneficial mutations and the rate of increase in fitness (Gillespie, 1994; Joyce *et al.*, 2008; Orr, 2010). That work does emphasize the importance of underlying symmetries in understanding the relation between pattern and process. However, the particular analyses from the past work have sometimes focused on rather detailed and specific patterns or assumptions (Joyce *et al.*, 2008). Those details could, in theory, matter a lot when analysing the actual patterns of mutation, selection and the increase in fitness.

The great value of the data and analysis from Wiser *et al.* (2013) is the simplicity of observed pattern. My main goal has been to match that observed simplicity to a theory that captures the underlying symmetry, in a way that could explain the primary cause of such simplicity of pattern. Of course, other general models besides extreme value theory could match the same observed pattern. But that is exactly my point. Simple, common patterns tend to be attractors for many underlying models. Extreme value theory does have a privileged position, shared only with the central limit theorem and its generalization to a broad family of related patterns (Frank, 2009, 2011; Frank & Smith, 2011). Many common patterns arise from those few families of special attractors.

Acknowledgments

National Science Foundation (USA) Grants EF-0822399 and DEB-1251035 support my research.

References

- Anderson, P. 1972. More is different. *Science* **177**: 393–396.
- Beisel, C.J., Rokyta, D.R., Wichman, H.A. & Joyce, P. 2007. Testing the extreme value domain of attraction for distributions of beneficial fitness effects. *Genetics* **176**: 2441–2449.
- Clauset, A., Shalizi, C.R. & Newman M.E.J. 2009. Power-law distributions in empirical data. *SIAM Rev.* **51**: 661–703.
- Embrechts, P., Kluppelberg, C. & Mikosch, T. 1997. *Modeling Extremal Events: For Insurance and Finance*. Springer Verlag, Heidelberg.
- Feynman, R.P. 1967. *The Character of Physical Law*. MIT Press, Cambridge, MA.
- Frank, S.A. 2009. The common patterns of nature. *J. Evol. Biol.* **22**: 1563–1585.
- Frank, S.A. 2011. Measurement scale in maximum entropy models of species abundance. *J. Evol. Biol.* **24**: 485–496.
- Frank, S.A. 2012. Natural selection. V. How to read the fundamental equations of evolutionary change in terms of information theory. *J. Evol. Biol.* **25**: 2377–2396.
- Frank, S.A. 2013. Input-output relations in biological systems: measurement, information and the hill equation. *Biol. Direct* **8**: 31.
- Frank, S.A. & Smith, E. 2011. A simple derivation and classification of common probability distributions based on information symmetry and measurement scale. *J. Evol. Biol.* **24**: 469–484.
- Gillespie, J.H. 1983. A simple stochastic gene substitution model. *Theor. Popul. Biol.* **23**: 202–215.
- Gillespie, J.H. 1984. Molecular evolution over the mutational landscape. *Evolution* **38**: 1116–1129.
- Gillespie, J.H. 1994. *The Causes of Molecular Evolution*. Oxford University Press, New York, NY, USA.
- Gutenkunst, R.N., Waterfall, J.J., Casey, F.P., Brown, K.S., Myers, C.R. & Sethna, J.P. 2007. Universally sloppy parameter sensitivities in systems biology models. *PLoS Comput. Biol.* **3**: e189.
- Jaynes, E.T. 2003. *Probability Theory: The Logic of Science*. Cambridge University Press, New York.
- Joyce, P., Rokyta, D.R., Beisel, C.J. & Orr, H.A. 2008. A general extreme value theory model for the adaptation of dna sequences under strong selection and weak mutation. *Genetics* **180**: 1627–1643.
- Kotz, S. & Nadarajah, S. 2000. *Extreme Value Distributions: Theory and Applications*. World Scientific, Singapore.
- Machta, B.B., Chachra, R., Transtrum, M.K. & Sethna, J.P. 2013. Parameter space compression underlies emergent theories and predictive models. *Science* **342**: 604–607.
- Orr, H.A. 2003. The distribution of fitness effects among beneficial mutations. *Genetics* **163**: 1519–1526.
- Orr, H.A. 2010. The population genetics of beneficial mutations. *Phil. Trans. R. Soc. B: Biol. Sci.* **365**: 1195–1201.
- Stanford, K. 2013. Underdetermination of scientific theory. In: *The Stanford Encyclopedia of Philosophy Winter 2013 edn* (E.N. Zalta, ed.). Stanford University, Palo Alto, CA. Retrieved from <http://plato.stanford.edu/archives/win2013/entries/scientific-underdetermination/>
- Wagner, G.P. 2010. The measurement theory of fitness. *Evolution* **64**: 1358–1376.
- Weyl, H. 1983. *Symmetry*. Princeton University Press, Princeton, NJ.
- Wiser, M.J., Ribeck, N. & Lenski, R.E. 2013. Long-term dynamics of adaptation in asexual populations. *Science* **342**: 1364–1367.

Received 3 February 2014; revised 25 March 2014; accepted 30 March 2014