

# RANDOMNESS, CHAOS AND PREDICTION IN EVOLUTIONARY THEORY

---

CHARLES H. PENCE , *Université catholique de Louvain*

**Abstract.** We often hear that evolutionary theory tells us that the history of life has been directed by the whims of randomness, or even that ‘we are here by chance’. At the same time, evolutionary theorists often construct models of evolution that are taken to be predictive, and geneticists and molecular biologists occasionally offer us extremely accurate predictions of molecular-level evolutionary change. How should we understand this interaction between prediction and randomness? I will explore here one particular kind of prediction – predictions on the basis of quantitative estimates of fitness – in light of both the data that we need to draw those predictions and some recent mathematical work on the impact of chaos on evolutionary models, with the aim of examining what we might still be able to say about the predictability of the future of life in an evolving world.

**Keywords:** Fitness, Prediction, Natural selection, Inference, Chaos

## 1. Introduction

Critics of evolutionary theory, especially those of a creationist bent, often argue that evolution is a ‘random’ process. An analogy especially prominent in the conservative-Christian area where I grew up likens the development of complex traits by evolution to the possibility that a tornado passing through a junkyard would miraculously assemble a jet aeroplane. Of course, even the most superficial understanding of evolutionary theory makes it clear that natural selection is the very opposite of a random process – as Darwin himself put it, with more than a little anthropomorphisation, natural selection is ‘daily and hourly scrutinising, throughout the world, every variation, even the slightest; rejecting that which is

*Philosophy, Politics and Critique* 2.2, 2025, 157–169

DOI: 10.3366/ppc.2025.0075

© Edinburgh University Press 2025

[www.euppublishing.com/ppc](http://www.euppublishing.com/ppc)

bad, preserving and adding up all that is good; silently and insensibly working, wherever and whenever opportunity offers, at the improvement of each organic being'.<sup>1</sup>

But to say that some of these refined products of evolution – that is, evolutionary adaptations – are non-random is a far cry from saying that *the course of evolution itself* is non-random, or, in different terms that will be my focus here, that the process of evolution is *predictable*. As is well known, mutation, which provides the 'raw material' on which selection acts, is a paradigmatically random process, and evolution is not an optimizing, omniscient designer, but rather a tinkerer that 'makes do' with the material at hand in any given biological scenario.<sup>2</sup> This and other features of evolution, are often cited to bolster the idea that evolution, whether at the micro- or the macro-level, is an unpredictable phenomenon.

At the same time, biologists regularly produce mathematical models of evolutionary change that offer us predictions. Minimally, quantitative calculations of the evolutionary fitness of traits – of their expected growth rates in the future, for instance – seem to be useful in predicting evolutionary outcomes, a sense of fitness that has been baptized 'predictive fitness' (or 'mathematical fitness') in the philosophy of biology literature.<sup>3</sup> An implicit contrast is to be drawn here with 'vernacular' or 'ecological' fitness, which captures the more general idea that organisms with traits that 'fit' the environment better will go on to greater success. While this latter notion is difficult to formulate precisely enough to provide clear foresight (after all, there are many ways to 'fit' with the environment), predictive fitness is taken to offer us quantified, expected future changes in traits or genes. In a more complex vein, work on the genetic basis of evolution has been used to predict future sites of evolutionary change, in an intricate interaction between genetic regulation, the biochemical function of particular genes, and population genetics.<sup>4</sup>

This apparent tension over the status of evolutionary prediction has long been noted by philosophers of science. As early as the 1950s, Michael Scriven argued (in *Science* magazine, no less) that despite the fact that we might wish that evolutionary theory supported grand predictions, 'its great commitment and its profound illumination are to be found in its application to the lengthening past, not the distant future: in the tasks of explanation, not in those of prediction'.<sup>5</sup> Rasmus Winther, in opposition, has claimed that selection does, indeed, exhibit just the kind of prediction that we would expect from a successful scientific theory – that is, risky predictions of surprising, novel phenomena.<sup>6</sup>

How ought we square this circle? What is the relationship between evolutionary predictability and the evolutionary randomness that seems to threaten it? In this paper, I'll explore that connection by focusing on one class of evolutionary predictions, already introduced in outline above: What can we infer from the claim that one trait is fitter than another in a particular model of mathematical fitness? Put more colloquially, just how *predictive* is *predictive fitness*? Two lines of argument are relevant here. First, I'll consider what sort of data we might be able to obtain to derive our estimates of predictive fitness in various types of evolving populations.

Then, I'll present some challenging results from evolutionary modelling that cast doubt on the predictability of at least several kinds of evolutionary change.

As we will see, the overall appraisal is mixed: while prediction will be possible in a few cases, it seems as though evolution makes a number of situations predictively intractable, either because we can never collect enough data about them, or because they are too random, or more precisely too chaotic, to permit prediction. Evolutionary predictions, at least of this sort, will be reasonably rare and hard-won.

## **2. On Prediction**

What kind of prediction are we making when we say that, for instance, one trait has a higher fitness value than another and, thus, we predict that it will increase in frequency over time?<sup>7</sup> The outlines of a general answer to this question have been offered by Richard Lewontin. What it means, he says, to 'have an evolutionary perspective on a system' is to be 'interested in the change of state of some universe in time'.<sup>8</sup> What we want is information about the future transformation of the system: what is the function that tells us how we get from the system's current state to its future states? In Newtonian mechanics, for instance, this would be an equation like  $F = ma$ , with which we are all familiar from secondary-school physics: if we know about the relevant forces and masses, Newton's second law tells us what will happen to an object as we move forward in time.

In realistic cases for complex sciences such as evolution, we cannot pursue the precise form of this transformation (as we perhaps could if we were working with Newtonian forces). Rather, in Lewontin's terms, we set some 'tolerance limits' on system states, such that any two states that are similar enough to one another will be regarded as indistinguishable, making the problem more tractable.<sup>9</sup> And these tolerance limits are almost always a matter of debate in evolutionary biology. Do we just want relative predictions about one trait doing better or worse than another? Or do we want precise, numerical predictions? When we combine this with the practical difficulty of determining the real-world values of the evolutionary parameters that govern the behaviour of populations, the problem is made all the more difficult.<sup>10</sup>

Let's zoom in on the kinds of predictions that I want to target here: predictions about the future state of an evolving system based upon the mathematical fitness values of traits. How do these predictions work? One more observation from Lewontin is salient. In the final, emphasized claim with which he concludes his discussion of prediction, he notes that

the sufficient set of state variables for describing an evolutionary process within a population must include some information about the statistical distribution of genotypic frequencies. It is for this reason that the empirical study of population genetics has always begun with and centered around the characterization of the genetic variation in populations.<sup>11</sup>

Because the stock-in-trade of population genetics is the description of the distribution of characters in a population – the distribution of gene frequencies – then the predictions that population genetics makes will have to involve those same statistical distributions. Predicting with fitness, then, is about finding the right kind of estimation of future population behaviour on the basis of our current knowledge of the distribution of characters.

A variety of authors have argued that providing these kinds of predictions is the primary reason that a fitness concept is worth having in evolutionary theory at all.<sup>12</sup> We can even see mathematical fitness used as a medium-to-long-range predictor in biological works, such as those from the Long-Term Evolution Experiment run by Richard Lenski and colleagues over a period of decades. Their experimental population of *E. coli* bacteria began without the ability to use citrate as a carbon source, but after several thousand generations, one lineage evolved this capacity, and with it, a significant mathematical fitness advantage (measured, in this case, by an increase in growth rate). Lenski and colleagues then explored the contingency and path dependence of this result – confident, at any rate, that the evolution of citrate metabolism could now be predicted in the history of this and other lineages.<sup>13</sup>

If, as I will argue, the predictions that fitness in fact offers us are of mixed quality at best, then the primary role of fitness must not be predictive, *pace* these authors.<sup>14</sup> By extension, the theoretical arguments and practical orientation toward fitness that they exhibit must be mistaken. But we must start with a more general question: what do we need in order to be able to draw quality predictions from models in population genetics? I'll focus on two features here, each of which gives rise to the opportunity for predictive failure. First, we need to be in possession of a sufficient sample of the possible relationships between the trait and all the various environments and interactions which it might encounter. Second, the models themselves need to be structured such that these predictions are formally coherent. As we'll see, neither of these features is guaranteed.

### 3. Getting to Fitness Values

Individual organisms and their traits, as Elliott Sober poetically phrased the matter, 'taste of life but once'.<sup>15</sup> At the extreme, attempting to make a prediction on the basis of the particular life actualized by one organism gives us the smallest possible 'sample size': an organism will only realize precisely one of the myriad possible lives that it could have lived, the rest consigned to the dustbin of unexplored possibilities by a host of influences, many of which we might, not unjustly, call 'random'.<sup>16</sup> Such an inference is therefore extremely unlikely to accurately predict the 'real' mathematical fitness value of any of that organism's traits. In natural populations, we may be able to sample a number of similar, or even clonal, organisms,<sup>17</sup> raising this sample size slightly, but it is still likely to be an unrepresentative sample of the space of possible lives, leading once again to prediction from insufficient data. For instance, these natural populations will only likely encounter a small set of

environmental conditions and other organisms. It is unlikely that these will form a representative sample of the set of all such interactions.

If we turn to the predictive fitness values of traits in larger populations instead of estimating from smaller groups, the outlook gets a bit better. When we think about traits as forming classes or groups of individuals, we will have more organisms available to measure, although the definitions of ‘trait’, ‘class’, ‘type’, and similar terms in this context remain controversial.<sup>18</sup> At a high enough level of abstraction, moving to general types rather than specific traits (‘brown fur’ as a type in mice, say), we might well be able to achieve medium-sized samples, from which we could argue that we have reasonably well estimated the fitness of the trait at issue.

The best possible sampling of the distribution of evolutionary outcomes will occur when we define fitness as the fitness of broader types, *and* we have a type with a vast, vast number of instances, *and* they are all available to measure, at least in the aggregate. For example, consider the ‘digests citrate’ type in the population of *E. coli* bacteria in the Long-Term Evolution Experiment. This type was expressed not just billions of times in the evolving laboratory population itself, but re-expressed billions of times more as the bacterium was thawed from the group’s historical library of colonies and tested to see how often the citrate type would reappear.<sup>19</sup> Here, plausibly, we may have a representative sample of the probability distribution associated with all of that type’s outcomes. But this kind of prediction only allows us to overcome evolution’s unpredictability in an extremely limited set of circumstances, one which will not be applicable to the majority of potential uses of predictive fitness.

It’s important to be clear about the way in which I’m claiming that these kinds of fitness predictions sometimes fail (Table 1). To be sure, biologists have an enormous ‘toolbox’ of models of most of evolution’s central concepts, including fitness.<sup>20</sup> Given a particular set of data, one of those models will be preferable – it will provide the best prediction available in the circumstances. But that’s not the kind of worry that I’m trying to indicate here. However clever we might be in constructing models, we in many situations may find ourselves unable to supply those models with the correct kinds of data, as we can only access small, unrepresentative samples of the total distribution of possible outcomes at issue. This makes it unlikely that predictive fitness can give us any real-world predictive power, outside of cases like type-fitness in experimental evolution. We simply won’t be able to collect enough data about the individuals or traits involved to be able to offer good predictions.

#### 4. Chaotic Population Dynamics

Things become even more complex when we take into account some recent work on chaotic population dynamics derived in mathematical biology, offering us yet another way of analysing how evolving systems might be said to be random. Most models of fitness take, as one of their background assumptions, the claim that evolutionary dynamics are non-chaotic (even if this assumption is not often made

**Table 1.** Inferential bases and sample sizes/conditions for various definitions of individual and trait predictive fitness.

<b>Fitness Measure</b>	<b>Inferential Basis</b>	<b>Sample Size/Character</b>
Trait fitness, including environmental and pleiotropic effects	One trait-history	Very small, unrepresentative
Trait fitness, including similar traits	A small number of trait-histories in similar environmental conditions	Small, likely unrepresentative
Type fitness, natural populations	A moderate number of type-histories in similar environmental conditions	Moderately-sized, possibly representative
Type fitness, experimental evolution	A huge number of type-histories in nearly identical environmental conditions	Large and representative, high-quality predictions

explicit). That is, organisms or populations that are very similar to each other – that are now close to one another in phenotypic space – will remain so as they evolve over time, unless the environment changes. This assumption is necessary for many of the basic mathematical models of fitness to take on stable values. It would be valid in all evolving populations if it were the case that, as Joel Cohen has argued, ‘under reasonable conditions, which are likely to be satisfied in demographic applications’, chaotic population dynamics were not present.<sup>21</sup> What we need, then, is some way to determine just how likely chaotic dynamics are in real-world populations. Will chaotic population behaviour regularly, or only rarely, frustrate our predictions?

As it turns out, a simulation approach to answer precisely this question has been developed by Michael Doebeli and Iaroslav Ispolatov. They begin by noting that evolution is almost always seen as a kind of optimizing process in phenotype space, ‘a dynamical system that converges to an equilibrium’, with unpredictability over long time scales ‘usually attributed to changes in the external environment causing shifts in evolutionary optima’.<sup>22</sup> Given enough time, we take it that evolution will not only find a solution to a given environmental challenge, but will find something approaching the *best* solution available (granted, of course, the constraints of development and the environment). Similar points have been made by philosophers in other contexts – for example, when discussing the nature of evolutionary contingency.<sup>23</sup> In a stable environment, one tends to expect stable and predictable results; we therefore would be inclined to think that it is environmental change, not any inherent randomness in evolutionary dynamics, that might pose problems for our ability to predict those results. As Doebeli and Ispolatov note, this kind of perspective can be tied fairly directly to implicit assumptions (about which

more later) that evolution is a process optimizing for (mathematical) fitness. But as their work shows us, that picture will often fail.

Their approach hinges on introducing to the traditional conception of fitness two features of evolution that are well understood, reasonably common in natural populations, but not often explicitly considered in mathematical models. The first is *density-dependent selection*. A relatively elementary evolutionary phenomenon, density-dependent selection occurs when the extent to which a given trait is advantageous depends upon how dense the population currently is.<sup>24</sup> Cichlid fishes of the same colour, for example, seem to compete intensely for habitat, thus giving a density-dependent advantage to males of rare colours. This advantage may even have been strong enough to drive speciation events.<sup>25</sup> Quite a few natural populations will undergo density-dependent selection, as it is implicated in population regulation, the core set of ecological processes which keep populations at fluctuating yet reasonably bounded sizes over time.<sup>26</sup> Brook and Bradshaw estimate that across a sample of some 1,198 species, between 75 and 92 percent exhibited evidence of density-dependent demographic processes.<sup>27</sup>

The second complicating factor is the *high dimensionality of phenotype space* in natural populations. We know that in most organisms, the number of phenotypic components – that is, elements of the organisms' displayed characters – which contribute to individual fitness is incredibly large. Further, thanks to recent theoretical work by biologists such as Sergey Gavrilets, we know that evolutionary outcomes on high-dimensional 'fitness landscapes' can often differ not just in degree from those on simpler, low-dimensional landscapes, but can differ in kind, producing novel varieties of system behaviour, including more rapid speciation and increased neutral evolution.<sup>28</sup>

The main contribution of Doebeli and Ispolatov's work, then, is the derivation and simulation of a general mathematical model that can take both of these features of real-world population dynamics into account, an important advance for capturing change in natural populations. They proceed to derive an equation for trait dynamics over time, in terms of the dimensionality of the phenotype space, the frequency of each trait, and a number of constants. If organisms do not compete at all, then no change is possible: the population remains at its initial equilibrium. The question, then, is under what circumstances – that is, under what choice of competition constants and for what dimensionalities of phenotype space – do we see long-term chaotic evolutionary dynamics? Evaluating the space of plausible parameters in this model should be able to give us a reasonable estimate of how likely chaotic population dynamics are in general.

Shockingly, for even relatively small dimensionalities (with respect to the extremely high values we might expect in natural populations), chaotic behaviour is nearly certain. In fact, 'the probability of chaos increases with the dimensionality  $d$  of the evolving system, approaching 1 for  $d \sim 75$ .'<sup>29</sup> Even for much smaller dimensionalities (on the order of 15), the trajectories become what is known as ergodic – that is, in the long run, they visit nearly every region of phenotype space.

It is important to realize just how different this model's results are from the traditional picture of long-term evolution. As Doebeli and Ispolatov's model does not take account of change in the external environment, it allows us to conclude that *even if the external environment is constant*, the assumption of density-dependent selection and a reasonably high-dimensional phenotype space results in chaotic dynamics. Further, the difference between any two closely spaced trajectories may not be detectable for a long period of time – that is, the trajectories may evolve down an indistinguishably similar path for quite some time before diverging wildly. They thus exhibit the *sensitive dependence on initial conditions* that is so common in chaotic systems. This is very much not the process of optimizing toward the single point in the space with highest fitness, as we might have expected.

If chaotic dynamics of this sort hold in at least some natural populations, as Doebeli and Ispolatov argue, it is unlikely that it makes sense to speak of *any* of the products of evolution in those populations as the result of an optimizing process, regardless of which quantity evolution is taken to optimize. And this is true *in spite of* the fact that, at every step, the only process driving population change is adaptive. Further, it is unlikely that it makes sense to say that *any* of the products of evolution, even those which appear to be currently highly adaptive, would have been predictable, thanks to the process's sensitive dependence upon its initial conditions.

If Doebeli and Ispolatov's model is widely applicable (about which more in a moment), is there any remaining argument to be made for prediction in cases of chaotic dynamics? The inference from chaos to unpredictability cannot be made *too* hastily. At best, it is *more likely* that actual outcomes will diverge exponentially from predicted values. As Charlotte Werndl has masterfully argued, what makes prediction different in chaotic systems is the fact that 'all sufficiently past events are approximately probabilistically irrelevant'.<sup>30</sup> Such a conclusion in no way makes it impossible to provide (in at least some sense) predictions for chaotic systems.

What is undoubtedly the case, however, is that in chaotic cases, a single prediction – like one we make on the basis of the mathematical fitness of a trait – is unlikely to reflect the future behaviour of a system.<sup>31</sup> Rather, to the extent that we are able to predict anything at all meaningful about the system's future dynamics, it is nothing more than *the chaotic behaviour itself*.

There are, assuredly, a few problems with the application of the Doebeli and Ispolatov results to natural populations that limit the scope of this trouble for predictive fitness. It is certainly true that almost all natural populations undergo at least some degree of density-dependent selection, and for almost all natural populations the dimensionality of phenotype space can be expected to be quite large. But Doebeli and Ispolatov's model indicates only that population dynamics will be ergodic in the long run – and this 'long run' is defined with respect to an arbitrary time parameter appearing in their model, not any biologically relevant unit like years or number of generations. It's thus unclear just how much time is required for chaotic behaviour to emerge, and hence how relevant these results are to observations in natural populations.



Recall, however, that Doebeli and Ispolatov's results hold *in the absence of* environmental changes. It seems plausible, therefore, that either (1) Doebeli and Ispolatov's results indicate that chaotic behaviour in natural populations is a reasonable expectation, or (2) changes in the environment occur often enough that it is unlikely that fitness predictions will be very useful anyway, as the rate of environmental change will make inferences from the current environmental state relatively meaningless. In either event, we have serious trouble for predictive fitness.

### 5. Going Too Far?

One might well object that the argument thus far has proven too much: surely it is not the case that mathematical fitness values are *always* predictively irrelevant? By and large, I agree; there are a number of refinements that should be added to the general case I have presented here. None of these adjustments, however, can offer us broad support for predictions on the basis of mathematical fitness.

First of all, biologists do indeed use various *proxies* for fitness – like the viability or fecundity of organisms at particular ages – to explore the relationship between changes in traits or changes in environments and population structure. They also combine measures of fitness with other information to make inferences about the impact of possible interventions in populations – what would happen if, say, young organisms became twice as fecund? But these are all *ceteris paribus* predictions – if we hold everything else equal, how might we expect the population to respond to a particular kind of intervention? We don't have successful prediction here, so much as a looser conception of fitness, plus information about the life histories of organisms, working in tandem to support claims about possible trajectories for evolving systems.

Further, if we limit ourselves to short-term comparisons within the same environment, we can certainly argue that an organism or trait with a higher fitness is likely to outcompete one with a lower fitness. But the importance of the results discussed here is precisely in the trouble that they make for the transition from these short-term steps, where we might be able to make those kinds of claims, to longer-term knowledge of the evolutionary future. It is here that the complexity and chaos of evolutionary change seems to cause real issues.

Finally, I should also note that I am by no means the first to make an argument of this sort. To take just one example, Jonathan Birch has persuasively argued against the usefulness of considering fitness – whether individual, inclusive or as developed in Alan Grafen's 'formal Darwinism' program<sup>32</sup> – as a quantity that is somehow 'maximized' by natural selection.<sup>33</sup> Birch considers some seven different putative maximization principles, and finds that none constitutes 'a maximization principle with biological meaning'.<sup>34</sup> This result, too, stands in support of the claims argued for here. Were it the case that there was a readily justifiable

fitness-maximization principle, then we might be able to construct predictions on that basis. In the absence of such a principle, however, yet another possible avenue for predicting the trajectory of evolution is foreclosed.


## 6. Conclusion

As we have seen, then, predictions of evolutionary change based upon the fitness values of traits turn out not to be very predictive after all. They are useful only when we can accumulate enough data to make the predictions meaningful, and only when chaotic population dynamics are absent. In some cases we may be reasonably confident that both of these requirements are satisfied, such as in experimental evolution. But in others, we simply *will not know* how to gauge the reliability of fitness predictions. As we've seen, the application of Doebeli and Ispolatov's model needs to be explored empirically, as does the extent to which we will be able to gather data on the fate of a given trait. The failure of such predictions, though, does seem to be much more likely than we might otherwise have thought. A kind of cautious agnosticism may be the order of the day.

These worries for predictive fitness also draw our attention to another long-standing issue in the philosophy of biology: the relationship between short-term and long-term evolutionary processes. Complete continuity is obviously not to be expected. Various kinds of discontinuous phenomena, driving a wedge between our causal descriptions of microevolution and the patterns of macroevolution, have been noted at least since the proposal of Stephen Jay Gould and Niles Eldredge's punctuated equilibrium hypothesis.<sup>35</sup> As Jon Wilkins and Peter Godfrey-Smith put the issue, 'many questions about the adaptive character of evolution depend crucially on the grain at which evolutionary processes are being described'.<sup>36</sup> Or, to quote Jonathan Birch, 'there is a logical gap between claims about short-term changes in gene frequency and claims about longer-term phenotypic evolution'.<sup>37</sup> We might hope to retreat, then, to claims about 'medium-scale' evolutionary processes, as a place where our insights about microevolutionary causal structure could plausibly be extrapolated. But it is precisely upon these medium-scale inferences on the basis of fitness that the results here cast doubt. If we were hoping for fitness to provide us with such predictions, those hopes may ultimately be frustrated.

Where, then, to go from here? There are manifold ways in which evolution might be said to be 'random', and, of these, some seem to seriously impact our ability to make trustworthy predictions about the future trajectory of evolution. Our tendency to think about evolutionary theory as a kind of straightforward optimization process that pushes populations to the maximum fitness value or the optimal organismic design runs the real risk of leading us astray. The complexity and chaoticity of evolutionary change constitute at least one sense of 'randomness' that may keep knowledge of the evolutionary future forever beyond our reach.

## ORCID

Charles H. Pence  <https://orcid.org/0000-0002-6836-6047>

## Acknowledgments

For comments on previous drafts of this work, thanks to Hayley Clatterbuck, Roberta Millstein, Grant Ramsey, and Daniel Swaim. Thanks to an audience at the Deep South Philosophy of Biology Conference, especially Marshall Abrams, Dan Burnston, Lindley Darden, Richard Richards, and Michael Ruse; an audience at the London School of Economics, especially J. McKenzie Alexander, Jonathan Birch, Luc Bovens, Christan List, and Bryan Roberts; an audience at the Mississippi State University Philosophy Colloquium for comments and discussion, especially John Bickle, Bart Moffatt, Kristin Boyce, Michael Bruno, William Kallfelz, and a few impressive undergraduate students whose names I failed to record; and an audience at ISHPSSB 2015 in Montreal which heard a *very* early version of this project, especially Marshall Abrams (again! he is a patient colleague), Hayley Clatterbuck, Roberta Millstein, Elliott Sober, and Joel Velasco.

## Notes

1. Charles Darwin, *On the Origin of Species*, 1st edition (London: John Murray, 1859), 54.
2. See, for instance, the classic presentation in Jacques Monod, *Chance and Necessity: An Essay on the Natural Philosophy of Modern Biology* (New York: Alfred A. Knopf, 1971), as well as Francesca Merlin, 'Evolutionary Chance Mutation: A Defense of the Modern Synthesis' Consensus View', *Philosophy and Theory in Biology*, 2 (2010), e103; and Francesca Merlin, 'Weak Randomness at the Origin of Biological Variation: The Case of Genetic Mutations', in Grant Ramsey and Charles H. Pence (eds.), *Chance in Evolution* (Chicago: University of Chicago Press, 2016), 176–95, on randomness in mutations.
3. Elliott Sober, 'The Two Faces of Fitness', in Rama S. Singh (ed.), *Thinking about Evolution: Historical, Philosophical, and Political Perspectives* (Cambridge, MA: The MIT Press, 2001), 309–21; Mohan Matthen and André Ariew, 'Two Ways of Thinking about Fitness and Natural Selection', *Journal of Philosophy*, 99.2 (2002), 55–83.
4. David L. Stern and Virginie Orgogozo, 'The Loci of Evolution: How Predictable Is Genetic Evolution?', *Evolution* 62.9 (2008), 2155–77.
5. Michael Scriven, 'Explanation and Prediction in Evolutionary Theory', *Science*, 130.3374 (1959), 481.
6. Rasmus Grønfeldt Winther, 'Prediction in Selectionist Evolutionary Theory', *Philosophy of Science*, 76.5 (2009), 889–901.
7. The answer to this question will be messy, as we are generalizing over biological practice. The goal here is to capture the usage of predictive fitness among philosophers of biology. I will return to the topic of biological practice later on. Also, for reasons of space, I will pass over the question of the 'units of selection' problem – whether it is more apt to think of individuals, traits, or populations as the things being selected by natural selection.
8. Richard C. Lewontin, *The Genetic Basis of Evolutionary Change* (New York: Columbia University Press, 1974), 6.

9. Lewontin, *Genetic Basis*, 8.
10. Lewontin, *Genetic Basis*, 10.
11. Lewontin, *Genetic Basis*, 16.
12. See, among others: Matthen and Ariew, 'Two Ways'; Costas B. Krimbas, 'On Fitness', *Biology & Philosophy*, 19.2 (2004), 185–203; Richard C. Lewontin and André Ariew, 'The Confusions of Fitness', *British Journal for the Philosophy of Science*, 55.2 (2004), 347–63; Massimo Pigliucci and Jonathan Michael Kaplan, *Making Sense of Evolution: The Conceptual Foundations of Evolutionary Theory* (Chicago: University of Chicago Press, 2006); Denis M. Walsh, 'Not a Sure Thing: Fitness, Probability, and Causation', *Philosophy of Science*, 77.2 (2010), 147–71.
13. Zachary D. Blount, 'History's Windings in a Flask: Microbial Experiments into Evolutionary Contingency', in Grant Ramsey and Charles H. Pence (eds.), *Chance in Evolution* (Chicago: University of Chicago Press, 2016), 244–63.
14. I have nothing profound in mind by 'quality' here, only something like 'reasonably likely to occur'.
15. Elliott Sober, 'Trait Fitness Is Not a Propensity, but Fitness Variation Is', *Studies in History and Philosophy of Biological and Biomedical Sciences*, 44.3 (2013), 337.
16. To simplify the presentation here, I am also leaving aside the question of what kinds of objects are the bearers of fitness, speaking for the moment only about the fitness of traits. For the importance of individual fitness, see Charles H. Pence and Grant Ramsey, 'A New Foundation for the Propensity Interpretation of Fitness', *British Journal for the Philosophy of Science*, 64.4 (2013), 851–81; Charles H. Pence and Grant Ramsey, 'Is Organismic Fitness at the Basis of Evolutionary Theory?', *Philosophy of Science*, 82.2 (2015), 1081–91.
17. Robert N. Brandon and Scott Carson, 'The Indeterministic Character of Evolutionary Theory: No "No Hidden Variables Proof" but No Room for Determinism Either', *Philosophy of Science*, 63.3 (1996), 315–37.
18. For just one facet of this complexity, see James DiFrisco and Grant Ramsey, 'Adaptationism and Trait Individuation', *Philosophy of Science*, 90.5 (2023), 1234–43.
19. Zachary D. Blount, Jeffrey E. Barrick, Carla J. Davidson, and Richard E. Lenski, 'Genomic Analysis of a Key Innovation in an Experimental *Escherichia coli* Population', *Nature*, 488.7417 (2012), 513–18.
20. C. Kenneth Waters, 'Okasha's Unintended Argument for Toolbox Theorizing', *Philosophy and Phenomenological Research*, 82.1 (2011), 232–40.
21. Joel E. Cohen, 'Comparative Statics and Stochastic Dynamics of Age-Structured Populations', *Theoretical Population Biology*, 16.2 (1979), 164.
22. Michael Doebeli and Iaroslav Ispolatov, 'Chaos and Unpredictability in Evolution', *Evolution*, 68.5 (2014), 1365.
23. John H. Beatty and Eric Cyr Desjardins, 'Natural Selection and History', *Biology & Philosophy*, 24.2 (2009), 231–46.
24. Mark A. Hixon and Darren W. Johnson, 'Density Dependence and Independence', in *Encyclopedia of Life Sciences* (Oxford: John Wiley & Sons, 2001).
25. Ole Seehausen and Dolph Schluter, 'Male-Male Competition and Nuptial Colour Displacement as a Diversifying Force in Lake Victoria Cichlid Fishes', *Proceedings of the Royal Society of London B: Biological Sciences*, 271.1546 (2004), 1345–53.
26. Mark A. Hixon, Stephen W. Pacala, and Stuart A. Sandin, 'Population Regulation: Historical Context and Contemporary Challenges of Open vs. Closed Systems', *Ecology*, 83.6 (2002), 1490–1508.

27. Barry W. Brook and Corey J. A. Bradshaw, 'Strength of Evidence for Density Dependence in Abundance Time Series of 1198 Species', *Ecology*, 87.6 (2006), 1445–51.
28. Pigliucci and Kaplan, *Making Sense of Evolution*, 175–206; Sergey Gavrillets, 'High-Dimensional Fitness Landscapes and Speciation', in Massimo Pigliucci and Gerd B. Müller (eds.), *Toward an Extended Evolutionary Synthesis* (Cambridge, MA: The MIT Press, 2010), 45–79.
29. Doebeli and Ispolatov, 'Chaos and Unpredictability', 1368.
30. Charlotte Werndl, 'What Are the New Implications of Chaos for Unpredictability?', *British Journal for the Philosophy of Science*, 60.1 (2009), 217.
31. As has been argued by Hayley Clatterbuck, 'A Defense of Low-Probability Scientific Explanations', *Philosophy of Science*, 87.1 (2020), 91–112, this is not to say that the chaotic dynamics do not *explain* these outcomes – even the improbable prediction drawn from fitness, were it to occur. Many theories in biology, in fact, obtain much explanatory weight from the occurrence of outcomes that they deem relatively improbable.
32. Alan Grafen, 'The Formal Darwinism Project in Outline', *Biology & Philosophy*, 29.2 (2014), 155–74.
33. Jonathan Birch, 'Natural Selection and the Maximization of Fitness', *Biological Reviews*, 91 (2016), 712–27; Jonathan Birch, 'Fitness Maximization', in Richard Joyce (ed.), *The Routledge Handbook of Evolution and Philosophy* (Abingdon, UK: Routledge, 2017), 49–63.
34. Birch, 'Natural Selection', 713.
35. Niles Eldredge and Stephen Jay Gould, 'Punctuated Equilibria: An Alternative to Phyletic Gradualism', in Thomas J. M. Schopf (ed.), *Models in Paleobiology*, (San Francisco: Freeman, Cooper, 1972).
36. Jon F Wilkins and Peter Godfrey-Smith, 'Adaptationism and the Adaptive Landscape', *Biology & Philosophy*, 24.2 (2009), 213.
37. Birch, 'Fitness Maximization', 53.