

Modeling the Evolution of Social Behavior
A Guide for the Perplexed

Richard McElreath and Robert Boyd

The University of Chicago Press Chicago and London

Richard McElreath and Robert Boyd

The University of Chicago Press, Chicago 60637

The University of Chicago Press, Ltd., London

© 2007 by The University of Chicago

All rights reserved. Published 2007

Printed in the United States of America

16 15 14 13 12 11 10 09 08 07 1 2 3 4 5

ISBN-13: 978-0-226-55826-4 (cloth)

ISBN-10: 0-226-55826-6 (cloth)

⊗ The paper used in this publication meets the minimum requirements of the American National Standard for Information Sciences—Permanence of Paper for Printed Library Materials, ANSI Z39.48-1992.

Contents

Preface	ix
1 Theoretician's Laboratory	1
1.1 The structure of evolutionary theory	3
1.2 The utility of simple models	4
1.3 Why not just simulate?	8
1.4 A model of viability selection	11
1.5 Determining long-term consequences	16
1.6 Nongenetic replication	27
2 Animal Conflict	37
2.1 The Hawk-Dove game	38
2.2 Retaliation	46
2.3 Continuous stable strategies	52
2.4 Ownership, an asymmetry	55
2.5 Resource holding power	58
2.6 Sequential play	60
3 Altruism & Inclusive Fitness	71
3.1 The prisoner's dilemma	72
3.2 Positive assortment	76
3.3 Common descent and inclusive fitness	78
3.4 Rediscovering Hamilton's rule	82
3.5 Justifying Hamilton's rule	97
3.6 Using Hamilton's rule	99
4 Reciprocity	123

4.1	The Axelrod-Hamilton model	124
4.2	Mutants and mistakes	132
4.3	Partner choice	145
4.4	Indirect reciprocity	150
4.5	Reciprocity and collective action	156
5	Animal Communication	173
5.1	Costly signaling theory	174
5.2	Cheap, honest signals	192
5.3	Signaling and altruism	201
5.4	Social learning	206
6	Selection among Groups	223
6.1	Three views of selection	225
6.2	Deriving the Price equation	228
6.3	Selection within and between groups	232
6.4	Dispersal	249
7	Sex Allocation	261
7.1	Fisher's theory of sex allocation	262
7.2	Reproductive value and Fisherian sex ratios	263
7.3	Using the Shaw-Mohler theorem	267
7.4	Biased sex ratios	271
7.5	Breaking the eigen barrier	281
8	Sexual Selection	295
8.1	Quantitative genetic models	298
8.2	Fisher's runaway process	305
8.3	Costly choice and sensory bias	309
8.4	Good genes and sexy sons	313
A	Facts about Derivatives	333
B	Facts about Random Variables	335
C	Binomial Expectations	337
D	Numerical Solution	343
E	Solutions to Problems	349

CONTENTS

vii

Bibliography

393

Index

408

Preface

Imagine a field in which nearly all important theory is written in Latin, but most researchers can barely read Latin and certainly cannot speak it. Everyone cites the Latin papers, and a few even struggle through them for implications. However, most rely upon a tiny cabal of fluent Latin-speakers to develop and vet important theory.

Things aren't quite so bad as this in evolutionary biology. Like all good caricature, it exaggerates but also contains a grain of truth. Most of the important theory developed in the last 50 years is at least partly in the form of mathematics, the "Latin" of the field, yet few animal behaviorists and behavioral ecologists understand formal evolutionary models. The problem is more acute among those who take an evolutionary approach to human behavior, as students in anthropology and psychology are customarily even less numerate than their biology colleagues.

We would like to see the average student and researcher in animal behavior, behavioral ecology, and evolutionary approaches to human behavior become more fluent in the "Latin" of our fields. Increased fluency will help empiricists better appreciate the power and limits of the theory, and more sophisticated consumers will encourage theorists to address issues of empirical importance. Both of us teach courses to this end, and this book arose from our experiences teaching anxious, eager, and hard-working students the basic tools and results of theoretical evolutionary ecology.

How to use this book

While existing theory texts are generally excellent, they often assume too much mathematical background. Typical students in animal behavior, behavioral ecology, anthropology, or psychology had and one semester of calculus long ago. It was hard and didn't seem to have much to do with their interests, and, as a result, they have forgotten most of it. Their algebra skills have atrophied from disuse, and even factoring a polynomial is only an ancient memory, as if from a past life. They have never had proper training in probability or dynamical systems. In our experience, these students need more hand-holding to get them started.

This book does more hand-holding than most, but ultimately learning mathematical theory is just as hard as learning a foreign language. To this end, we have some suggestions. In order to get full use of this book, the reader should:

1. Read the book in order. Each chapter builds on the last. After the first three chapters, it will be easier to jump ahead to topics of individual interest.
2. Work through examples within each chapter. Math is like karate: it must be learned by doing. If you really want to understand this stuff, you should work through every derivation and understand every step. Particularly involved derivations have been set in boxes, however, so that you can read the chapters at two levels, for general understanding and for skill development.
3. Work all the problems. At the end of each chapter, we provide a set of problems that elaborate upon the material in the chapter and allow the reader to practice the techniques we describe. The solutions to these problems are found at the end of the book. We encourage readers to attempt all the problems before looking at the solutions. We know these problems are solvable with only the tools presented in this book because our

own students have solved them. Some of the problems are subtle, however, and we encourage you to collaborate with your friends and fellow students when you attempt them.

4. Be patient, both with the material and yourself. Learning to understand and build formal evolutionary models is much like learning a language. This book is a first course, and it can help the reader understand the grammar and essential vocabulary. It is even a passable pocket dictionary. However, only practice and use of the language will build fluency. Students sometimes jump to the conclusion that they are stupid because they do not immediately understand a theoretical paper of interest. This is not justified. No one quickly learns to comprehend and produce mathematical arguments, any more than one can quickly become fluent in Latin.

The foundation of the foundations

Some reviewers of this book complained that some of the material is old fashioned. True, nearly all of Chapter 2 is 30 years old. The derivation of Hamilton's rule in Chapter 3 is hardly new either, at least to professional theorists. Most of Chapter 4 is old hat to mathematical social scientists, although biologists are unlikely to know much of it.

We think the vintage of the models is irrelevant for two reasons. One reason is that these models still have currency in important debates. But even if this material were no longer relevant to debates within the field (eventually this should be true), it would still be worth learning. Modern physicists do the bulk of their day-to-day work in mechanics using Lagrangian dynamics, a system more general and powerful than the classical Newtonian dynamics most of us learned in secondary school or college. Nonetheless, all physics programs teach Newtonian methods first because they are useful and easy to understand. In the same way, the simplest

game-theoretic and evolutionary models are worth knowing, even if, some day, they will no longer be applied in the day-to-day work of evolutionary theorists.

While this book provides a foundation for understanding the construction and analysis of formal models, it does not treat any of the topics exhaustively. At the end of each chapter, we provide a brief Guide to the Literature that directs interested readers to deeper treatments of the material.

Why so much algebra?

This book has lots of algebra. Much more, in fact, than many more-advanced texts in evolutionary theory. The reason is that beginning students need it. If you really want to learn this material, you have to follow the derivations one mathematical step at a time. Students that haven't done much math in a while often have a shaky grasp on how to do the algebra; they make lots of mistakes, and, even more important, don't have much intuition about how to get from point A to point B in a mathematical argument. All this means that they can find themselves stuck, unable to derive the next result, and unsure of whether they don't understand something fundamental or the obstacle is just a mathematical trick or algebraic error. Many years of teaching this material convinces us that the best remedy is to show lots of intermediate steps in derivations.

Acknowledgments

A number of people gave us feedback and advice during the long evolutionary process of writing this book. We have not always followed their advice, but we always appreciated it. These include Sam Bowles, Ford Denison, Thomas Dudey, Herb Gintis, Rick Grosberg, Mark Grote, Joe Henrich, Rufus Johnstone, Hanna Kokko, Mike Loeb, John McNamara, Karthik Panchanathan, Aimée Plourde Andrew Pamiankowski Peter Rodman, Joan Silk, Pontus Strimling, Masanori

Takezawa, Peter Todd, Eva van den Broek, Annika Wallin, and Franz Weissing.

We are particularly grateful to all of our students who endured various primitive forms of the book, at UCLA and UC Davis. Their reactions, both positive and negative, formed the most valuable feedback possible.

Chapter 1

The Theoretician's Laboratory

Mathematical models and the tools used to analyze them constitute the theoretician's laboratory. Simple mathematical models are experiments aimed at understanding the causal relationships that drive important natural phenomena. Theoreticians in evolutionary biology use a variety of tools to study such models, divining their secrets to reveal how interactions that take place over long time spans shape the evolution of behavior. These models are almost always too simple to make accurate predictions or even accurately represent how any real behavior evolves. Nonetheless, they have proven to be extremely valuable because they help us understand processes too complex to grasp by verbal reasoning alone.

Consequently, evolutionary biology has been greatly influenced by the use of such "experiments." In order to read the primary literature in evolutionary biology or animal behavior, you have to be able to really understand mathematical models. Unfortunately, few programs adequately prepare even advanced graduate students to read and understand mathematical theory. Even when research is largely empirical, predictions are often derived from mathematical

models, and students of animal behavior—be the animals humans or otherwise—are handicapped by not fully understanding the theory from which the empirical work derives. What makes it worse is that having a good education in mathematics or probability, while certainly helpful, isn't enough to understand evolutionary models. One needs to acquire a unique toolbox of ideas for dealing with formal population dynamics.

As economists, anthropologists, and others outside biology have been drawn to this body of theory over the last few decades, the problem has become more acute. Now social scientists of many flavors are struggling with mathematical theories they never encountered in their own fields. Often, we are told that someone has shown some result mathematically, and we will just have to accept it. Sometimes, this leads to uncritical adoption of theory. Other times, it leads to uncritical rejection. Neither is an acceptable state of affairs. Both kin selection (Chapter 3) and reciprocal altruism (Chapter 4), topics central to the evolution of social behavior, are cases where the gulf between theorists and consumers of theory is far too wide.

This book aims to help make the student or professional researcher in biology or the social sciences conversant in the language of the evolutionary theory of social behavior. It assumes that the student can do algebra and has had an introduction to calculus but maybe has forgotten much of it. It covers many of the basic topics: contests and conflicts, kin selection, reciprocity, signaling, multilevel selection, sex allocation, and sexual selection. At the same time, readers will be introduced to many of the typical mathematical tools that are used to analyze evolutionary models. At the end of each chapter, we provide a set of problems that elaborate upon the material in the chapter and allow the reader to practice the techniques we describe. The solutions to these problems are found at the end of the book. Those who work through the book and do all the problems should have a reasonable grasp on the fundamental concepts that underly modern biology's

1.1. THE STRUCTURE OF EVOLUTIONARY THEORY³

understanding of the evolution of behavior, be able to intelligently and critically read papers in journals like *Animal Behaviour* or *Proceedings of the Royal Society, Series B*, and with a little more practice maybe start building models of their own.

1.1 The structure of evolutionary theory

When Darwin left for his voyage around the world on the *Beagle*, he took with him the first volume of Charles Lyell's *Principles of Geology*. Somewhere in South America, he received the second volume by post. Lyell is famous for, among other things, never accepting Darwin's account of evolution by natural selection, apparently because of his religious beliefs. There is a sad irony here because Lyell's work played a crucial role in the development of Darwin's thinking. In fact, in some ways Lyell's principle of uniformitarianism is as central to Darwinism as is natural selection.

Before Lyell, it was common to explain the features of the earth's geology in terms of past catastrophes: floods, earthquakes, and other cataclysms. In contrast, Lyell tried to explain what he observed in terms of the cumulative action of processes that we could observe every day in the world around us—the sinking of lands, the building up of sediments, and so on. By appreciating the accumulated small effects of such processes over long time spans, great changes could be explained.

Darwin took this idea and applied it to populations of organisms. Darwin was a good naturalist and knew a lot about the everyday lives of plants and animals. They mate, they give birth, they move from one place to another, and they die. Darwin's great insight was to see that organisms vary and that the processes of their lives affect which types spread and which vanish. The key to explaining long-run change in nature was to apply Lyell's principle of uniformitarianism to

populations. By keeping track of how the small events of everyday life change the composition of populations, we can explain great events over long time scales.

Biologists have been thinking this way ever since Darwin, but it is still news in some fields. Are people products of their societies or are societies products of people? The answer must be “both,” but theory in the social sciences has tended to take one side or the other. In evolutionary models, this classical conflict between explanations at the level of the society (think Durkheimian social facts) and explanations at the level of individuals (think micro-economics) simply disappears. Population models allow explanation and real causation at both levels (and more than two levels) to exist seamlessly and meaningfully in one theory. We don't have to choose between atomistic and group-level explanations. Instead, one can build models about how groups of individuals can create population-level effects which then change individuals in powerful ways. This aspect of evolutionary theory gives it great power to explain the evolution of behavior in both people and other animals.

1.2 The utility of simple models

The models you will analyze in this chapter, and in fact every model we consider in this book, are much simpler than the real systems that they represent. The populations are infinite (or at least very large), the organisms that populate them have highly stylized trait inheritance and reproduction, the generations are discrete and nonoverlapping, stochastic effects are largely nonexistent, and gene interactions are wholly absent.

Simple evolutionary models never come close to capturing all the detail in any real species or situation. Instead, game theory trades specificity and detail for tractability and clarity. Models are like maps—they are most useful when they contain the details of interest and ignore others. A map of

subways helps a person to navigate by subway. A map of streets helps on the surface. A map with both subways and streets would be too crowded and confusing for either task. A passage from Lewis Carroll's *Sylvie and Bruno Concluded*¹ illustrates this point:

“What do you consider the largest map that would be really useful?”

“About six inches to the mile.”

“Only six inches!” exclaimed Mein Herr. “We very soon got six yards to the mile. Then we tried a hundred yards to the mile. And then came the grandest idea of all! We actually made a map of the country, on the scale of a mile to the mile!”

“Have you used it much?” I enquired.

“It has never been spread out, yet,” said Mein Herr: “The farmers objected: they said it would cover the whole country, and shut out the sunlight! So now we use the country itself, as its own map, and I assure you it does nearly as well.”

Game-theoretic models are simple maps for understanding the consequences of a small number of key assumptions. Despite their deliberate myopia, and perhaps as a result of it, they have proven quite useful in a number of fields, especially in the study of social behavior. Their simplicity limits their power, but it also allows them to do things for us that the real world cannot. For example, we can fit them in our pockets.

The models we will explore in this book are only as complicated as they must be to capture the structure of their problems. They always omit factors which are known to matter in the real world. Such omissions are intentional. It is accepted practice to control for variables in psychology or economics experiments, variables which are known matter in many important situations. But ignoring some (most!) variables allows experimenters to examine the effects of other variables in isolation. Simple models are much the same.

They are not meant to be replicas of the real world, but rather to help us understand it, one piece at a time.

Simple models can aid our understanding of the world in several ways.

Existence proofs. There is usually a very large number of possible accounts of any particular biological or social phenomenon. And since much of the data needed to evaluate these accounts are lost or impossible to collect, it can be challenging to narrow down the field of possibilities. But models which formalize these accounts tell us which are internally consistent and when conclusions follow from their premises. With purely verbal arguments about evolutionary processes, it is all too often the case that our conclusions do not follow from our assumptions. Unaided reasoning about the mass effects of many weak forces operating over many generations has proved to be hazardous. Formalizing our arguments helps us understand which stories are possible explanations.² They provide proof that some candidate set of processes *could* explain the observations of interest.

Aid communication. Formal models are much easier to present and explain. In almost every field in which theory has been important—philosophy, the natural sciences, economics—formal systems for communicating and inspecting theory have arisen. This is partly for the reasons just explained, but formal models are also important because they give us the ability to clearly communicate what we mean. The looseness of verbal models sometimes allows scholars to argue for years about their implications. Such wrangling is much rarer in mathematical disciplines because you have to precisely specify how every piece of the theory relates to every other piece. There is comfort in vagueness, and formal theory allows for little comfort. Sometimes, being forced to spell out an idea in detail is enough to make the formal exercise worth our time.

Counterintuitive results. Simple models often lead to surprising results. If such models always told us what we thought they would, there would be little point in constructing them. Instead, as any practicing theorist can tell you, models often surprise us, producing effects and interactions we hadn't imagined prior to constructing the model. These unanticipated effects lead to additional theory construction and data collection. They take our work in directions we would have missed had we stuck to verbal reasoning, and they help us to understand features of the system of study that were previously mysterious.

Prediction. Simple formal models can be used to make predictions about natural phenomena. In most cases, these predictions are qualitative. We might expect more or less of some effect as we increase the real-world analogue of a particular parameter of the model. For example, a model of voter turnout might predict that turnout increases in times of national crisis. The precise amount of increase per amount of crisis is very hard to predict, as deciding how to measure "crisis" on a quantitative scale seems like a life's work in itself. Even so, the model may provide a useful prediction of the direction of change. These sorts of predictions are sometimes called *comparative statics*. And even very simple models can tell us whether some effect should be linear, exponential, or some other functional form. All too often, scientists construct theories which imply highly nonlinear effects, yet they analyze their data with linear regression and analysis of variance. Allowing our models of the actual phenomenon to make predictions of our data can yield much more analytic power.³ In a few cases, such as sex allocation, models can make very precise quantitative predictions that can be applied to data.

But keep in mind that even when direct prediction from a model seems unlikely, models can direct our attention to things in the real world we should be measuring but had previously ignored. New data will then lead us to revise our

models, leading in turn to more data collection. The relationship between models and data is not a one-way street. Data inspire and challenge theories just as often as theories inspire data collection. Or at least they should.

1.3 Why not just simulate?

There is a growing number of modelers who know very little about analytic methods. Instead, these researchers focus on computer simulations of complex systems. When computers were slow and memory was tight, simulation was not a realistic option. Without analytic methods, it would have taken years to simulate even moderately complex systems. With the rocketing ascent of computer speed and plummeting price of hardware, it has become increasingly easy to simulate very complex systems. This makes it tempting to give up on analytic methods, since most people find them difficult to learn and to understand.

There are several reasons why simulations are poor substitutes for analytic models.

Equations talk. Equations—given the proper training—really do speak to you. They provide intuitions about the reasons an evolutionary system behaves as it does, and these reasons can be read from the expressions that define the dynamics and resting states of the system. Analytic models therefore *tell us* things that we must infer, often with great difficulty, from simulation results. Analytic models can provide proofs, while simulations provide only a collection of examples.

Sensitivity analysis. It is difficult to explore the sensitivity of simulations to changes in parameter values. Parameters are quantities that specify assumptions of the model for a given run of the simulation—things like population size,

mutation rate, and the value of a resource. In analytic models, the effects of these changes can be read directly from equations or by using various analytic techniques. In simulations, there are no analogous expressions. Instead, the analyst has to run a large number of simulations, varying the parameters in all combinations. For a small number of parameters, this may not be so bad. But let's assume a model has four parameters of interest, each of which has only 10 interesting values. Then we require 10^4 simulations. If there are any stochastic effects in the model, we will need maybe 100 or 1000 or 10,000 times as many.

To see how time-consuming this can be, suppose we have a fairly fast computer that can complete a thorough single run of a simulation in 1 second. This may seem like a long time to run a simulation, but a nondeterministic, agent-based simulation can take a long time to run out enough generations to get a good picture of any steady state or dynamic cycles. Now, we will need $10^4 = 10,000$ different combinations of parameter values to entirely map out the parameter space, at the given intervals. Suppose we require 1000 simulations at each combination, to reliably estimate the mean outcome in each case. Now we are looking at 10^7 simulations. At 1 second each, that's almost 116 days of computer time. No doubt computers will one day be fast enough to do 100 or more such simulations per second, but until then computer time is a major constraint.

In practice, people often ignore large portions of the parameter space, and so, many fewer simulations may be needed. There are often obvious and good reasons to do so. But even when large portions of the parameter space can be safely neglected, managing and interpreting the large amounts of data generated from the rest of the combinations can be a giant project, and this data-management problem will remain no matter how fast computers become in the future. Technology cannot save us here. When simple analytic methods can produce the same results, simulation should be avoided, both for economy and sanity.

Computer programs are hard to communicate and verify. There is as yet no standard way to communicate the structure of a simulation, especially a complicated “agent-based” simulation, in which each organism or other entity is kept track of independently.⁴ Often, key aspects of the model are never mentioned at all. Subtle and important details of how organisms reproduce or interact can remain very murky indeed. In contrast, analytic models have benefitted from generations of notational standardization, and even unmentioned assumptions can be read from the expressions in the text. Thus it is much easier for other researchers to verify and reproduce modeling results in the analytic case. Bugs are all too common, and simulations are rarely replicated, so this is not a minor virtue.

Overspecification. The apparent ease of simulation often tempts the modeler to put in every variable which might matter, leading to complicated and uninterpretable models of an already complicated world. Surprising results can emerge from simulations, effects that we cannot explain. In these cases, it's hard to tell what exactly the models have taught us. We had a world we didn't understand and now we have added a model we don't understand.

If the temptation to overspecify is resisted, however, simulation and analytic methods complement each other. Each is probably most useful when practiced alongside the other. There are plenty of important problems for which it is simply impossible to derive analytic results. In these cases, simulation is the only solution. And many important analytic models can be specified entirely as mathematical expressions but cannot be solved, except numerically (we present an example in Chapter 8). For these reasons, we would prefer formal and simulation models be learned side by side. However, most people find learning formal modeling hard enough without learning computer programming on top of it. So while we will sometimes comment on the complementary value of

simulation, for the most part it's pencil to paper, chalk to board, and full steam ahead from here on out.

1.4 A model of viability selection

Enough abstraction, let's get to work. In this section, we will build a simple model of evolutionary changes that result from variation in probability of survival—a form of natural selection that populations geneticists call *viability selection*. Along the way, we'll see how to go from a specification of the life cycle of an organism to deducing the long-term effects on the genes in the population. At the same time, we'll derive some useful expressions that we will use again and again and also learn some useful ways to visualize and interpret evolutionary models.

In any evolutionary model, the minimum we need is:

1. A description of the *population*: This describes how many individuals there are and how the social organization is structured.
2. A set of *heritable variants*: In genetic evolution, this usually includes different genotypes; in cultural evolution we speak of cultural variants, information stored in brains and transmitted from one generation to the next.
3. A *life cycle*: A description of the events in individual lives that affect the survival and proliferation of heritable variants.

The first step in building a model is to derive a mathematical representation of how the events in the life cycle change the distribution of heritable variants in the population over one time period (usually, but not always, one biological generation). Then we want to figure out what happens over the long run, as the events of the life cycle take place over many time periods. As we will see, there are a variety of techniques for accomplishing the second step.




Figure 1.1: Life cycle assumed in the viability selection model. At a given time t , zygotes in the population mature into adults, if they survive. Adults then mate and produce zygotes at time $t + 1$ who undergo the same process.

1.4.1 Selection for survival

Assume there is a population of n individuals, where n is large enough that we can ignore sampling variation. These individuals are haploid, meaning that there is no sex or recombination. There are two genotypes, A and B . Generations are discrete, meaning that individuals go through their life cycle in lock step—everyone is born at the same time, and everyone reproduces at the same time. Suppose n zygotes exist at some time t . Some zygotes manage to survive the slings and arrows of natural selection and mature into adults— A types are more likely to survive than B types. Surviving adults then reproduce, creating a population of zygotes at a time $t + 1$. It's often useful to represent the life cycle using a diagram like that shown in Figure 1.1.

As with each chapter in this book, we summarize definitions of the symbols used in this chapter in a box at the end of the chapter. Formal theory often requires more symbols than can be easily remembered. You can refer to this box as you work through the model in order to make sure you understand each step.

In most of the models we explore in this book, we keep track of the proportions or frequencies of different genotypes or strategies in a population. These frequencies are numbers between zero and one that simply tell us what fraction of the population is of a certain type. They are examples of what mathematicians who study dynamical systems call

state variables, quantities that describe the state of the system and are sufficient to predict the state at future times. In population-genetic models the number of state variables is often one less than the number of alternative alleles or strategies in the population. For example, if a population has only two types, A and B , then the frequency of A , p_A , tells us the portion of the population of type A . The proportion that is type B could be expressed as p_B or simply $(1 - p_A)$, since there are only two alleles in this example. In a system with three alternative types, we would need two state variables to specify all the frequencies: p_1 , p_2 , and $1 - p_1 - p_2$.

In this simplest of evolutionary models, only one evolutionary force is at work. Individuals with some genotypes are more likely to survive than others. In order to deduce the long-term effects of these events, we need to specify how individual genotypes survive from one time period to the next. In essence, we need to perform a census of the genotypes at time t and at time $t + 1$. Recall that there are two genotypes, A and B , in the population. Let p be the frequency of the A genotype in the population at time t . This is simply

$$p = \frac{\text{number of A zygotes}}{n},$$

where n is the number of individuals currently in the population. The above expression means that the number of A zygotes is equal to np . Consequently, the number of B zygotes is $n(1 - p)$. This means we are now keeping track of the differential evolution of types A and B in this population with a single number, p . Thus from stage to stage in the life cycle, the value of p may change, depending upon how events affect the different alleles. The variable p is our only state variable. By convention, p' (pronounced "p prime") represents the value of our state variable during the next stage of the life cycle. In calculus, the $'$ symbol sometimes means the derivative of a function. In this book, it will mean the value of the variable in the next time period or stage in the life cycle.

So far, we have only counted up the numbers of different genotypes among the zygotes, to get our state variable. No biology so far. Next, these zygotes will mature, and some will survive to adulthood, and this is where the biology comes in. Let $V(A)$ and $V(B)$ be the probabilities that a zygote of each type survives to adulthood (we will use V 's throughout the book to refer to payoffs to types). For example, when $V(A) = 0.5$, half of all zygotes with the A genotype survive to be adults. With this assumption, we can do a census for the population of adults. The number of adults with genotype A will be

$$\text{number of } A \text{ adults} = npV(A).$$

The number of B adults is then $n(1-p)V(B)$. The frequency of genotype A among adults, p' , is given by

$$\begin{aligned} p' &= \frac{\text{number of } A \text{ adults}}{\text{number of } A \text{ adults} + \text{number of } B \text{ adults}} \\ &= \frac{npV(A)}{npV(A) + n(1-p)V(B)}. \end{aligned}$$

After selection, the remaining adults mate and reproduce. In this model, let's keep reproduction simple. Each individual reproduces asexually, producing z zygotes, regardless of genotype. When this is true, p'' , the frequency of genotype A among zygotes at time $t + 1$, is given by

$$p'' = \frac{(\text{number of } A \text{ adults})z}{(\text{number of } A \text{ adults})z + (\text{number of } B \text{ adults})z} = p'.$$

Since all genotypes produce the same average number of offspring, reproduction in itself does not change the frequencies of A and B genotypes.

Given the work above, we can write an expression for the frequency of genotype A after one generation:

$$p'' = \frac{pV(A)}{pV(A) + (1-p)V(B)}. \quad (1.1)$$

We normally call an expression of this kind a *recursion*, and it allows us to apply the per-generation effects of evolutionary

Box 1.1 Derivation of viability selection difference equation

Begin by subtracting the frequency of type A at the start of the generation, p , from each side of the recursion:

$$\Delta p = p'' - p = \frac{pV(A)}{pV(A) + (1-p)V(B)} - p.$$

Multiply the top and bottom of the far-right term by $pV(A) + (1-p)V(B)$:

$$\Delta p = \frac{pV(A)}{pV(A) + (1-p)V(B)} - \frac{p(pV(A) + (1-p)V(B))}{pV(A) + (1-p)V(B)}.$$

After simplifying, we obtain

$$\Delta p = \frac{p(1-p)(V(A) - V(B))}{pV(A) + (1-p)V(B)}.$$

We commonly write the denominator $pV(A) + (1-p)V(B)$ and expressions like it as \bar{W} , which is the average fitness in the population. This gives us a final difference equation in the main text.

forces over any number of generations. To see why, consider a case where we know the frequency of genotype A in generation 1 (let's call this p_1). Then we can calculate the frequency of genotype A in generation 2 (p_2) as

$$p_2 = \frac{p_1 V(A)}{p_1 V(A) + (1 - p_1) V(B)}.$$

The frequency of genotype A in generation 3 (p_3) is likewise calculated by substituting p_2 into the recursion.

Often, it is useful to represent these events using a *difference equation*. Instead of specifying the new frequency after one generation, a difference equation yields the *change* in the frequency after one generation. In Box 1.1, we demonstrate how to convert the recursion above into the difference equation

$$\Delta p = p(1-p) \frac{V(A) - V(B)}{\bar{w}}.$$

This recursion and difference equation are commonly called *replicator dynamics* for viability selection. We will see them again and again, so understanding what they say is important. This difference equation reveals something interesting about how natural selection (currently the only force in our model) changes genotype frequencies. The first part, $p(1-p)$, is the variance in genotypes in the population (you will prove this later in the book). When either genotype is common, this product is small. The variance is maximized when both genotypes are equally common ($p = 0.5$). This means that natural selection is strongest when variance is maximized. Natural selection is a culling process, and it changes genotypes most quickly when there is more to cull. At the limit, when there is no variation, natural selection cannot change genotype frequencies at all. We will see variance terms like this crop up again and again.

The second part, $\frac{V(A)-V(B)}{W}$, is simply the proportional increase or decrease of genotype A compared to genotype B . If genotype A is more fit than genotype B , this term will be positive, and the frequency of genotype A will increase each generation as long as there is still some variation in genotypes ($p(1-p) > 0$). Likewise, when $V(A) < V(B)$, the frequency of genotype A will decrease each generation.

A final concern is what happens when fertility, the z 's in the model above, differs among the types. We derive a general combined viability/fertility recursion in Box 1.2.

1.5 Determining long-term consequences

The recursion is a mathematical representation of how events in the lives of individuals change the genetic composition of the population over one generation. Evolutionists want to figure out what will happen in the long run if these processes are repeated generation after generation. There are a number of ways to do this. For very simple recursions like the one above, it is possible to calculate (or rather, guess

Box 1.2 Selection on fertility and viability

What happens when different types produce different numbers of zygotes? This would be selection on fertility, and is in fact what many of the models in this book assume. If we let $z(A)$ and $z(B)$ be the average number of zygotes produced by each type, the viability recursion (Equation 1.1) becomes

$$p'' = \frac{pV(A)z(A)}{pV(A)z(A) + (1-p)V(B)z(B)}.$$

This implies that the fitness of each type is the product of its survival and reproductive success. Thus a combined viability/fertility selection recursion can be written using $W(A) = V(A)z(A)$ and $W(B) = V(B)z(B)$:

$$p'' = \frac{pW(A)}{pW(A) + (1-p)W(B)} = p \frac{W(A)}{\bar{w}}.$$

The combined difference equation is naturally

$$\Delta p = p(1-p) \frac{W(A) - W(B)}{\bar{w}}.$$

at) an explicit solution for the number or frequency of genotype A at any time t . However, this will not be the case for most of the other models we examine. For even slightly more complicated recursions, explicit solutions are almost always impossible. Therefore, in this section we present several ways to go about deriving long-term consequences. First, we explore an explicit solution. In this case, the explicit solution is possible and even of some interest to us. Second, we demonstrate how to visualize the dynamics of these systems and derive outcomes graphically. Third, we explain how to derive long-term consequences in general by linearizing the recursion.

Box 1.3 Explicit solution to viability recursion

Notice that the number (not the frequency) of A alleles at time $t + 1$ is just

$$\begin{aligned} n_{A,t+1} &= (\# \text{ zygotes per individual}) \times n_{A,t} \times (A \text{ type survival}) \\ n_{A,t+1} &= zn_{A,t}V(A). \end{aligned}$$

The number of B alleles is irrelevant in this computation, since the fitness of both types is a constant (it isn't affected by density in any way). So the number of A alleles at time $t + 2$ is just

$$\begin{aligned} n_{A,t+2} &= zn_{A,t+1}V(A) = z\{zn_{A,t}V(A)\}V(A) \\ &= n_{A,t}(zV(A))^2. \end{aligned}$$

and at time $t + 3$ it is

$$\begin{aligned} n_{A,t+3} &= zn_{A,t+2}V(A) = z\{n_{A,t}(zV(A))^2\}V(A) \\ &= n_{A,t}(zV(A))^3. \end{aligned}$$

You can probably see by now that at any time t the number of A alleles will be

$$n_{A,t} = n_{A,0}(zV(A))^t,$$

where $n_{A,0}$ is the number of A alleles in the first time period. The number of B alleles at any time t is similarly

$$n_{B,t} = n_{B,0}(zV(B))^t.$$

So the frequency of A alleles at any time t is just

$$p_t = \frac{n_{A,t}}{n_{A,t} + n_{B,t}} = \frac{n_{A,0}V(A)^t}{n_{A,0}V(A)^t + n_{B,0}V(B)^t}.$$

Finally, dividing the numerator and denominator by a factor $n_{A,0}V(A)^t$ yields the function in the main text (Equation 1.2).




Figure 1.2: The long-term consequences of selection, as derived from Equation 1.2. The frequency of A alleles increases slowly at first, most quickly when the variance is maximized at $p = 0.5$, and again slows down before reaching 1.

1.5.1 Explicit solution

An explicit solution is an expression that tells us the frequency of each type at any time t . We explain in Box 1.3 how to find the explicit solution in this case. The function giving the frequency of A alleles at any time t is

$$p_t = \frac{1}{1 + \left(\frac{n_{B,0}}{n_{A,0}}\right) \left(\frac{V(B)}{V(A)}\right)^t}. \quad (1.2)$$

This is the logistic growth equation commonly used in ecology to model population growth.

Let's plot the frequency over time, given some values for $V(A)$ and $V(B)$. Plotting the frequency over time is often useful for understanding the ongoing evolutionary process. Let the initial frequency of genotype A be 1% ($p_0 = 0.01$), $V(A) = 0.8$, and $V(B) = 0.5$. Figure 1.2 shows the value of p_t for each generation, using the explicit solution in Equation 1.2. Natural selection acts very slowly at first, since there is not much variation in the population. As p increases, selection acts more strongly. Finally, selection slows down again as genotype B s are largely removed from the population.

Eventually, in a finite population, all genotype B s will disappear, and $p = 1$ on the infinite horizon.

We haven't modeled any regulation of the population size here, and in general we won't. Population regulation is interesting, and sometimes it matters a lot exactly how it works (we take this up in Chapter 3), but we're going to focus our attention on mainly other aspects of the biology.

1.5.2 Equilibrium analysis

Explicit solution is hardly ever possible. So, for most models we have to do with a second best option, *equilibrium analysis*. This is the process of deriving values of the state variables (in this case there is only one state variable, p) for which the system does not change from one time period to another and determining if the system moves back to any of these values when perturbed a bit. *Equilibria* are states at which the system gets stuck, so once it gets to an equilibrium state, it stays there. Equilibria are *stable* if the system returns to them once perturbed a small amount. These perturbations can arise from mutation, migration, or other forces. We have to derive the equilibria and determine their stability because not all equilibria are stable. Both kinds of equilibria are of great interest. Stable equilibria are candidates for long-term evolutionary outcomes. Unstable equilibria are interesting, as they usually define the boundaries between *basins of attraction* of two or more stable equilibria. A basin of attraction of a stable equilibrium is the set of initial conditions which converge to that equilibria. When there is more than one equilibrium, we need to know which is the most likely evolutionary outcome, and the size of the basin of attraction is one (but not the only) measure of this. (It can also be that no equilibrium is stable, and the system changes forever. Such behavior is important in ecology but less so in the models in this book.) We'll have examples of all these things before long, so hang on while we work through it.

1.5. DETERMINING LONG-TERM CONSEQUENCES 21


Equilibria in our model here can be found by deriving the frequencies at which $p' = p$. We will label these equilibrium frequencies \hat{p} . It is often easiest to find these equilibria by setting the difference equation (Δp , remember) equal to zero, which implies that we are finding the value of p which leaves the frequency of p the same from one time period to the next. The equilibrium in our system lies where

$$\Delta p = \hat{p}(1 - \hat{p}) \frac{V(A) - V(B)}{\bar{w}} = 0.$$


If any of the terms in this equation equals zero, the whole thing is zero, and we are at an equilibrium. The first two equilibria are easy to spot. When $p = 1$ or $p = 0$, the entire expression is zero. When either genotype exists alone, natural selection cannot change the genotype frequencies. The only other way for Δp to equal zero is when $V(A) = V(B)$. When both genotypes have the same probability of surviving to adulthood, natural selection does change the genotype frequencies. This situation isn't very interesting here, however, as it simply tells us that natural selection does not act in a population when all genotypes have the same fitness. In the following chapters (and one of your homework problems) we will study cases in which fitness depends on the composition population, and in such cases, this equilibrium will be very interesting to us. But not yet.

Taking the two meaningful equilibria, $\hat{p} = 0$ and $\hat{p} = 1$, let's ask when they are stable. Figure 1.3 plots the recursion from the model we derived above for all values of p from zero to one, for $V(A) > V(B)$. The thick line plots p' as a function of p as given by the recursion, and the dashed line is simply the graph of the set of points for which $p' = p$. Whenever the thick line and the dashed line intersect, there is an equilibrium. In this case, there are equilibria at zero and one.

Now let's investigate what happens when the system is at $\hat{p} = 0$ and we perturb it a bit away, so that p is slightly greater than zero (Figure 1.3(a)). Will the system go back to $\hat{p} = 0$ or will it run away? There is a simple graphical way



(a)



(b)

Figure 1.3: Graphical dynamics when $V(A) > V(B)$. (a) *An unstable equilibrium.* The system runs away toward the other equilibrium when perturbed a bit away from the equilibrium at $\hat{p} = 0$. (b) *A stable equilibrium.* In this case, the population returns to the equilibrium when perturbed a bit away from the equilibrium at $\hat{p} = 1$.

of answering this question, using graphs like the the one in Figure 1.3(a). Let's pick a value of p along the x -axis that is slightly greater than zero and draw a line upwards until it meets the curve for p' . We just traced one iteration of the recursion. To find where p goes from this point, we extend our line to the diagonal line where $p' = p$, then change direction again and meet the curve for p' , as shown in Figure 1.3. We could continue this process all the way to $\hat{p} = 1$, which is where the system is heading. Thus, the equilibrium at $\hat{p} = 0$ is unstable. What about the equilibrium at $\hat{p} = 1$? Using the same technique (Figure 1.3(b)), we see that the system returns to $\hat{p} = 1$, so this equilibrium is stable.

This technique works for any continuous recursion. We'll encounter several more complex recursions very soon, and it is left as an exercise for the reader to verify that the graphical method here and the analytical method we're about to encounter lead to the same conclusions.

1.5.3 Linear stability analysis

The formal method for determining the stability of an equilibrium is to compute the derivative of the recursion and evaluate it at the equilibrium. If the resulting value is less than 1 and greater than -1 , the equilibrium is stable. Why is this? Computing the derivative of a function tells us its rate of change at each point. In this context, the rate of change is the change in p' as a function of p . It is the slope of the line tangent to the recursion for a given value of p .

Consider now an imaginary recursion that is just the straight line $p' = p$. At every point on this function, the slope is 1. Every point is an equilibrium, and so there is no deterministic change at any point p . Thus you can see that when the slope is 1, there is no deterministic result after a small perturbation. The system will stick wherever it ends up after the perturbation.

Suppose at some point the slope is less than 1, however. This means the recursion passes *under* the line $p' = p$ on the

right side of the equilibrium and *over* the line on the left side of it (Figure 1.4(a)). Now positive perturbations lead to diminishing values, until the system returns to the equilibrium. The same argument explains why negative perturbations also return to the equilibrium. Since the slope of this recursion is less than 1, the equilibrium is stable. Technically, the slope must also be greater than -1 in order for the equilibrium to be stable. Can you figure out why, using the graphical method?

For nonlinear recursions, we use the same tactic. If perturbations away from the equilibrium are small enough, then a linear estimation of the function at the equilibrium point is all we need to know in order to tell whether the equilibrium is stable (Figure 1.4(b)). In essence, for a curved recursion, we want to know if it looks like the recursion in Figure 1.4(a) in the region of equilibrium under question. If the slope in this region is less than 1 and greater than -1 , then it's stable, at least for small perturbations in the region of the equilibrium.

So how do we do this mathematically? Here's the recipe:

1. Take the derivative of the general recursion (Equation 1.1 in this case) with respect to p . This gives us the slope of the recursion at any point p .
2. Find the value of the derivative when $p = \hat{p}$. This gives us the slope of the line tangent to the recursion at the equilibrium point \hat{p} .
3. If this value is less than 1 and greater than -1 , the equilibrium is stable.

Step 1. Let's take the recursion we derived previously and find its derivative. We want

$$\frac{dp'}{dp} = \frac{d}{dp} \left(\frac{pV(A)}{\bar{w}} \right) = \frac{V(A)V(B)}{\bar{w}^2}.$$

In Box 1.4, we remind the reader how to find this result.

1.5. DETERMINING LONG-TERM CONSEQUENCES 25





Figure 1.4: Linear estimation of evolutionary dynamics. (a) A *linear recursion*. The slope of the recursion at the equilibrium is less than 1, so it is stable, as shown by the graphical dynamics. (b) A *nonlinear recursion*. The slope of the recursion at the equilibrium (shown by the line tangent to the curve at $p = 0.5$) is less than 1, so this equilibrium is also stable.

Box 1.4 Derivative of viability recursion

To calculate the derivative $\frac{d}{dp} \frac{pV(A)}{\bar{W}}$, recall the quotient rule:

$$f(x) = \frac{g(x)}{h(x)} : \frac{d}{dx} f(x) = \frac{\frac{d}{dx} g(x) \times h(x) - g(x) \times \frac{d}{dx} h(x)}{h(x)^2}.$$

Thus in our case, $f(p)$ is the entire recursion and

$$\begin{aligned} g(p) &= pV(A), \\ h(p) &= \bar{W} = pV(A) + (1-p)V(B). \end{aligned}$$

The derivative of $pV(A)$ is just $V(A)$. This is true only because $V(A)$ does not depend upon p . If it did, the derivative would be more complicated. The derivative of $h(p)$ is similarly

$$\frac{d}{dp} h(p) = V(A) - V(B).$$

Combining the above expressions gives

$$\begin{aligned} \frac{d}{dp} \left(\frac{pV(A)}{\bar{W}} \right) &= \frac{\{V(A)\}\bar{W} - pV(A)\{V(A) - V(B)\}}{\bar{w}^2} \\ &= V(A) \frac{\{pV(A) + (1-p)V(B)\} - p\{V(A) - V(B)\}}{\bar{w}^2} \\ &= \frac{V(A)V(B)}{\bar{w}^2}. \end{aligned}$$

Step 2. Let's consider the first equilibrium, at $\hat{p} = 0$. We find the value of the derivative above when $p = \hat{p} = 0$:

$$\left. \frac{dp'}{dp} \right|_{p=0} = \frac{V(A)V(B)}{V(B)^2} = \frac{V(A)}{V(B)}.$$

Step 3. Then we ask when the slope is less than 1 and greater than -1 . Since neither $V(A)$ nor $V(B)$ can be negative, the slope is always greater than -1 . What about less than 1?

$$\frac{V(A)}{V(B)} < 1$$

when

$$V(A) < V(B).$$

When $V(A) > V(B)$, the slope will be greater than one, and so the equilibrium at $\hat{p} = 0$ is unstable. When the opposite is true, $V(A) < V(B)$, the quotient is less than one, and so it is stable. Happily, this makes sense: when the fitness of the A allele is greater, it will increase in frequency after it is introduced in small numbers. If instead B is more fit, selection will remove any rare A alleles after introduction. A simple calculation shows a symmetrical result for the other equilibrium, at $\hat{p} = 1$.

1.6 Nongenetic replication

In some species, many important behavioral variants are transmitted by social learning rather than by genetic transmission. We might wonder if the replicator dynamics we derived above are applicable to such cases. Here, we will introduce you to a very useful tool called a *mating table*, as well as derive the dynamics for one form of social learning. We will see that social learning dynamics can be very similar to viability selection dynamics for simple genetic models.

A mating table is used in population-genetic models to calculate the outcomes of different possible matings in a population. For any particular set of parents, the mating table gives the probabilities of each type of offspring. Using these probabilities, we can then write a set of recursions for the system. The nice thing about mating tables is that they are very general tools. We can use them for almost all problems which will concern us in this book. And like many tools, they break a complex problem into smaller manageable components.

Here, we will construct a “mating” table that calculates the probabilities of different behaviors being imitated, given that an individual encounters a certain set of models. Suppose individuals behave, receive payoffs, and then compare their own payoff to their estimate of the payoff of another,

random individual. If the other individual has a higher payoff, the individual imitates the other individual's behavior. The probability that an individual acquires another behavior is proportional to the difference between the two payoffs.

We begin a mating table by writing down the types of possible "encounters" we are interested in. Let's assume here that there are only two possible alternative behaviors, A and B . Then there are four possible combinations: A and A , A and B , B and A , and B and B . For each of these possible combinations of self and other, we now write the probability that this combination occurs. Let p be the frequency of behavior A . Since the other individual is chosen at random, the probability that the other individual practices behavior A is p and the probability that he practices behavior B is $1 - p$. The chance that our focal individual has behavior A is likewise p , and the chance of B is likewise $1 - p$. This gives us the familiar probabilities for each type of interaction, shown in Table 1.1.

When the number of a certain type in the population is expressed as a frequency, we can use this frequency to compute the probability of different types of encounters. Given that only types A and B exist in a population and that the frequency of type A is p , the chances of each type of two-person encounter are

$$\begin{aligned}\Pr(A, A) &= p^2 \\ \Pr(A, B) &= 2p(1 - p) \\ \Pr(B, B) &= (1 - p)^2.\end{aligned}$$

This is exactly like random mating in The Hardy-Weinberg model, which many of you learned about during your first year in college. The conditional probabilities are even easier, assuming random interaction (which we often will). Given that an individual is type A , the chance she meets an individual of type A ($\Pr(A|A)$) is just p . Likewise, the chance she meets a type B ($\Pr(B|A)$) is just $(1 - p)$.

Now we write a column for each type the focal individual

can end up as. If this were a usual genetic model, we'd write down a column for each genotype of the offspring. In this case, we write a column for each behavioral alternative. In these columns, we write the probabilities that the interaction on that row produces the given behavioral alternative after social learning. In order to write these probabilities, we need to be a bit more specific about what we are assuming.

We need an expression that produces in any A and B interaction, when the payoffs to behaviors A and B are equal, an equal chance of the individual acquiring either. If the payoffs are equal, the individual will be indifferent about which behavior he should imitate. We could achieve indifference by entering a $1/2$ in each cell for each A/B row (there are two A/B rows in the table). But then we need to bias this $1/2$ by the difference in payoff between behavior A and behavior B . Assume that an individual who is in a pair with both an individual with behavior A and one with behavior B acquires behavior B with probability

$$\Pr(B|A, B) = \frac{1}{2} + \beta\{V(B) - V(A)\}.$$

$V(A)$ and $V(B)$ are the payoffs to behaviors A and B , respectively. The parameter β scales the term in braces so that the total probability is always between zero and one. If we make beta small, this is easy enough. Essentially, β tells

Table 1.1: Mating table Step 1: Kinds of interactions and the probability of each.

Self	Other	Probability
A	A	p^2
A	B	$p(1-p)$
B	A	$(1-p)p$
B	B	$(1-p)^2$

Table 1.2: Mating table Step 2: Probabilities of acquiring each behavior. To write a recursion for the frequency of A , p , we multiply the probability in each row by $\Pr(A|\text{self, other})$ in each row. Then we add up each of these row products to get the total expected frequency of A 's in the next time period.

Self	Other	$\Pr(\text{self, other})$	$\Pr(A \text{self, other})$	$\Pr(B \text{self, other})$
A	A	p^2	1	0
A	B	$p(1-p)$	$\frac{1}{2} + \beta\{V(A) - V(B)\}$	$\frac{1}{2} + \beta\{V(B) - V(A)\}$
B	A	$(1-p)p$	$\frac{1}{2} + \beta\{V(A) - V(B)\}$	$\frac{1}{2} + \beta\{V(B) - V(A)\}$
B	B	$(1-p)^2$	0	1

us how strongly imitation is biased by difference in payoff. Now look what happens when $V(B) > V(A)$. Then the total probability is greater than $1/2$, so odds are that the individual acquires the more profitable behavior. Likewise, $\Pr(A|A, B)$ is given by

$$\Pr(A|A, B) = \frac{1}{2} + \beta\{V(A) - V(B)\}.$$

Where did these expressions come from? We made them up, as descriptions of the outcomes of social learning. Part of the creative art in game theory is imagining such formal expressions for outcomes that capture our intentions and make the math simpler. The expressions above fulfill our desiderata of being random choice when there is no payoff difference between the individuals and increasing bias in imitation toward the higher payoff individual when they do differ. Let's fill out the mating table now, by entering these probabilities in the correct cells (Table 1.2). The probabilities in the last two columns in the table must sum to one for each row. A little algebra confirms that $\frac{1}{2} + \beta\{V(A) - V(B)\} + \frac{1}{2} + \beta\{V(B) - V(A)\} = 1$.

We've written the entire mating table, so now we can construct a recursion for the frequency of either behavior. We accomplish this by multiplying the probability of each type of interaction (from the third column) by the probability of the outcome type we are counting (either the fourth or fifth

Box 1.5 Deriving a difference equation from a mating table

We take each probability of a pairing in Table 1.2 and multiply that by the probability that the pairing produces an individual of type A :

$$\begin{aligned} p' &= \Pr(A, A) \Pr(A|A, A) + \Pr(A, B) \Pr(A|A, B) \\ &\quad + \Pr(B, A) \Pr(A|B, A) + \Pr(B, B) \Pr(A|B, B) \\ p' &= p^2(1) + 2p(1-p) \left(\frac{1}{2} + \beta\{V(A) - V(B)\} \right) + (1-p)^2(0). \end{aligned}$$

We've added together the two middle rows to produce the term that begins $2p(1-p)$. To turn this into a difference equation, we just subtract p from both sides and simplify:

$$\begin{aligned} \Delta p &= p' - p \\ &= p^2 + 2p(1-p) \left(\frac{1}{2} + \beta\{V(A) - V(B)\} \right) - p \\ &= p(1-p)2\beta\{V(A) - V(B)\}. \end{aligned}$$

column), for all rows in the table. In Box 1.5, we show how to follow this procedure to compute the difference equation for behavior A :

$$\Delta p = p(1-p)2\beta\{V(A) - V(B)\}.$$

This expression is often simply called *the replicator dynamic*, and it has been derived independently many times, in more or less general forms, in both economics and biology.⁵ This is very similar to the viability selection dynamics we derived earlier. The change in behavior A depends upon the variance in behaviors in the population and the direction, and magnitude of change depends upon the payoffs to the different behaviors, just as in the simple haploid genetic model. Provided the population is very large, it turns out that many different assumptions about payoff-sensitive replication lead to similar replicator dynamics. This fact will be useful to us, since it will allow us to use the same game-theoretic tools, at least with some squinting, for models of both genetic and learned

behaviors. Keep in mind, however, that it is always possible to derive unique replicator dynamics for a given model by using a mating table. We are not stuck with any particular replicator dynamic, although many different dynamics may lead to the same or very similar equilibria.

Guide to the Literature

Philosophy of model building. Wimsatt's 1987 essay on models is helpful for many, students and professionals alike. Practical advice on model formulation and application can be found in Hilborn and Mangel's 1997 book, *The Ecological Detective*. Both give reasons why the best model for a scientific purpose may not be the most realistic model. **Population genetics and ecology.** Roughgarden's 1979 text is now in print again (Roughgarden 1996) and remains a useful and clear introduction to formal evolutionary ecology, especially when it comes to the goals and procedures of analysis. Alan Hastings' 1997 primer on population ecology covers inter-specific and age-structured models that we essentially ignore. **Social learning dynamics.** The payoff-biased imitation recursion in this chapter is commonly employed in economics now. Many other assumptions about imitation are possible. The psychology of how people learn is surprisingly poorly understood. There are many successful models, but most tend to make very similar predictions. See Camerer 2003, Chapter 6, for a nice review and discussion. Highly non-linear learning, like majority-rule conformity, produces quite different replicator dynamics than that we derived here. See Boyd and Richerson 1985 for models of other imitation rules.

Problems

1. Linear dynamics. Use the graphical method to sketch the trajectory for the recursion $x' = x + a(x - 1)$ for $a = -1.5$, $a = -2.0$, and $a = -2.5$.

2. Two-way mutation. Natural selection isn't the only force which changes the frequencies of alleles, of course. Let's derive the dynamics for a simple model in which mutation slowly adds and subtracts alleles. Imagine a very large population with two alleles at a single locus, labeled 1 and 2. Let p represent the frequency of allele 1. Further more, let m_{12} be the proportion of allele 1's that become allele 2's each generation. Similarly, m_{21} is the proportion of allele 2's that become allele 1's each generation. These numbers are usually called mutation rates. Assume that mutation occurs during gamete production. The frequency of allele 1 after mutation is given by

$$p' = p - m_{12}p + m_{21}(1 - p).$$

A fraction m_{12} of allele 1's become allele 2's, and so are subtracted each generation. A fraction m_{21} of the allele 2's become allele 1's, and so are added each generation.

Show that this system has a single equilibrium at

$$\hat{p} = \frac{m_{21}}{m_{12} + m_{21}}.$$

Also prove that this equilibrium is stable.

3. Social learning and replicator dynamics. Imagine an organism capable of social learning that acquires a certain behavior by selecting two models at random and imitating with some probability the one with the more successful behavior. Suppose there are two alternative behaviors and that one is better on average, and this difference is observable, but with some error. Precisely, assume that individuals make mistakes in comparing payoffs at a rate e such that they acquire the more successful behavior $1 - e$ of the time and mistakenly acquire the less successful behavior e of the time. When both models are the same, assume individuals copy either model at random. Construct the mating table for this model and use it to show that the change in the frequency of the more profitable trait, p , is

$$\Delta p = p(1 - p)(1 - 2e).$$

4. When fitness is not a constant. Sometimes, the payoff to a certain behavior or the fitness of a certain allele depends upon the frequency of one or more behaviors or alleles in the population. A well-known example of this is the human sickle-cell allele, which has highest fitness in the heterozygote genotype.

Let w_{AA} , w_{AS} , and w_{SS} be the fitnesses of each genotype. For a diploid system with two alleles A and S , the change in the frequency of the A allele after selection is given by

$$\Delta p = p(1 - p) \frac{W_A - W_S}{\bar{w}}.$$

- (a) Write expressions for W_A , W_S , and \bar{W} .
 (b) Show that this system has an equilibrium at

$$\hat{p} = \frac{w_{AS} - w_{SS}}{2w_{AS} - w_{AA} - w_{SS}}.$$

(c) Using Excel (or any other program you like), plot the dynamics of this system for all values of p , assuming first that $w_{AA} = 0.9$, $w_{AS} = 1$, and $w_{SS} = 0.5$ and then that $w_{AA} = 0.9$, $w_{AS} = 0.75$, and $w_{SS} = 0.5$. That is, produce two graphs.

(d) Use the graphical method to prove that the internal equilibrium is stable only when w_{AS} is greater than both w_{AA} and w_{SS} . Provide a verbal argument as to why this is true. How does the system behave when w_{AS} is smaller than both w_{AA} and w_{SS} ?

5. When environments vary with time. Consider a haploid population in which there are two genotypes A and B . This population lives in an seasonal environment with a wet season and a dry season. Individuals are born at the beginning of each season, reproduce during that season, and die before the next season begins. The A genotype has fitness 1 in both seasons, while B has fitness 0.2 in the dry season and fitness 2 in wet season. Show that the only stable equilibrium is a population in which all individuals are A . How can you reconcile this with the fact that B has higher average fitness?

6. Horizontal cultural transmission, again. Consider a population of individuals who live forever. They have one of two culturally transmitted behaviors, A and B . Each time period individuals accumulate payoffs $V(A)$ and $V(B)$, where $V(A) > V(B)$. Individuals are then paired at random and observe the behavior and payoff of the individual with which they are paired. If the other individual has a higher payoff, they switch with a probability proportional to $\beta(\text{own payoff} - \text{other's payoff})$. Let p be the frequency of A . Show that the change in p over on time period is $\Delta p = \beta(V(A) - V(B))$.

Notes

¹1893.

²The philosopher of biology Robert N. Brandon 1990 argues that such “how-possibly” stories play a crucial role in evolutionary biology.

³See Hilborn and Mangel 1997 for an accessible introduction to marrying simple models to data analysis.

⁴Turchin 2003.

⁵Gintis 2000 shows a derivation for any number of alternative behaviors.

Box 1.6 Symbols used in Chapter 1

- A and B Alternative genotypes in the population
- t Index of the time or generation
- n Number of individuals in the population at the current time
- p Frequency of A in the population
- p' Frequency of A in the next stage of the life cycle
- \hat{p} Equilibrium value of p
- Δp Change in the frequency of A from time t to time $t + 1$
- $V(A)$ Probability a type A individual survives to adulthood
- $V(B)$ Probability a type B individual survives to adulthood
- z Average number of zygotes produced by either type of individual
- \bar{W} Average fitness, defined as $pV(A) + (1 - p)V(B)$
- $z(A)$ Average number of zygotes produced by type A
- $W(A)$ Fitness of type A , defined as $V(A)z(A)$
- $\Pr(x, y)$ Probability of events x and y occurring together
- $\Pr(x|y)$ Probability of event x , given event y
- $\Pr(x|y, z)$ Probability of event x , given events y and z
- β A constant that determines the strength of imitation