

Game Theory in Biology

Concepts and Frontiers

John M. McNamara and Olof Leimar



OSEE

Oxford Series in Ecology and Evolution

GAME THEORY IN BIOLOGY

Oxford Series in Ecology and Evolution

Edited by Paul H. Harvey, Robert M. May, H. Charles J. Godfray, and Jennifer A. Dunne

The Comparative Method in Evolutionary Biology

Paul H. Harvey and Mark D. Pagel

The Cause of Molecular Evolution

John H. Gillespie

Dunnock Behaviour and Social Evolution

N. B. Davies

Natural Selection: Domains, Levels, and Challenges

George C. Williams

Behaviour and Social Evolution of Wasps: The Communal Aggregation Hypothesis

Yosiaki Itô

Life History Invariants: Some Explorations of Symmetry in Evolutionary Ecology

Eric L. Charnov

Quantitative Ecology and the Brown Trout

J. M. Elliott

Sexual Selection and the Barn Swallow

Anders Pape Møller

Ecology and Evolution in Anoxic Worlds

Tom Fenchel and Bland J. Finlay

Anolis Lizards of the Caribbean: Ecology, Evolution, and Plate Tectonics

Jonathan Roughgarden

From Individual Behaviour to Population Ecology

William J. Sutherland

Evolution of Social Insect Colonies: Sex Allocation and Kin Selection

Ross H. Crozier and Pekka Pamilo

Biological Invasions: Theory and Practice

Nanako Shigesada and Kohkichi Kawasaki

Cooperation Among Animals: An Evolutionary Perspective

Lee Alan Dugatkin

Natural Hybridization and Evolution

Michael L. Arnold

The Evolution of Sibling Rivalry

Douglas W. Mock and Geoffrey A. Parker

Asymmetry, Developmental Stability, and Evolution

Anders Pape Møller and John P. Swaddle

Metapopulation Ecology

Ilkka Hanski

Dynamic State Variable Models in Ecology: Methods and Applications

Colin W. Clark and Marc Mangel

The Origin, Expansion, and Demise of Plant Species

Donald A. Levin

The Spatial and Temporal Dynamics of Host-Parasitoid Interactions

Michael P. Hassell

The Ecology of Adaptive Radiation
Dolph Schluter

Parasites and the Behavior of Animals
Janice Moore

Evolutionary Ecology of Birds
Peter Bennett and Ian Owens

The Role of Chromosomal Change in Plant Evolution
Donald A. Levin

Living in Groups
Jens Krause and Graeme D. Ruxton

Stochastic Population Dynamics in Ecology and Conservation
Russell Lande, Steiner Engen, and Bernt-Erik Sæther

The Structure and Dynamics of Geographic Ranges
Kevin J. Gaston

Animal Signals
John Maynard Smith and David Harper

Evolutionary Ecology: The Trinidadian Guppy
Anne E. Magurran

Infectious Diseases in Primates: Behavior, Ecology, and Evolution
Charles L. Nunn and Sonia Altizer

Computational Molecular Evolution
Ziheng Yang

The Evolution and Emergence of RNA Viruses
Edward C. Holmes

Aboveground–Belowground Linkages: Biotic Interactions, Ecosystem Processes, and Global Change
Richard D. Bardgett and David A. Wardle

Principles of Social Evolution
Andrew F. G. Bourke

Maximum Entropy and Ecology: A Theory of Abundance, Distribution, and Energetics
John Harte

Ecological Speciation
Patrik Nosil

Energetic Food Webs: An Analysis of Real and Model Ecosystems
John C. Moore and Peter C. de Ruiter

Evolutionary Biomechanics: Selection, Phylogeny, and Constraint
Graham K. Taylor and Adrian L. R. Thomas

Quantitative Ecology and Evolutionary Biology: Integrating Models with Data
Otso Ovaskainen, Henrik Johan de Knecht, and Maria del Mar Delgado

Mitochondrial Ecology
Geoffrey E. Hill

The Evolutionary Biology of Species
Timothy G. Barraclough

Game Theory in Biology: Concepts and Frontiers
John M. McNamara and Olof Leimar

Game Theory in Biology

Concepts and Frontiers

JOHN M. McNAMARA

School of Mathematics, University of Bristol, UK

OLOF LEIMAR

Department of Zoology, Stockholm University, Sweden

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© John M. McNamara and Olof Leimar 2020

The moral rights of the authors have been asserted

First Edition published in 2020

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2020938654

ISBN 978-0-19-881577-8 (hbk.)

ISBN 978-0-19-881578-5 (pbk.)

DOI: 10.1093/oso/9780198815778.001.0001

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

Acknowledgements	xi
1 Setting the Scene	1
1.1 Introduction	1
1.2 Frequency Dependence	2
1.3 The Modelling Approach	6
1.4 Scope of the Field and Challenges	7
1.5 Approach in This Book	9
2 Central Concepts	13
2.1 Actions, States, and Strategies	13
2.2 The Phenotypic Gambit	16
2.3 Invasion Fitness	19
2.4 Evolutionary Endpoints	22
2.5 Fitness Proxies	23
2.6 From Strategies to Individuals	25
3 Standard Examples	27
3.1 Contributing to the Common Benefit at a Cost	27
3.2 Helping Others: The Prisoner's Dilemma Game	30
3.3 The Tragedy of the Commons	31
3.4 Biparental Care: The Parental Effort Game	32
3.5 Contest Over a Resource: The Hawk–Dove Game	36
3.6 The Evolution of Signalling: From Cue to Signal	41
3.7 Coordination Games	44
3.8 Produce Sons or Daughters? The Sex-Allocation Game	45
3.9 Playing the Field	48
3.10 Dispersal as a Means of Reducing Kin Competition	48
3.11 Threshold Decisions	50
3.12 Assessment and Bayesian Updating	56
3.13 Exercises	60

4	Stability Concepts: Beyond Nash Equilibria	63
4.1	Evolutionarily Stable Strategies	64
4.2	Adaptive Dynamics	69
4.3	Evolution to a Fitness Minimum	74
4.4	Replicator Dynamics	79
4.5	Games Between Relatives	81
4.6	Exercises	87
5	Learning in Large Worlds	91
5.1	Reinforcement Learning	92
5.2	Learning and the Hawk–Dove Game	96
5.3	Learning in a Game of Joint Benefit of Investment	99
5.4	A Dominance Game	102
5.5	Approaches to Learning in Game Theory	107
5.6	Exercises	109
6	Co-evolution of Traits	111
6.1	Stability in More than One Dimension	112
6.2	Role Asymmetries	114
6.3	The Evolution of Anisogamy	118
6.4	Evolution of Abilities and Role Specialization	122
6.5	Learning and Individual Specialization	126
6.6	Co-evolution of Prosociality and Dispersal	130
6.7	Co-evolution of Species	133
6.8	Concluding Comments	137
6.9	Exercises	138
7	Variation, Consistency, and Reputation	141
7.1	Variation has Consequences	141
7.2	Variation and the Stability of Equilibria	143
7.3	Taking a Chance	145
7.4	Signalling and the Handicap Principle	148
7.5	Reputation	150
7.6	Indirect Reciprocity	153
7.7	Differences Select for Social Sensitivity	157
7.8	Markets	160
7.9	Choosiness, Assortment, and Cooperation	164
7.10	Commitment	166
7.11	Exercises	169
8	Interaction, Negotiation, and Learning	173
8.1	Interaction over Time	173
8.2	Information and the Order of Choice	174

8.3	Credible Threats and Strategic Commitment	177
8.4	Negotiation between Partners	180
8.5	Evolution of Cognitive Bias	185
8.6	Social Dominance	188
8.7	Assessment in Contests	193
8.8	Outlook: Games with Interaction over Time	198
8.9	Exercises	201
9	Games Embedded in Life	203
9.1	Self-consistency	203
9.2	The Shadow of the Future, and the Past	204
9.3	Resident Strategy Affects Future Opportunities	206
9.4	Dependence on Future Actions	209
9.5	Territorial Defence and the Desperado Effect	216
9.6	State-dependent Ideal Free Distributions	221
9.7	Is it Worth it?	226
9.8	Exercises	228
10	Structured Populations and Games over Generations	231
10.1	Invasion Fitness for Structured Populations	233
10.2	Offspring Quality versus Number	236
10.3	Reproductive Value Maximization	240
10.4	Sex Allocation as a Game over Generations	242
10.5	The Fisher Runaway Process	246
10.6	Maximizing Lifetime Reproductive Success	252
10.7	Dispersal	255
10.8	Evolutionary Analysis in Structured Populations	257
10.9	Exercises	258
11	Future Perspectives	261
11.1	Phylogeny	263
11.2	Behavioural Mechanisms in Large Worlds	266
11.3	Ontogeny and the Acquisition of Behaviour	271
	Appendix A: Summary of Notation	273
	Appendix B: Solutions to Exercises	275
	References	305
	Index	325

Acknowledgements

We started this book 3 years ago. During that time the book has evolved and improved, helped enormously by the feedback given to us by the following friends and colleagues: Redouan Bshary, Sean Collins, Sasha Dall, Tim Fawcett, Lutz Fromhage, Andy Higginson, Alasdair Houston, Patrick Kennedy, Eva Kisdi, Sean Rands, Brian Robathan, Susanne Schindler, Harry Suter, Pete Trimmer, and Chris Wallace. We are very grateful for their efforts. Thank you!

Part of the modelling work presented in this book was supported by the Swedish Research Council (grant 2018-03772).

Setting the Scene

1.1 Introduction

There is an extraordinarily diverse range of traits and behaviours in the natural world. Birds sing, flock, and migrate, male peacocks display extravagant tails, fish change sex, and ants ‘farm’ aphids. When faced with such phenomena it is natural to ask why. The term ‘why’ can have different meanings. In this book we are concerned with explanations in terms of the past action of natural selection. The action of natural selection tends to favour traits and behaviours that enhance the number of offspring left by an individual, i.e. it produces adaptations to the environment in which organisms live. For example, the aerodynamic shape of a bird’s wing can be seen as an adaptation to increase flight efficiency. Many explanations in biology are similarly in terms of adaptation to the abiotic environment.

However, it is not just the abiotic environment that matters to organisms, but often equally or even more the characteristics of other organisms. The fitness of one individual, involving survival and reproductive success, often depends on how other population members behave, grow, and develop, i.e. depends on the strategies of others. Thus, the best strategy of one individual can depend on the strategies of others. When this is true, we say that the action of natural selection is frequency dependent. For example, oak trees in a wood shade each other. If surrounding trees are tall, then a tree must also grow tall to get more light. This growth requires the tree to divert resources that could otherwise be expended on seed production. The result is that in a wood it is best for an oak to be tall if other oaks are tall, but shorter if they are shorter. The shape of the crown is also predicted to be influenced by this competition (Iwasa et al., 1985).

Biological game theory attempts to model such frequency-dependent situations, and this theory is the subject of this book. The approach of the models is to predict the endpoints of the process of evolution by natural selection, rather than to describe the details of evolutionary trajectories. The overall objective is to understand the selective forces that have shaped observed behavioural and developmental strategies and make predictions that can be tested empirically.

The theory is concerned with the observable characteristics of organisms. These characteristics include both behaviour and morphological and physiological attributes, and are referred to as the phenotype of the organism. We do not try to model the details of the genetics that underlie the phenotype. Instead we usually

assume that the genetic system is capable of producing strategies that do well in their natural environment; i.e. producing phenotypes that are adapted to the environment. This allows the theory to characterize evolutionary outcomes in phenotypic terms. The explanation of tree height in terms of trees shading each other can be understood in phenotypic terms, even though we might be ignorant of the details of the underlying genetics. This 'phenotypic gambit' is further explained in Section 2.2.

This book aims to present the central concepts and modelling approaches in biological game theory. It relates to applications in that it focuses on concepts that have been important for biologists in their attempts to explain observations. This connection between concepts and applications is a recurrent theme throughout the book. The book also aims to highlight the limitations of current models and to signpost directions that we believe are important to develop in the future.

Game theory is applied to a huge range of topics in biology. In Chapter 3 many of the standard games are described, with others being introduced in later chapters. We do not, however, attempt an encyclopedic account of the applications of game theory in biology; our focus is on game-theoretical concepts. Nevertheless, we illustrate these concepts in a range of situations in which the biology is important, incorporating central topics in life-history theory, ecology, and general evolutionary biology along the way.

1.2 Frequency Dependence

Examples of frequency dependence abound in biology. Here we list several to give the reader some sense of the diverse situations in which frequency dependence occurs.

Contesting a resource. When two individuals are in competition for the same food item or mate each might adopt a range of levels of aggression towards their competitor. The optimal level of aggression will depend on how aggressive an opponent is liable to be. If the opponent will fight to the death it is probably best not to be aggressive, while if the opponent is liable to run away if attacked it is best to be aggressive and attack.

Alternative mating tactics. In many species, different males in a population employ different mating tactics. In one widespread pattern some males attempt to defend territories and attract females to their territories, while other males attempt to sneak matings with females that enter these territories. This happens in Coho salmon, where the larger 'hook' males defend territories while the smaller 'jacks' attempt to sneak matings (Gross, 1996). But is it best to be a hook or a jack? It seems plausible that if almost all males are hooks and hold territories, then it is best to be a jack since there will be lots of opportunities to sneak. In contrast, if there are many jacks per hook then sneakers will be in competition with each other and so it is best to be a hook.

Search for resources. If members of a population are searching for suitable nest sites, individuals should be choosy and reject poor sites. However, they cannot afford to be

too choosy, since if they delay too long the better sites will have been taken by others. But how quickly such sites disappear depends on the choice strategy of others. Thus the best choice strategy of one individual depends on the strategies employed by other population members.

Parental effort. Consider two parent birds that are caring for their common young. How much effort should each expend on provisioning the young with food? The more food the young receive the greater their survival prospects, although there are diminishing returns for additional food at high provisioning rates. However, parental effort is typically costly to the parent expending the effort. There are various reasons for this. Increased foraging effort might increase the time that the bird exposes itself to predation risk, and so increase its probability of dying. Foraging might tend to exhaust the bird, leading to an increased probability of disease and hence death overwinter. It might also reduce the time available for other activities such as seeking other mates. In short, each parent faces a trade-off: increased effort provisioning the current brood increases its reproductive success from this brood but reduces its reproductive success in the future. Typically the benefits of increased effort have diminishing returns but the costs (such as the probability of mortality) accelerate with effort. Faced with this trade-off, if a partner expends low effort on parental care it is best to expend high effort, as has been found for several species of birds (Sanz et al., 2000). In contrast, if the partner expends high effort, the young are already doing reasonably well and it is best to expend low effort to reduce costs. In other words, parental effort is frequency dependent.

Joint ventures. In the above example there is a common good (the young) but individual costs. Except for the case of lifelong monogamy, this means that there is a conflict of interest between the parents—each would prefer the other to expend the high effort so that they benefit themselves but do not pay the cost. This conflict of interest applies to many joint ventures. For instance, in a group of animals foraging under predation risk, vigilance benefits the whole group, but is costly to vigilant group members, since their feeding rate is reduced. Cooperative hunting is another example, where all group members benefit from prey capture, but the costs and risks of capturing and subduing prey go to the most active hunters.

Resource specialization. The beak size of a seed-eating bird is often best at exploiting seeds of a certain size. In an environment in which there is a range of seed sizes, if most individuals are specializing in one seed size the competition for these seeds will reduce their availability, so that it is better to have a different beak size and specialize in seeds of a different size. Under these circumstances there may be selection to evolve a range of beak sizes (cf. the evolutionary branching example of Section 4.3).

Timing of migration. For migratory birds, arriving early at the breeding grounds is advantageous in terms of getting the best nest sites, but it also involves risks. For American redstarts, for instance, male birds that arrived early at the breeding grounds in northern Michigan gained higher-quality territories, and early females produced earlier and heavier nestlings, with higher chances of survival to adulthood (Smith and Moore, 2005). A typical outcome of this competition for breeding sites for migratory

birds is that individuals in good condition, who are better able to withstand the risks of early migration, tend to arrive earlier (Kokko, 1999).

Sex allocation. If a female can allocate limited resources between producing sons or producing daughters, her optimal allocation depends on the allocation strategies of other females in the population. At the heart of this dependence is the fact that every offspring produced has exactly one mother and one father. This implies that the total reproductive success of all males in a population must equal the total success of all females. Thus in a population in which other females are producing mostly sons, there will be more males than females and each son will, on average, have lower reproductive success than each daughter. It is then best for our focal female to produce daughters. Conversely, if the allocation strategy of other females results in a population that is female biased, it is best to produce sons (Section 3.8).

This example is of interest because it raises the question of what is fitness. Note that we are measuring the reproductive success of a female in terms of the numbers of matings of her offspring rather than her number of offspring (see Sections 3.8 and 10.4).

Producers versus scroungers. There are many gregarious bird species where flocks of birds search for food. In a flock, some individuals specialize in searching for and finding food, while others watch the searchers and move in to exploit any food source found. Thus a flock consists of ‘producers’ and ‘scroungers’. Figure 1.1 shows




Fig. 1.1 Data from a producer–scrounger field experiment with great tits in Wytham Woods near Oxford (Aplin and Morand-Ferron, 2017a). Visits as producer or scrounger at puzzle-box feeders (Aplin et al., 2015) were recorded for individual birds. Because of frequency dependence, a mixture of producing and scrounging was seen, also for individual birds, but they tended to specialize as either producer or scrounger (a). There was more scrounging in bigger replicate sub-populations (b), which is consistent with the idea that scrounging is more rewarding in high-density populations. Illustration of great tits by Jos Zwarts. Published under the Creative Commons Attribution-Share Alike 4.0 International license (CC BY-SA 4.0).

data from a field experiment recording producer–scrounger behaviours in great tits. In many other situations certain individuals invest in something that becomes a common good of the group, while others exploit the fruits of these investments. For example, females in some bird species deposit their eggs in the nest of other females, so gaining the benefit of the investment of other females in nest building and care. As we outline in Section 3.1, bacteria can acquire iron by producing and releasing siderophores. Siderophores bind to environmental iron and can then be captured by any local bacteria that are present as well as the original releasing bacterium. Siderophore production is thus an investment in a common good that can be exploited by ‘free riders’ that just reap the benefits. There are strong analogies with many human activities in this context (Giraldeau et al., 2017). Exploiting the labour of others is liable to be very profitable when exploiters are rare, but will be less advantageous as a strategy when producers are rare (cf. the example of alternative mating tactics in salmon).

Cooperative hunting. Many social animals hunt in groups or packs, potentially being more efficient at capturing prey that are bigger and harder to subdue. Wolves, lions, and African wild dogs are among the species using this method. Cooperative hunting only works if several group members take part in it, and the rewards should go preferentially to active hunters for the behaviour to be favoured. If this is so it is referred to as a synergistic effect. In general, there is synergism if investing in an activity or a project is only worthwhile if others also invest and the benefits go primarily to those that invest. Synergism can be an important explanation for cooperation.

Warning colouration. Certain organisms are potential prey but are actually unsuitable as prey, for instance because they contain poisonous substances or they are otherwise dangerous or unpleasant to attack. Sometimes they have evolved a striking appearance that advertises their quality as unsuitable prey. The black and yellow stripes of stinging wasps and the striking colouration of monarch butterflies, which contain poisonous cardenolides, are examples of this. Predators learn through experience to avoid attacking such prey, which means that the warning signal is only helpful to prey if it is sufficiently common. There is thus synergism through positive frequency dependence for warning colouration.

Some of the above situations, such as the contest over a resource and parental effort, are concerned with individuals interacting pairwise. In analysing these situations one is dealing with a two-player game. At the other extreme, in the sex allocation example a female is not directly interacting with another female. Instead, her optimal allocation to daughters versus sons only depends on the ratio of males to females in the population, and this ratio is a product of the range of allocation strategies of all other females. The female is then ‘playing the field’, and the resultant game is referred to as a playing-the-field game. The situation of trees in a wood is somewhat intermediate, since a given tree is mostly affected by its local neighbours.

Despite the above seemingly simple classification, we will argue that pairwise competition over a resource is often not really a simple two-player game. Again

consider two individuals contesting a food item. The optimal level of aggression of one individual not only depends on the aggressiveness of the opponent, but also depends on how easy it will be to access other sources of food. If other sources could also be contested, then the availability of alternative food depends on how other population members contest these items. One cannot then treat a single contest for food in isolation, and it is not really a two-player game (Section 9.4).

In most of the above cases that involve a continuous trait the best trait value of any given single individual decreases as the trait value in the rest of the population increases. We refer to this as negative frequency dependence. For a discrete trait such as in the jack versus hook example, a given strategy becomes less profitable as it becomes more common, and we also refer to this as negative frequency dependence. In contrast, there is positive frequency dependence in the example of trees in a wood, as the mean height of trees increases so does the best height for any individual tree. The example of warning colouration also has positive frequency dependence, as do some forms of group hunting.

1.3 The Modelling Approach

Evolutionary game theory aims to predict the evolutionary outcome in situations where frequency dependence exists. In particular it asks whether there is some final endpoint of the evolutionary process, and if so what that endpoint will be. For example, when two parents are caring for their common young, do we predict stable levels of care to eventually evolve? If so, at this stable endpoint is there equal care by both parents or mainly care by one sex? Will the conflict of interest between the parents over care be resolved in a way that results in high levels of care, or will there be low levels with the young suffering as each parent attempts to get its partner to do the bulk of the caring? In the case of sex allocation do we expect an equal number of males and females in populations that have reached an evolutionary endpoint? Also, can we predict how the sex of offspring might depend on local environmental conditions—for example can we explain why in some parasitic wasps females deposit sons in small hosts and daughters in large hosts? We deal with this last question in Section 3.11.

Our focus is on what evolves. By ‘what evolves’ we mean the strategy that evolves. We give a formal definition of strategies in Section 2.1, but we can think of a strategy as a rule that specifies how behaviour and other aspects of the phenotype depend on circumstances. For example, the growth strategy of a tree might be something as simple as to always grow to a given size, but could be more complex, specifying how the allocation of resources to growth depends on light levels and nutrient intake.

As we have explained, although the book is about evolutionary processes, we are mainly concerned with evolutionary endpoints. A formal specification of exactly what is meant by an endpoint is given in Section 2.4 and Chapter 4. Loosely, we can think of an endpoint as having the following property. A strategy will be referred to as the resident strategy when almost all population members follow this strategy. Then a necessary condition for a strategy x^* to be a stable endpoint of the process

of evolution by natural selection is that no single individual with a different strategy can do better than x^* , when x^* is the resident strategy. We will need to specify exactly what we mean by ‘do better’. In game theory in biology, invasion fitness provides a suitable performance measure. This concept is outlined in Section 2.3. For most of the situations analysed in this book, the invasion fitness of a strategy is just the mean number of surviving offspring left by an individual following the strategy. The simple measure of mean number of offspring is not adequate when not all offspring are the same, because the ability of these offspring to leave offspring themselves can then vary with their state. We defer the detailed presentation of this more complex case until Chapter 10.

If x^* is the resident strategy and no other strategy has greater invasion fitness than the resident strategy, we refer to x^* as a Nash equilibrium (Section 2.4). The Nash equilibrium condition is necessary but not sufficient for stability (Chapter 4). The typical approach to modelling a specific situation is first to specify the circumstances in detail, for example the actions that are possible, the time order of events, and so on. We then specify which strategies are possible. For the simplest models, for which an analytic approach is feasible, we can then evaluate how the invasion fitness of each strategy depends on the resident strategy in the population. This allows us to find all Nash equilibria. We then investigate further stability properties of each Nash equilibrium.

Often it is cumbersome to deal directly with invasion fitness and instead we work with a simpler performance measure (such as the rate of energy intake in some situations), which we know leads to the same Nash equilibria. The use of such fitness proxies is described in Section 2.5.

For many models an analytic approach is infeasible. One option is then to resort to evolutionary simulations. In a simulation a large population of virtual organisms interact, produce offspring, and die, and the offspring inherit the traits from their parent(s) with the occasional mutation producing a different trait. The simulation then follows the population forward over many generations until traits settle down to roughly equilibrium values, which are usually close to the values at a Nash equilibrium.

1.4 Scope of the Field and Challenges

To appreciate what game theory has achieved in biology, and to understand the challenges that still remain, some insights into the history of the field are helpful. Ideas with a game-theoretical flavour were developed already in the early 18th century by mathematicians interested in probability (see Bellhouse and Fillion, 2015), including by de Montmort and Nicolaus Bernoulli. They were interested in a card game called ‘Le Her’. The crucial point was that for certain positions in the game it appeared that there was no single best thing for a player to do. Rather the player should randomize between different plays, in this way preventing the possibility that an opponent could predict the play and take effective countermeasures. We can recognize this as an

instance of frequency dependence. R. A. Fisher, who was a prominent statistician as well as an evolutionary geneticist, took an interest in the matter, publishing a work concluding that indeed randomization can be the ‘solution’ to the problem of what is the best strategy for a player of *Le Her* (Fisher, 1934). Fisher was fascinated by the idea of possible important uses of randomization—most familiar is his insistence that treatments should be randomized in experiments—and when he much later wrote about polymorphism and natural selection (Fisher, 1958), he returned to the idea. By that time he was familiar with game theory as developed by mathematicians and economists, and he proposed that the outcome of natural selection could be that a population is polymorphic, effectively randomizing between phenotypes. The concrete example he had in mind was polymorphic Batesian mimicry in certain butterflies. The variation in appearances could prevent predators from quickly learning to attack a single prey appearance, which again is an instance of frequency dependence. However, Fisher did not develop the idea further.

It was Maynard Smith and Price (1973) who were the first to systematically use game theory in biology. They posed the question of why animal contests over resources, such as over a territory or a dominance position in a social group, are usually of the ‘limited war’ type, being settled without serious injury. They used the so-called Hawk–Dove game (Section 3.5) to conclude that, because of frequency dependence, the evolutionary endpoint could be a mixture or polymorphism of direct aggression and display behaviour; thus to some extent limiting direct aggression. This was a breakthrough partly because of the elegance and simplicity of the Hawk–Dove game, but even more because they presented a method to analyse important frequency-dependent situations in biology. Concerning animal contests, their work was just a starting point, with many further analyses of possible reasons for limited aggression appearing, for instance of the effects of assessment of relative fighting ability by contestants (Parker, 1974).

We describe parts of this work (Sections 3.5, 8.7, 9.4, and 9.5), which certainly has improved the understanding of animal contests. Nevertheless, the original question identified by Maynard Smith and Price (1973) is up to now only partially answered using worked-out models. An example where current models are lacking is aggressive interactions in social groups. This points to challenges for game theory in biology, and our aim is to address these in the book. One challenge lies in the sometimes intractable complexities of analysing representations of the information individuals gain about each other; for instance, information individuals acquire when they form dominance hierarchies in social groups. An equally important challenge is to make contact between models and the strategies and decision processes used by real animals. Our approach, as we outline in the next section, is to introduce behavioural mechanisms into models, including mechanisms inspired by animal psychology and neuroscience. Frequency dependence still plays a role in such models, but evolutionary outcomes are prescriptions of how to represent and respond to events, rather than a randomization between two or more actions.

The analysis of sex allocation is one of the major empirical successes of game theory in biology. The origin of ideas about why sex ratios should be close to 1:1

goes back to work by Düsing in the late 19th century, but it was Fisher (1930) who presented a highly influential formal model (Section 3.8). We have noted that frequency dependence is a crucial ingredient in such analyses (Section 1.2). Much of the theory and tests of sex allocation is surveyed by Charnov (1982), and we describe some of these models in the book (Sections 3.8 and 10.4).

As we have indicated, the roots of game theory as a formal tool to analyse behaviour lie in economics. There are many situations in economics where frequency dependence matters. For example, the best amount to produce or the best price to charge for a product depends on the amounts or prices chosen by other firms. These questions were analysed in early applications of game theory in economics and are referred to as Cournot and Bertrand competition. The ideas have also played a role in making biologists think in terms of markets (see Section 7.8). Other economic applications include hiring decisions in labour markets, the structuring of auctions, and the effects of social norms on behaviour. The definition of a Nash equilibrium in biology is mathematically the same as the standard Nash equilibrium condition of economic game theory. There are, however, differences in approach in the two areas (Box 2.1).

Perhaps the biggest influence from economics on biological game theory is an interest in cooperation, often presented in terms of one or several rounds of the so-called Prisoner's Dilemma (see Sections 3.2, 4.5, 7.3, 7.6, and 8.8). As information about other individuals is likely to be just as important in cooperative as in competitive interactions, these present another major challenge to game theory in biology. Our approach is to model such situations in ways that correspond in a reasonable way to biological reality (Sections 3.2 and 8.8), for instance using continuously varying investments for parents that jointly rear their young, and then to incorporate behavioural mechanisms into the model, for instance mechanisms that can be interpreted as negotiation between partners (Section 8.4). We discuss these matters in Section 8.8, pointing to successes and challenges for the current state of affairs in game theory.

1.5 Approach in This Book

Here we give an overview of major issues that we deal with in this book, indicating where in the book they occur. The central concepts of biological game theory, such as invasion fitness, the Nash equilibrium, and the phenotypic gambit are expanded on in Chapter 2. That chapter also deals with fitness proxies and the problem of embedding a game that occupies a small part of the life of an organism into a whole-life perspective, a topic we return to in Chapter 9.

A variety of game-theoretic models dealing with a range of phenomena have previously been developed. In Chapter 3 we present many of what have become the standard models. Some of these standard models deal with cooperation and the contribution to a common good, including parental care. We also introduce the simplest model of animal conflict over a resource: the Hawk–Dove game. Many

animals signal to others, and we present a simple model showing that signals can evolve from cues, later returning to the question of why signals should be honest in Section 7.4. We also present the standard model of sex allocation, a topic we later return to in Section 10.4. Most of these simple models assume that all individuals are the same, so that if they take different actions this is because their choice has a random component. In reality it is likely that individuals differ in aspects of their state such as size or fighting ability, and different behaviours are a result of these difference. At the end of Chapter 3 we illustrate how such state-dependent decision making can be incorporated into models. The effects of state differences are harder to analyse when the state of offspring is affected by the state and action(s) of their parent(s). We defer the description of some standard models that have these features until Chapter 10, where we outline the theory needed to analyse inter-generational effects. We apply this theory to the problem of sex allocation when offspring tend to inherit the quality of their mother (Trivers and Willard theory) and to the case where female preference for a male ornament (such as tail length) and the ornament co-evolve (the Fisher process).

A population that has reached an evolutionary stable endpoint is necessarily at a Nash equilibrium. The converse is, however, not true. There are Nash equilibria that are not stable. Furthermore, there can be endpoints that would be stable if they could be reached, but cannot be reached by the evolutionary process. In Chapter 4 we return to general theory. There we set out the conditions needed for a strategy to be an Evolutionarily Stable Strategy (ESS); conditions that are stronger than those required for a Nash equilibrium, and which ensure stability. We also consider evolutionary trajectories using the concepts of adaptive dynamics and convergence stability (Section 4.2).

Evolution is concerned with the spread of genes. In most of the analyses in this book we can translate ideas about the spread of genes directly into the idea that at evolutionary stability each population member maximizes a suitable payoff. However, care must be taken when relatives affect each other, since relatives share genes. We describe the analysis of games between relatives in Section 4.5.

1.5.1 Challenges

The approach taken in the standard game-theoretic models often rests on idealized assumptions. This is important and helpful in providing easily understandable and clear predictions, but biologists might rely on models without careful examination of the consequences of the assumptions and limitations of the models. We believe the ideas used in the field need to be re-evaluated and updated. In particular, game theory needs to be richer, and much of the remainder of the book is concerned with ways in which we believe it should be enriched. Here we outline some of these ways.

Co-evolution. There is a tendency to consider the evolution of a single trait keeping other traits fixed. It is often the case, however, that another trait strongly interacts with the focal trait (Chapter 6). Co-evolution of the two traits can for instance bring about

disruptive selection causing two morphs to coexist or giving rise to two evolutionarily stable outcomes. These insights might not be gained if the traits are considered singly.

Variation. There is usually considerable variation in natural populations. Many existing game-theoretical models ignore this and assume all individuals in a given role are the same. However, variation affects the degree to which individuals should be choosy over who they interact with and the value of expending effort observing others. Variation thus leads to co-evolution of the focal trait with either choosiness or social sensitivity. These issues are crucial for effects of reputation and the functioning of biological markets (Chapter 7). Variation is also crucial for phenomena such as signalling. We believe that models need to be explicit about the type and amount of variation present, and to explore the consequences.

Process. In many game-theoretical models of the interaction of two individuals each chooses its action without knowledge of the choice of their partner. Furthermore, neither alters its choice once the action of its partner has been revealed (a simultaneous or one-shot game). In reality most interactions involve individuals responding to each other. The final outcome of the interaction can then be thought of as resulting from some interaction process. The outcome can depend strongly on the nature of this process (Chapter 8). Since partners vary, the interaction often involves gaining information about the abilities and intentions of other individuals, emphasizing the importance of variation and learning (Chapters 5 and 8).

Timescales. Games can occur over different timescales. Many, such as a contest over a food item, are concerned with a brief part of the whole lifetime of the contestants. We may then ask to what extent it is possible to isolate the contest from the rest of the lifetime. We often isolate games by assuming suitable payoffs, but as we discuss in Chapter 9 this is not always possible and a more holistic view of an organism's lifetime is necessary if our model is to be consistent. Some games, such as those involving the inheritance by males of a trait that is attractive to females, occur over more than one generation (Chapter 10).

Behavioural mechanisms and large worlds. The number of subtly distinct circumstances encountered in the real world is vast. Evolution shapes behavioural mechanisms that must deal with this vast range. It is not realistic to suppose that this produces strategies that respond flexibly and appropriately to every possible circumstance. The behavioural mechanism needed to implement such a strategy could probably not evolve in the first place, and even if it did it would involve so much neuronal machinery that the maintenance costs of the machinery would outweigh the benefit. For this reason we expect the evolution of strategies that are implemented by psychological and physiological mechanisms of limited complexity, which perform well on average but may not be exactly optimal in any circumstance. Most game-theory models take what can be called a small-worlds perspective: they deal with optimal strategies in simple situations. The real world is large (Chapter 5) and models need to reflect this if we are to make more realistic predictions. To do so models need to explicitly consider psychological and physiological mechanisms, accepting that these are not exactly optimal.

To elaborate on this point, game-theory models in biology have traditionally not assumed particular behavioural mechanisms, but rather assumed that organisms can use any kind of information available to them in an optimal manner. Such models are certainly valuable when they can be achieved, because they allow a full accounting for the action of natural selection. However, limiting oneself to such models has two important drawbacks. One drawback is that these models can easily become too challenging to achieve, in practice preventing modellers from analysing many important situations where interacting individuals have partial information about each other. The difficulty is caused by the complexity of a complete representation of an individual's information about its world, including its social environment (technically, a complete representation would involve probability distributions in high-dimensional spaces). The other drawback is, as mentioned, that real organisms might not use information in an unconstrained manner, but rather rely on particular behavioural mechanisms. Models that do not allow for these mechanisms might then provide a weaker connection between model results and empirical observation.

It is difficult for theory to *a priori* predict the type of mechanism that will evolve, partly because many mechanisms can work well, but also because the effect of phylogeny is important. For most purposes we suggest that we should base models on observed mechanisms, such as known mechanisms for learning. We can then investigate how parameters of these mechanisms might be tuned in different environments. We can also ask how limitations of these mechanisms might change our predictions of observed behaviour, compared with the possibly Panglossian predictions of small-worlds models (e.g. Sections 8.5 and 11.2).

Because the world is complex and unpredictable we might expect learning to be important. Individuals engaging in a game will often enter the game not knowing the abilities of themselves or others, and may not know the game structure or the payoffs. Learning in games has hitherto received insufficient attention in biology. In Chapters 5 and 8 we explore the effect of a mechanism called actor–critic learning that arose in machine learning, but has links with known psychological mechanisms and neuroscience. For instance, we use this approach to model social dominance and individual recognition in groups (Section 8.6), which are phenomena that traditionally have been difficult for game theory to handle. In doing so we hope to promote more research on realistic learning mechanisms.

In general, we argue that game theory needs to be richer, by incorporating insights and approaches from neighbouring disciplines, including ideas from neuroscience, experimental psychology, and machine learning. We present more of our views on these issues in Sections 8.8 and 11.2.

2

Central Concepts

Game theory in biology is built on certain central concepts, which we introduce here. Among these are strategies (Section 2.1), which are regarded as being genetically determined and are hence a key concept in that they are passed on to future generations. However, we need some justification of why we can deal with strategies rather than having to deal with genes directly. In biological game theory this is usually discussed in terms of the phenotypic gambit (Section 2.2). We also need some measure of the performance of a strategy and this is provided by the concept of invasion fitness (Section 2.3). As we explain, this is a measure of the per-capita rate of increase in the number of individuals following the strategy when the strategy is rare. We are then in a position to formulate a necessary condition for the resident strategy in a population to be evolutionarily stable in terms of invasion fitness (Section 2.4). This stability condition is central to the book. It can, however, be helpful in many situations to formulate the stability condition in terms of a fitness proxy instead of working directly with invasion fitness (Section 2.5). Finally, many games of interest occur over a short part of the lifetime of an individual. In Section 2.6 we describe how one can relate the actions chosen in a part of the life of a single individual to invasion fitness, which concerns the growth of a cohort of individuals over many generations.

2.1 Actions, States, and Strategies

Strategies are simply rules that specify the action chosen as a function of state. In this section we unpack this statement, clarifying what we mean by an action and what we mean by a state.

2.1.1 Actions

As the examples in Section 1.2 illustrate, when we talk about strategies and actions this is meant in the widest sense, including behaviour, growth, and the allocation of the sex of offspring. The analysis in a specific context starts by specifying the range of possible actions each organism might take in that context. Example of actions include:

- For a reproducing female, two possible actions might be to produce a son or produce a daughter.

- In a contest over a resource such as a food item or potential mate, three possible actions might be to display to the opponent, to attack the opponent, or to run away.
- For a developing individual, two possible actions might be to disperse or remain at the birth site.
- For a foraging animal, the available actions might be to search for food or to rest.

The above examples all involve a finite set of actions. In many other examples possible actions lie on a continuum, so that there is a graded response available to the organism. For example:

- For a tree, an action might specify the allocation to growth of foliage as opposed to seed production.
- For a pathogen, an action might be its level of virulence.
- For a feeding animal, an action might be its level of vigilance while consuming food.

2.1.2 States

The action chosen by an individual is often contingent on the state of the individual. The idea of a state variable is very broad—basically it is a description of some aspect of the current circumstance of an organism. Examples of state variables include:

- *Energy reserves.* For a foraging animal, we might expect the decision to rest or search for food to depend on the animal's energy reserves. Similarly, the level of vigilance might depend on energy reserves.
- *Size.* In a contest over food an individual might base its decision on whether to attack an opponent on the opponent's size relative to its own size.
- *Opponent's last move.* The action in a contest may also depend on the previous behaviour of the opponent. In this sense we regard aspects of this previous behaviour as state variables.
- *Role.* In some interactions individuals have clearly defined roles that are known to the individuals. For example, in biparental care there are two roles, male and female. When an intruder challenges the owner of a territory there are two roles, owner and intruder. When there are clearly defined roles actions may depend on role, resulting in totally different evolutionary outcomes compared with the case in which there are no role asymmetries (Section 6.2).
- *Environmental temperature.* Environmental temperature is an important state variable for many plants, for example influencing vernalization and seed germination. In some reptiles environmental temperature affects the sexes differentially during development. Some temperatures favour male offspring over female, while the reverse is true at other temperatures. This probably accounts for the fact that in some species the sex of offspring is not genetically determined; rather whether males or females develop depends on the temperature (Section 3.11).
- *Social status.* This is an important state variable in many social species.

In general any aspect of physiology or any information that the organism possesses can act as a state variable. The action taken by an individual could potentially depend

on the combination of values of its state variables. The action could also depend on time of day or year.

2.1.3 Strategies

Strategies are genetically determined rules that specify the action taken as a function of the state of the organism. For example, in a contest between two individuals over a resource, a strategy might specify whether to escalate a contest, display to the opponent, or retreat depending on perceived differences in fighting ability and the level of aggression already shown by the opponent. In temperature-dependent sex determination in reptiles, a strategy might be: if the temperature is greater than 30°C then develop as a female, if less than 30°C develop as a male. For a foraging animal a strategy might specify whether to forage or rest, depending on current energy reserves, the amount of food in the gut, and time of day. For a growing animal, a strategy might specify the relationship between environmental conditions and the time at which the animal switches from growth to reproduction. In this and other life-history contexts, such a reaction norm is just a strategy.

Genes code for strategies, but it should not be thought that there is a simple relationship between them. Models often assume that there are a few genetic loci that determine a strategy, and a simple relationship between the alleles at these loci and the strategy the individual employs. We should not regard such genotype–phenotype maps as being realistic. The phenotype is typically the product of developmental processes that are affected by many genes in a manner that may not be straightforward. However, this does not mean that the simple models should be dismissed (see the next section).

Strategies are the basic units through which selection acts. The reproductive success of an individual will typically depend on the strategy that the individual employs. The differential reproductive success of strategies affects the population frequency of alleles at genetic loci that code for the class of strategies that are possible.

In modelling a specific biological scenario it is important to keep the distinction between strategies and actions, being clear as to exactly what is regarded as a strategy. Strategies specify the action to take in all possible circumstances. Thus an organism following a given strategy may or may not take a specific action; it will depend on the state it finds itself in. The distinction needs to be kept clear because it concerns which characteristics are genetically determined and hence the level at which selection acts. To illustrate this, suppose that we are interested in the level of parental effort that evolves in a species in which both parents care for their common young. We might then regard the effort of a parent as her/his action. At the simplest level a strategy specifies a pair of efforts: the effort of an individual if female and the effort if male. Since strategies are genetically determined, this would be equivalent to assuming that male effort and female effort are genetically determined. In many cases this might be too simplistic, since there is evidence in many species that each parent adjusts her/his effort in response to the parental effort of the partner. We might instead regard a strategy as a sex-specific pair of negotiation rules where each rule specifies how own

effort adjusts in response to the effort of the partner. From this perspective it is the negotiation rules that are genetically determined, not the efforts. Efforts are the result of the negotiation process. As we will see (Section 3.4 and 8.4), the predicted parental efforts differ between these two perspectives. When individuals learn which actions to take (Chapters 5 and 8), the behavioural mechanisms that guide learning correspond to strategies.

The behavioural response of an animal to stimuli is often influenced by the previous experience of the animal. So for example, the way in which an animal deals with a stressful situation as an adult may depend on its experiences when young. This might seem to contradict the notion that strategies are genetically determined. However, just as for the parental care case, it is important to analyse behaviour at the correct level. For the stress example, a strategy is not the response to stress as an adult but the rule that specifies how response to stress as an adult depends on early life experience. Similarly, there has recently been an interest in non-genetic inheritance, in particular, in epigenetic marks that are passed from mother to offspring and can affect the adult phenotype. In this context, a strategy is a genetically determined rule that specifies how the adult phenotype depends on the epigenetic marks and other influences.

2.2 The Phenotypic Gambit

To specify the evolutionary dynamics of strategies, one needs to make assumptions about the genetic determination. The approach in game theory is to first examine very simple dynamics, and to add specific details about the genetic determination of strategies only to the extent that this is important. A simple assumption is to think of each strategy as corresponding to a genotype. This could for instance be implemented as asexual inheritance of the strategy. Using the approach need not, however, imply a restriction to this single case, but rather that one grants that evolutionary endpoints can be the same for different kinds of genetic determination of strategies, including the very common case of diploid multilocus genetics. Alan Grafen coined the term ‘phenotypic gambit’ for this general idea (Grafen, 1984), and it is very widely used.

As we will describe, a fitness proxy is a function that assigns a measure of performance to each strategy (Section 2.5). If there is no frequency dependence one can think of a landscape of possible strategies, with the fitness proxy measuring the local height. The idea of the gambit is that evolution tends to increase this performance measure, so that the evolutionary dynamic is some form of hill climbing and that endpoints are strategies that maximize the fitness proxy. In other words, evolution takes the population onto a local peak. With frequency dependence, one needs to take into account that the fitness landscape depends on the strategies of population members, and so changes as these strategies change. One then has the idea that evolution takes the population to a hilltop in a landscape that is generated by the strategies in the population, when the population sits at this hilltop.

Even if some aspects of genetics are needed to describe a phenomenon, the basic idea of the gambit can still be used. For instance, for populations with two sexes and

where each individual has one mother and one father, it is possible to work out fitness functions for sex allocation strategies (see Section 3.8), and typically one finds a sex ratio of 50% females as an endpoint. The phenotypic gambit then means that one assumes that a fairly wide range of evolutionary dynamics will lead to this endpoint.

Evolutionary change towards a predicted endpoint is illustrated by a selection experiment on Atlantic silverside fish (Fig. 2.1). In this species sex is determined by a combination of the temperature during a sensitive period, implemented as a genetically determined threshold above which an individual becomes male, and temperature-insensitive sex-determining genes. The sex-determining mechanisms present from the start in the laboratory population will influence the particular sex-ratio trajectory over the generations, but with a sufficient supply of different kinds of genetic variation, it is reasonable that the end point of a 1:1 sex ratio is approached. As seen in Fig. 2.1, this is also what happened in the experiment by Conover et al. (1992).

It is thus possible to adapt the phenotypic gambit slightly by taking major qualitative features of genetic inheritance into account, for instance by incorporating that each individual has one mother and one father into game theory analyses of sex ratios. Another example is to include effects of genetic correlations between co-evolving traits (see Section 6.1). Selection that changes one trait can then lead to a correlated




Fig. 2.1 Data from a sex-ratio evolution experiment with Atlantic silverside (*Menidia menidia*) fish (Conover and Van Voorhees, 1990; Conover et al., 1992). The species has combined temperature-dependent and genetic sex determination. The figure shows the mean and 95% confidence intervals of the proportion of females in a laboratory population (redrawn from fig. 1A in Conover et al., 1992). The population was collected as embryos from South Carolina. It was kept at a high temperature of 28°C, resulting in a skewed starting sex ratio. Over the generations, the population approached a predicted sex ratio of 0.5. Illustration of Atlantic silverside by William Converse Kendall, from the Freshwater and Marine Image Bank, University of Washington.

response in another trait. A widely used approach, which is in the spirit of the phenotypic gambit, is to assume some specific but simple genetic determination, such as haploid or single-locus diploid genetics, and to limit consideration to the invasion of a rare mutant allele into a genetically monomorphic population (see next section). Using this one can, for instance, extend the analysis to games between relatives (see Section 4.5).

The gambit can fail. Even though the phenotypic gambit is a helpful idealization, it should not be taken for granted. An interesting example where it fails is the alternative male mating types in the ruff, a species of wading bird (Fig. 2.2). The male type is determined by the genotype at a single Mendelian locus, but the alleles at this locus are actually supergenes, and were originally formed through an inversion of a block of 125 genes on a chromosome (Küpper et al., 2016). This event occurred around 3.8 million years ago and was followed by modification of the supergene content, including a rare recombination event between the original and inverted alleles 500,000 years ago (Lamichhaney et al., 2016), giving rise to three alleles. During the breeding season, male homozygotes for the non-inverted allele develop an elaborate head and neck plumage and strive to defend small territories. They are referred to as ‘independents.’ Heterozygotes between this allele and the oldest inverted allele, which




Fig. 2.2 (a) Illustration of ruff plumage variation, by Jos Zwarts. Published under the Creative Commons Attribution-Share Alike 4.0 International license (CC BY-SA 4.0). (b) Survival of different morph-determining genotypes in the ruff. Observed (dark grey) and Mendelian expected (light grey) number of surviving male plus female offspring when both father and mother have a heterozygous genotype AB, where A is the allele for the ‘independent’ male morph and B is either the allele for the ‘faeder’ male morph (left side) or the ‘satellite’ male morph (right side). Data were obtained from a pedigree for a captive genotyped population, and come from table 1 in Küpper et al. (2016). Note that both the faeder and satellite alleles are lethal when homozygous (genotype BB), and this is the case both for males and females.

are referred to as ‘faeders’, have a female-like appearance and sneak copulations. Heterozygotes with the recombined allele are called ‘satellites’, and develop a head and neck breeding plumage, but less colourful than those of the independents. They do not defend territories, but associate with independents, displaying on—and increasing the attractiveness of—territories, and attempt to sneak copulations. The ruff male breeding plumage is quite variable (Fig. 2.2a), but differs between the genetic types.

As an accidental consequence of the original inversion, a gene needed for successful development was disrupted, making homozygotes for inverted alleles, either faeder or satellite, lethal (Fig. 2.2b). The heterozygotes also suffer to some extent from the maladaptive consequences of the inversion, but male heterozygotes compensate by a high mating success, possibly because of high success in sperm competition (Küpper et al., 2016). The upshot is that the population frequencies of the ruff mating types are not at a Nash equilibrium (Section 2.4), and the genetic underpinnings have clearly maladaptive features. It might seem surprising that the long evolutionary history of the mating system has not removed these maladaptive features, but the empirical evidence for their origin and current existence is strong. It is worth noting that the ruff is a widespread and successful bird species. For the applicability of game theory in biology, examples such as this should be uncommon, which seems to be the case. Finally, even though inversions and supergenes may imply genetic constraints, they can still play important roles in evolution.

2.3 Invasion Fitness

Consider a large population in which almost all individuals employ a given strategy x . As a shorthand for this statement we will say that x is the resident strategy. We assume that for every resident strategy x , the population settles down to a unique demographic equilibrium (that can depend on x). Here, by the term ‘demographic equilibrium’ we mean that the population size has settled down and the proportion of population members in each possible state has also settled down. This equilibrium then sets the background in which we consider the fate of any mutants.

Now suppose that we introduce one or a few individuals following an alternative strategy x' into a population with resident strategy x , where the resident population has reached its demographic equilibrium. Since we have in mind that strategy x' might have arisen through mutation from the resident strategy, we refer to x' as a mutant strategy. We will define the invasion fitness of the mutant strategy $\lambda(x', x)$ as an appropriate measure of the per-capita growth in the number of individuals that follow the mutant strategy, while mutants are still rare compared with residents. For convenience, we measure this growth rate as the per-capita annual increase.

The precise definition of this ‘appropriate measure’ depends on circumstance. Before defining it we need to consider how stochastic events affect population members. At one extreme all population members might experience the same environmental factors such as weather, population size, and the size of prey populations.

We refer to random fluctuations in these factors as environmental stochasticity. Economists refer to this form of stochasticity as aggregate risk since all population members experience the same fluctuations. At the other extreme, different population members might find different numbers of prey items due to good and bad luck when foraging. Fluctuations that affect different population members independently (for given environmental conditions) are referred to as demographic stochasticity, which economists would call idiosyncratic risk.

To study invasion fitness for a mutant strategy we census the mutant population at time points with a suitable spacing, for instance annually at the start of each year, which we now use for concreteness. Let $N(t)$ denote the total number of x' mutants present at the start of year t . We are interested in how this number changes over time. In the very simplest case we assume the following conditions:

1. The population is asexual.
2. There are discrete, non-overlapping generations with a generation time of 1 year. Thus those individuals that are born at census time t reproduce and die over the course of the year, leaving surviving offspring at census time $t + 1$.
3. All mutants are in the same state at each annual census time.
4. There is no environmental stochasticity, so that all years have the same environmental conditions. This does not mean that there are no seasonal effects, but if there are, they are the same each year.

For this scenario consider a mutant individual following strategy x' that is present at the census time in year t . We can refer to the surviving offspring left at the next annual census time by this individual as its recruits. The number of recruits depends on the good and bad luck that the mutant experiences. Let $\lambda(x', x)$ be the mean number of recruits that also follow the strategy x' , where in forming the mean we are averaging over demographic stochasticity. Then if there are $N(t)$ mutants following strategy x' in year t there will be approximately $N(t + 1) = \lambda(x', x)N(t)$ mutants in year $t + 1$. This approximation depends on $N(t)$ being sufficiently large so that we can use the law of large numbers to average over the (independent) good and bad luck experienced by different mutant individuals. Thus $\lambda(x', x)$ is the per-capita rate of increase in mutant numbers when mutants have become common enough to average over demographic stochasticity but are still rare compared with the resident population. We define $\lambda(x', x)$ to be the invasion fitness of the mutant. When considering whether a mutant strategy x' can invade into a resident x we should note that there is an initial phase where there are few mutant individuals. In this initial phase the mutant might go extinct because of demographic stochasticity, even if the number of mutants tends to increase on average. Our assumption is that there is some probability for mutant numbers to become large enough so that our approximation based on averaging over demographic stochasticity is valid. This initial randomness might thus require several mutants appearing before our invasion analysis can be applied. Assuming that this is so, if $\lambda(x', x) < 1$ mutant numbers will not increase further, while if $\lambda(x', x) > 1$ mutant numbers increase while rare. The background environment in which the cohort of mutant individuals live is determined by the resident strategy.

The stipulation that mutants should be rare compared with residents ensures that the background remains approximately constant when the growth rate is measured. If mutants become common, we would expect their per-capita rate of increase to change due to frequency- and density-dependent effects.

In the above motivation for invasion fitness it was important that all recruits were in the same state (condition 3); i.e. the population is unstructured. Suppose that this is not the case. For example, suppose that individuals differ in their energy reserves or parasite load at an annual census time. Then different x' mutants would have different potentials to leave surviving offspring. Furthermore, the state of any recruit might be correlated with the state of its parent. Under such circumstances it is not sufficient just to count numbers of recruits when defining invasion fitness. Instead, one needs to consider projection matrices that specify how parental state affects the number and state of descendants left the following year (Chapter 10). Invasion fitness is then taken to be the per-capita rate of increase when the distribution of states of the cohort of mutants has settled down to its limiting steady-state distribution, which is the leading eigenvalue of the projection matrix (Section 10.1). For most of our applications we will not need this machinery.

Our assumption also specifies that the population is asexual (condition 1). With two sexes the population is necessarily structured at the annual census time; there are males and females. Despite this, it is often sufficient just to take the number of female recruits left by a mutant female as invasion fitness. However, in the sex-allocation game of Section 3.8 this expedient fails. In that game a female has N recruits and must decide whether these recruits should be sons or daughters. The numbers of recruits left by a male depends on the sex ratio in the resident population, and can be different from the numbers of recruits left by a female. In this case, in principle one needs the full machinery of projection matrices. However, the analysis can be simplified as it can be shown (Section 10.4) that number of grandchildren left by a female acts as a fitness proxy (Section 2.5).

We further assumed that the population has discrete non-overlapping generations (condition 2). Suppose instead that an individual can live for several years. Then the population is typically structured at an annual census time as individuals of different ages, which can have different potentials to survive and reproduce. Nevertheless, one can simplify analyses when age is the only state variable, because all new recruits (age 0) are in the same state. To do so, one defines the lifetime reproductive success (LRS) of an individual as the total number of its recruits (discounted by relatedness if the population is sexually reproducing). As we explain in Section 2.5, mean LRS is what we refer to as a fitness proxy. We use this fitness proxy to analyse the evolutionarily stable scheduling of reproduction over the lifetime of an organism in Section 10.6.

In the cases we examine in this book condition 4 always holds. When condition 4 fails to hold, so that there is environmental stochasticity, then our argument that the actual change in mutant numbers approximately equals the mean change fails. For the definition of fitness in this more general scenario, see Metz et al. (1992) and chapter 10 in Houston and McNamara (1999).

2.4 Evolutionary Endpoints

As mentioned, rather than looking at evolutionary dynamics in detail, a large part of game theory in biology focuses on the endpoints of evolutionary change. By an endpoint we have in mind a strategy (or distribution of strategies) that is stable over time. In particular, no new mutant strategy that might arise can invade into the population and change its composition. In this section we introduce the concept of a Nash equilibrium, which provides an important necessary condition for evolutionary stability.

Let x be a resident strategy. We can consider a mutant strategy that is identical to this resident strategy. The invasion fitness of this mutant, $\lambda(x, x)$, is the per-capita annual growth of the resident population. For a population at its demographic equilibrium we must have $\lambda(x, x) = 1$. Now let x' be a mutant strategy that is different from the resident strategy. In the previous section we argued that if $\lambda(x', x) > 1$ then an individual mutant has a positive probability to invade, changing the composition of the population so that x would not be evolutionarily stable. Thus a necessary condition for stability is that $\lambda(x', x) \leq \lambda(x, x) = 1$. Motivated by this we define x^* to be a Nash equilibrium strategy if

$$\lambda(x', x^*) \leq \lambda(x^*, x^*) \quad (= 1) \quad \text{for all } x'. \quad (2.1)$$

We note that invasion fitness is often measured in terms of $r = \log \lambda$, the logarithm of the relative annual growth rate. We would then have $r(x', x^*) \leq r(x^*, x^*) (= 0)$ for all x' .

As we have remarked, the Nash condition is exactly the same as that in economic game theory. However, there are some differences in the approach to game theory in biology and economics (Box 2.1).

Condition (2.1) is necessary for evolutionary stability, but not sufficient. In particular, in Section 4.1 we consider what happens when there are mutants with equal fitness to residents, introducing the idea of an Evolutionarily Stable Strategy (ESS). As we will see, the condition for a strategy to be an ESS is stronger than that required for a strategy to be a Nash equilibrium. Thus an ESS is necessarily a Nash equilibrium.

We can reformulate the Nash equilibrium condition as follows. A strategy \hat{x} is called a best response to a resident strategy x if this strategy has fitness at least as great as any other mutant. That is

$$\lambda(\hat{x}, x) = \max_{x'} \lambda(x', x). \quad (2.2)$$

Then condition (2.1) is equivalent to the condition that x^* is a best response to itself. There may be more than one best response to a given resident strategy. However, in many of the examples we analyse every resident strategy x has a unique best response, which we denote by $\hat{b}(x)$. When this is so, condition (2.1) is equivalent to the condition that $\hat{b}(x^*) = x^*$.

The above analysis is concerned with evolutionary stability for a monomorphic population; i.e. at stability resident population members all employ the same strategy.

Box 2.1 Economic versus biological game theory

In economic game theory the agents making decisions could be individuals, firms, or governments. In biology they are individuals, but these are only important in that they pass on the genes determining behaviour to future generations, so it is the genes rather than the individuals that are important in evolutionary terms. Approaches in the two areas differ in two major respects.

Payoffs. In biology the payoff of an action is an increment in invasion fitness (e.g. the change in mean lifetime number of offspring). We expect evolution to produce individuals that maximize invasion fitness by maximizing their payoffs.

In economics there are clear performance criteria, such as monetary return, in many applications. In others, where for example social standing is important, the performance criterion is not usually specified in advance. Instead each individual's utility is constructed from their observed behaviour. If behaviour satisfies certain consistency conditions, such as transitivity of choice, then one can construct a utility function such that observed behaviour maximizes expected utility. The utility of an action then acts as the payoff in a game.

In summary, in biology payoffs are normative—they specify what an organism should do to maximize invasion fitness. In economics, utilities (if they exist) are purely descriptive of what individuals actually do (Kacelnik, 2006; Houston et al., 2014).

Justification of the Nash equilibrium concept. The Nash equilibrium is also the equilibrium concept in economic games. Early justification of this centred around the rationality of players. For example, in order to argue that two players would each play their Nash equilibrium strategies it was necessary to assume not only that both were rational, but that each knew the other was rational, and each knew that the other knew each was rational. More recently, attention has switched to dynamic models of choice, where individuals experience payoffs and update their actions accordingly—the focus then is whether a particular learning rule will eventually lead to strategies being in Nash equilibrium.

In contrast, in biology the dynamic process is over generations and is dictated by natural selection. If this dynamic process converges we expect it to be to a Nash equilibrium.

It is also possible to have evolutionarily stable mixtures of strategies. The population is then said to be polymorphic.

Finally, when considering the invasion fitness of a mutant x' in a resident x^* population, it is helpful to think of the mutant as a modification of the resident strategy. Potentially, mutants at different genetic loci can modify the resident strategy, so it is not necessary that there is a single gene or allele coding for each strategy.

2.5 Fitness Proxies

It is often more convenient to work with a fitness proxy rather than working directly with invasion fitness $\lambda(x', x)$. We will refer to a function $W(x', x)$ of x' and x as a fitness proxy if

$$W(x', x) < W(x, x) \iff \lambda(x', x) < \lambda(x, x) \quad \text{for all } x', x, \quad (2.3)$$