

Springer Theses

Recognizing Outstanding Ph.D. Research

Peter Ashcroft

The Statistical Physics of Fixation and Equilibration in Individual-Based Models

Springer Theses

Recognizing Outstanding Ph.D. Research

Aims and Scope

The series “Springer Theses” brings together a selection of the very best Ph.D. theses from around the world and across the physical sciences. Nominated and endorsed by two recognized specialists, each published volume has been selected for its scientific excellence and the high impact of its contents for the pertinent field of research. For greater accessibility to non-specialists, the published versions include an extended introduction, as well as a foreword by the student’s supervisor explaining the special relevance of the work for the field. As a whole, the series will provide a valuable resource both for newcomers to the research fields described, and for other scientists seeking detailed background information on special questions. Finally, it provides an accredited documentation of the valuable contributions made by today’s younger generation of scientists.

Theses are accepted into the series by invited nomination only and must fulfill all of the following criteria

- They must be written in good English.
- The topic should fall within the confines of Chemistry, Physics, Earth Sciences, Engineering and related interdisciplinary fields such as Materials, Nanoscience, Chemical Engineering, Complex Systems and Biophysics.
- The work reported in the thesis must represent a significant scientific advance.
- If the thesis includes previously published material, permission to reproduce this must be gained from the respective copyright holder.
- They must have been examined and passed during the 12 months prior to nomination.
- Each thesis should include a foreword by the supervisor outlining the significance of its content.
- The theses should have a clearly defined structure including an introduction accessible to scientists not expert in that particular field.

More information about this series at <http://www.springer.com/series/8790>

Peter Ashcroft

The Statistical Physics of Fixation and Equilibration in Individual-Based Models

Doctoral Thesis accepted by
The University of Manchester, UK

Author

Dr. Peter Ashcroft
Institute for Integrative Biology
ETH Zürich
Zürich
Switzerland

Supervisor

Dr. Tobias Galla
School of Physics and Astronomy
The University of Manchester
Manchester
UK

ISSN 2190-5053

Springer Theses

ISBN 978-3-319-41212-2

DOI 10.1007/978-3-319-41213-9

ISSN 2190-5061 (electronic)

ISBN 978-3-319-41213-9 (eBook)

Library of Congress Control Number: 2016943809

© Springer International Publishing Switzerland 2016

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by Springer Nature

The registered company is Springer International Publishing AG Switzerland

Parts of this thesis have been published in the following journal articles:

P. Ashcroft, P. M. Altrock, and T. Galla, Fixation in finite populations evolving in fluctuating environments. *J. R. Soc. Interface* **11**, 20140663 (2014).

P. Ashcroft, F. Michor, and T. Galla, Stochastic tunneling and metastable states during the somatic evolution of cancer. *Genetics* **199**, 1213 (2015).

P. Ashcroft, A. Traulsen, and T. Galla, When the mean is not enough: Calculating fixation time distributions in birth–death processes. *Phys. Rev. E* **92**, 042154 (2015).

Supervisor's Foreword

A physicist's view of the world used to be rather deterministic 200 years ago. If the state of all particles and the forces acting between them were known, Laplace hypothesised, a 'being' vast enough to project forward these laws could calculate all future states of the universe. This concept of a Laplacian daemon is of course hopelessly out of date. The discovery of quantum mechanics has fundamentally changed our view of the world; the way we think about determinism and predictability is rather different than at the times of Newton and later Laplace. Quantum physics describes randomness which is (as far as we know) intrinsic to nature, and which cannot be overcome, no matter how hard we try.

Perhaps as important as such aleatoric uncertainty is a second source of randomness in mathematical models of the world around us. Towards the late nineteenth century revolutionary ideas about how one might describe systems with large numbers of interacting particles were developed by heroic figures such as Boltzmann and Gibbs. The theory of 'statistical mechanics' was born. These concepts are founded on the idea that the detailed trajectories of each and every particle in a large system are not all that interesting. Instead it is the behaviour of the system as a whole that we care about. The logical consequence is to study ensembles of particles, and their 'statistics'. Unlike in Newtonian physics we no longer ask: Where is this particular particle going to be at a later time? Instead we ask: If the initial distribution of particles is this, what is the probability to find a given particle in a certain area of space, or with a given speed?

This leads to stochastic descriptions of the laws of physics, the equations governing the dynamics are now subject to noise. This is so-called 'epistemic noise', it originates from the way we model those ensembles of particles. Leaving quantum physics aside, epistemic uncertainty could be eliminated by making a more detailed model and including all forces and interactions in a Laplacian sense. But as statistical physicists we decide not to, because these are not the relevant questions. What is relevant is how global large-scale behaviour emerges from microscopic interaction, not the microscopic trajectories as such.

In order to do this a whole new theory—the theory of stochastic processes—had to be invented. Much of the twentieth century was spent on trying to understand and classify the equilibrium states interacting particle systems reach in the long term. This work is now largely complete, and the focus has moved to systems out of equilibrium. These are systems which do not settle down, they are subject to driving, fluxes, and coupling to the external surroundings. No coherent theory exists for the physics far from equilibrium, but at the same time many pressing challenges rely on progress in this field. This includes turbulence, plasma fusion, active matter, quantum materials, and most notably the physics of life. Biology is inherently out of equilibrium and based on transport of nutrients, energy, the absorption of light, sudden changes, large deviations, the dynamics of evolution and changing environments. It is no surprise that physicists have been able to make remarkable contributions, and that ideas from statistical physics and the theory of stochastic processes have delivered important advances.

Peter joined this adventure in 2012, his thesis focuses on the dynamics of fixation in models of interacting individuals. Peter has investigated several problems at the boundary of theoretical physics and biology. The thesis contains the study of an evolutionary model of populations in switching environments, relevant for example for antibiotic treatment in colonies of bacteria. He has also analysed metastable states in a model of cancer initiation, and the relation of so-called mixing times and the dynamics of fixation in birth–death processes. Chapter 6 of his thesis presents a pedagogical account of the so-called WKB method, a technique from semi-classical physics used to study phenomena including epidemics, ecosystems and, in Peter's thesis, models of cancer evolution.

While I mention the words ‘cancer’, ‘bacteria’ and ‘evolution’ we should be clear: this is a thesis in statistical physics. It contains mathematics and long equations, things are complicated and subtle. The problems Peter looks at are motivated in biology, but the true beauty of this thesis is in the beauty of the underlying mathematical structures and the theoretical concepts and ideas used to unveil them.

Santiago de Compostela, Spain
April 2016

Dr. Tobias Galla

Abstract

Individual-based models have been applied to study a broad spectrum of problems across multiple disciplines, such as the spread of epidemics or the outcome of social dilemma. They are used to investigate the macroscopic effects that arise from the microscopic dynamics of interacting individuals. Fixation describes the taking over of the population by a single type of individual or species. It is a prominent feature in the field of population genetics, which interprets many biological scenarios of evolution. Equilibration describes the process of reaching a heterogeneous steady state. In this thesis we analyse these macroscopic features through techniques derived from statistical physics and the theory of stochastic processes.

Birth–death processes are used to describe the interaction of two types of individual in a population, such as competing strains of bacteria. These interactions are often specified using the framework of evolutionary game theory. The environment in which the population evolves can have a crucial impact on selection. In systems where the environment switches between multiple states we develop a general theory to calculate the fixation time statistics of a mutant individual in a population of wild-types, as well as the stationary distributions when mutations are present in the dynamics. In some birth–death processes, and in particular those described by evolutionary game theory, the mean fixation time contains only limited information. By diagonalising the master equation that describes the process, we are able to obtain closed-form expressions for the complete fixation time distributions.

Individual-based models can also be used to describe the accumulation of mutations in a cell. This has important consequences for the initiation and progression of cancer. We find that such systems exhibit metastable states in the dynamics, and we can exploit the separation of timescales between relaxing to the quasi-stationary state and reaching fixation to characterise these phenomena. In this scenario we employ the WKB method to describe the population-level dynamics. Although this method has been used to describe numerous stochastic processes, a clear and coherent description is lacking in the literature. Through the use of multiple examples, including the aforementioned cancer initiation model, we carefully explain the multitude of constructs and equations that result from the application of this method.

The analytical characterisation of the evolutionary dynamics that are observed in these stochastic processes has resulted in a greater understanding of fixation and equilibration. This thesis promotes the benefits of analytical, or even semi-analytical methods, and on a more general level contributes towards a more complete understanding of evolutionary processes.

Acknowledgments

I would first like to thank my supervisor Tobias Galla, whose infectious enthusiasm and invaluable guidance have ensured this experience has been wholly fulfilling. My gratitudes are also with Alan McKane for maintaining a stimulating environment for me to work in, and with my collaborators Philipp Altrock, Franziska Michor, and Arne Traulsen for the support they provided and the opportunities they presented. A special mention is reserved for Yen Ting Lin for his much-appreciated feedback and constructive criticisms.

Thanks to all the occupants of office 7.26 who are too numerous to name, but you know who you are, and to the postdocs who have been and gone during my Ph.D.

Thanks to my niece Mia, whose timely birth in the middle of writing this thesis provided some much-needed distraction, and to the rest of my family for their support.

Finally, I want to give a really big thank you to Stacey, who has consistently put up with my unending drivel about the wonders of science. Without your support and understanding I would not be who I am today. Here's to the next adventure!

Contents

1	Introduction	1
	References	7
2	Technical Background	11
2.1	Definition: Stochastic Processes	11
2.2	Chapman–Kolmogorov and Master Equations	12
2.3	Birth–Death Process	14
2.4	Fixation Probability and Mean Fixation Times	16
2.5	Equilibration	23
2.6	Deterministic Dynamics and Stability Analysis	26
2.7	Continuous State-Space Approximations	29
2.8	Evolutionary Game Theory	30
2.9	Numerical Simulation Methods	34
	References	36
3	Finite Populations in Switching Environments	39
3.1	Introduction	39
3.2	Model	41
3.3	Mathematical Framework	43
3.4	Switching Between Two Games	49
3.5	Mutation–Selection Equilibria	56
3.6	Summary	60
	References	62
4	Fixation Time Distributions in Birth–Death Processes	63
4.1	Introduction	63
4.2	Model	64
4.3	Mathematical Framework	65
4.4	Physical Interpretation	72

4.5	Application to Evolutionary Games	77
4.6	Equilibration Processes in Systems with Mutation	78
4.7	Efficiency of the Method	86
4.8	Summary	87
	References	88
5	Metastable States in a Model of Cancer Initiation	91
5.1	Introduction	91
5.2	Model	94
5.3	Deterministic Analysis	96
5.4	Types of Stochastic Behaviour	101
5.5	WKB Analysis	105
5.6	Summary	122
	References	124
6	The WKB Method: A User-Guide	127
6.1	Introduction	127
6.2	Toy Model with One Degree of Freedom	129
6.3	The WKB Method in One Dimension	134
6.4	Four-State Toy Model	140
6.5	The WKB Method in Higher Dimensions	142
6.6	Examples	152
6.7	Summary	155
	References	156
7	Conclusions	159
7.1	General Discussion	159
7.2	Summary of Results	160
7.3	Outlook	162
	References	164

About the Author



Dr. Peter Ashcroft graduated as a Master of Mathematics and Physics from the University of Manchester in 2012. He then studied for his Ph.D. in theoretical physics under the supervision of Dr. Tobias Galla in Manchester. During his Ph.D., Peter was especially interested in the theory of stochastic processes and in the mathematical modelling of biological phenomena. He completed his Ph.D. in 2015, and currently is a postdoc in theoretical biology at ETH Zürich where, amongst other projects, he is investigating the dynamics of blood formation and disease.

Chapter 1

Introduction

Over the next 200 or so pages I will explore how tools and concepts developed within theoretical physics can be applied to problems in other sciences. Although more emphasis in this thesis will be directed to biological applications, the successes of this field can also be seen in social science [1, 2], economics [3, 4], and many other disciplines where so-called complex systems are a prominent feature.

My motivation for working in this area is the freedom you have to explore these numerous disciplines, and the satisfaction that arises from solving a long-standing problem by approaching it from an unconventional point of view. Interactions with academics from these various backgrounds has provided hours of intellectual conversation and brainstorming that have greatly enhanced my knowledge of the world outside of physics. But ultimately the main reason for joining this area, and choosing to continue my career in this field, is because the analysis is fun! The benefits of the approaches I use lie not only in their predictive power, but they are enjoyable, satisfying, stimulating and infuriating in equal measures.

The success of theoretical physics across multiple disciplines comes from its ability to break down objects to their fundamental constituents. Analysis of the inner workings then allows the practitioner to obtain a more complete understanding of the world. An experimentalist works with the real-world system, or a synthetic *in vitro* analogue. Their understanding of this system is achieved through the collection and analysis of data. Theorists, however, obtain an understanding by considering a representation of the real-world system, which I will refer to as a model.

For biological systems an exact model representation is often impossible due to the inherent complexity of many interacting entities. If a model is almost as complicated as the experimental system, it will be just as intractable. In the end you would have the same data set, but generated *in silico*, and no new insight or understanding will have been gained. As the level of abstraction from the real world increases, so does the level of tractability. The balance between accuracy and tractability is a choice to be made by the modeller. In the case of this thesis, Occam's Razor prevails;

I will focus on the simplest models which reproduce observed behaviours, but can be applied to a wide range of problems. These models can highlight the underlying mechanisms that result in the observed phenomenon, something that may not be immediately obvious from simply conducting an experiment.

One of the most profound examples of this in the biosciences is the explanation of the regular structures on the coats of animals [5]. The colouration was known to be caused by melanin in the skin, but there was no explanation for the origin of the pattern of this colouration in animals such as zebras and leopards. The seminal work of Alan Turing (1912–1954) provided part of the answer. Turing proposed that diffusive chemicals can settle into a stable, spatially-inhomogeneous state through the excitation of the now-called Turing instability [6]. Although the true mechanism is more complex than the idea proposed by Turing [7], the same basic principles were applied to reproduce observed animal coat patterns [8].

The class of systems in which my interest lies is not the continuous reaction–diffusion systems as studied by Turing, but systems that contain a finite number of discrete, interacting ‘particles’ or individuals. Such systems are ubiquitous in nature, where particles could represent proteins, molecules, cells, bacteria, animals or people. The dynamics of the particles can be governed by events such as production (birth), degradation (death), predation or infection, to name but a few. The discreteness of the particles, and the nature of the dynamics, are responsible for the observed stochasticity; that is, there is an intrinsic source of randomness in these systems, often referred to as demographic noise.

The discreteness of the particles, and with it the intrinsic stochasticity, is retained when modelling these systems. However, information about the behaviour of every individual particle is not necessary. Instead, the simplifying assumption that two particles of the same type are indistinguishable is made. The behaviour of the system can then be described by the statistics of the group of particles. This procedure is the basis of statistical mechanics, and the approach is poetically summarised by James Clerk Maxwell (1831–1879):

And here I wish to point out that, in adopting this statistical method of considering the average number of groups of molecules selected according to their velocities, we have abandoned the strict kinetic method of tracing the exact circumstances of each individual molecule in all its encounters.

It is therefore possible that we may arrive at results which, though they fairly represent the facts as long as we are supposed to deal with a gas in mass, would cease to be applicable if our faculties and instruments were so sharpened that we could detect and lay hold of each molecule and trace it through all its course.

James Clerk Maxwell, *The Theory of Heat* [9].

Here Maxwell is referring to the original derivation of the Maxwell–Boltzmann distribution, which describes the distribution of speeds of molecules of a contained ideal gas [10].¹

¹Ludwig Boltzmann (1844–1906) later derived this result from the kinetic theory of gases [11].

To model the particles in the discrete systems *in mass* they are treated like molecules of a gas. The interactions then take a form which is similar to that of chemical reactions. These reactions are dependent on the number of reactants (molecules) available and the rate at which the gas molecules interact [12]. These models are referred to as individual-based models, and they have been applied to study epidemic outbreaks [13], social dilemma [14], predator–prey interaction [15], and the list can go on and on. This thesis, however, is not dedicated to a particular system or application. Instead I will investigate particular phenomena that are observed in a variety of stochastic systems. These are:

Fixation: The process of a single type of individual taking over the whole population.

The term originates from the field of population genetics, where the fixation of an allele was a central topic [16–19]. In this case fixation occurs when all other alleles are irreversibly lost from the gene pool, and only a single *fixed* allele remains. The terminology is now used outside of population genetics and the study of gene frequencies, for example to describe the eradication of a disease or reaching a social consensus.

Equilibration: The process of reaching a stable stationary state. If fixation is not possible in a system, as is the case if individuals can change their type stochastically, then the success of a type of individual is no longer characterised by the probability that it takes over the population. Instead success can be measured by its relative concentration at long times. This is described by the stationary probability distribution. The time to approach this stable state is also of interest.

These two effects are closely linked; if a system fixates then no more dynamics can occur and hence the fixated state is stationary. They are also related if fixation takes a very long time, such that the system can initially relax into a quasi-stationary state before fixation occurs. These links will be investigated closely in Chaps. 4 and 5.


A concrete understanding of the effects of fixation and equilibration, and the interplay between them, will greatly contribute to our understanding of the process of evolution. This field of investigating evolution through mathematical approaches has been dubbed evolutionary dynamics, and it describes the change of populations over time subject to spontaneous mutation, selection, and random events [14, 20]. Different types of individual in the population, which we will sometimes call phenotypes in line with the biological literature, can emerge spontaneously by mutation, i.e. through errors during reproduction of the pre-existing *wild-types*. In many cases, wild-type and mutant individuals are characterised by heritable differences in behavioural traits or strategies [14]. Selection acts on different (pheno)types and their associated traits to change the population composition.

One of the great successes of evolutionary dynamics is the quantitative analysis of cancer, which is a genetic disease and according to Cancer Research UK, “*1 in 2 people in the UK born after 1960 will be diagnosed with some form of cancer during their lifetime*” [21]. Mathematical investigations have contributed profoundly to our understanding of “*the emperor of all maladies*” [22]. Numerous studies throughout the 20th century have addressed the kinetics of cancer initiation and progression

[23–28]. In Ref. [23], it was first proposed that “*several successive mutations in the same cell [...] would be necessary [for cancer to initiate]*”. Empirical observations of mortality rates across a range of cancer types agreed with this hypothesis [24]. For some varieties of cancer it was shown by Alfred Knudson (1922–) that tumours can be induced by as few as two mutations, corresponding to the inactivation of both copies of a specific tumour suppressor gene (TSG) [26]. The data that confirmed this hypothesis is presented in Fig. 1.1. This is data for the diagnosis of tumours, or retinoblastomas, in the eyes of children. Knudson hypothesised that if the tumours required two mutations, we would observe a quadratic incidence rate. However, if the child had inherited a defective gene, the incidence curve should be linear and there is a much larger probability that the cancer will be present in both eyes, which is referred to as bilateral. The data clearly favours Knudson’s interpretation, and this is the celebrated *two-hit hypothesis* [26].

The age of stochastic modelling of cancer initiation began in earnest with the introduction of the branching process, as shown in Fig. 1.2 [28]. Similar models have been used extensively to describe various aspects of carcinogenesis [29, 30],

Fig. 1.1 Fraction of cases of retinoblastoma not yet diagnosed as a function of the children’s age. The one-hit (bilateral) curve is $\log S = -t/30$, and the two-hit (unilateral) curve is $\log S = 4 \times 10^{-5}t^2$, where S is the fraction of cases not diagnosed and t is the children’s age in months. This figure is from Ref. [26]






Fig. 1.2 The branching process model of cancer initiation introduced in Ref. [28]. On the *left* is shown the three possible fates of an un-mutated cell, S : the cell can die (*upwards arrow* to state D), reproduce without mutation (*downwards arrow* to two S cells), or reproduce with mutation (*horizontal arrow* to one S and one I). The mutation occurs with rate μ_1 , and type- I cells harbour one mutation. In turn the type- I cells follow the same process of dying or reproducing with or without mutation. The mutation occurs at rate μ_2 and gives rise to a malignant cell M . This figure is from Ref. [28]

and the branching process itself has received significant analytical attention [31, 32]. Other stochastic models have also been investigated to describe the initiation of cancer, including the linear process which mimics the spatial structure of some tissues [33], and well-mixed Moran-type models which represent cells in a tissue of fixed size [34]. This latter class of models is investigated in Chap. 5 of this thesis, as described below.

Another prosperous area of evolutionary dynamics has arisen through the consideration of interactions between individuals. If, for example, one type of individual can produce a promoter which benefits the population, then more individuals of this type will lead to more production and a stronger population. However, this production usually comes at a cost to the producer, and the population is vulnerable to exploitation from individuals who do not produce, but still reap the reward. These individuals are *cheaters*, and this scenario is the celebrated public goods game [35]. Considering these types of interactions which are dependent on the state of the population leads to rich dynamical behaviour. The mathematical framework for handling such cases is evolutionary game theory [14, 36–39], and the rich behaviours that it predicts have been observed in experiments of biological evolution [40–43].

This thesis is dedicated to improving our understanding of the processes of fixation and equilibration through evolutionary dynamics. In the next chapter I will introduce the framework that will be used to analyse these systems of discrete, indistinguishable particles. Particular attention will be devoted to so-called birth–death processes, which are highly tractable yet ubiquitous in evolutionary dynamics [14]. These processes describe the temporal behaviour of a population containing two interacting types of individual, referred to as the wild-type and the mutant. I will also introduce evolutionary game theory in more detail as it is used throughout this thesis to illustrate the effects of fixation and equilibration.

In Chap. 3 I will extend the aforementioned birth–death process to study the impact of a changing environment on the evolution of a finite population of fixed size [44]. The rates at which the birth and death events occur are dependent on the state of the environment, which follows an independent random process. A general theory is developed to describe the fixation probability of a mutant in a population of wild-types, as well as the time taken for fixation to occur. The theory is then applied to evolutionary games for which the payoff structure varies in time. It is found, surprisingly, that the mutant can exploit the environmental noise; a dynamic environment that switches between two states can lead to a probability of fixation that is higher than in any of the individual environmental states. In this chapter I will also investigate the stationary distribution of the population when mutations are present in the dynamics, and prescribe approximations of the stationary distribution which are valid under different environmental dynamics.

The birth–death process will also be the central feature of Chap. 4. In this chapter the novelty comes not from an extension to the model, but from finding an exact solution to describe the distribution of fixation times. As the title of the associated article suggests, sometimes the mean is not enough to give a good representation of the statistics of arrival times [45]. This may be the case if the distribution is broad and skewed. The distribution can be expressed in terms of the spectrum of the birth–death process, and the analysis leads to different representations as forward-only processes in eigenspace. These allow efficient sampling of fixation time distributions and will be a powerful tool to use in model-reduction schemes. Again evolutionary games are considered as an exemplary application. In this chapter I will also highlight the median fixation time as a possible analogue of the time to stationarity in systems with small mutation rates and no absorbing states, whereas the mean fixation time has no such interpretation. This provides a crucial link between the effects of fixation and equilibration.

In Chap. 5 I will leave behind evolutionary games and turn my attention to the initiation of cancer. This work follows on from numerous computational and mathematical investigations of the phenomenon of stochastic tunnelling, where an intermediate mutant in a sequence does not reach fixation in a population before generating a double mutant [34, 46–51]. The field of stochastic tunnelling still lacked a comprehensive analytical description, since theoretical predictions of fixation times are only available for cases in which the second mutant is advantageous. The starting point for the investigation is the same stochastic model as used in these previous studies. By systematically analysing the deterministic dynamics of infinite populations, the parameter regimes captured by existing approaches were shown to be qualitatively different to those not captured. The analysis reveals the existence of quasi-equilibria when the final mutant is not the most advantageous in the sequence. The escape from these states is driven by the intrinsic noise, and the location of these states affects the probability of tunnelling occurring. Existing methods no longer apply in these regimes. Instead it is the escape from the quasi-equilibria that is the key bottleneck; fixation is no longer limited by the emergence of a successful mutant lineage. In these parameter regimes I employ the Wentzel–Kramers–Brillouin (WKB) method from mathematical physics to compute the time to fixation.

In Chap. 5 the WKB approach was used as an ‘off-the-shelf’ tool; I primarily used the method discussed in Ref. [52] to tackle the problem at hand. In Chap. 6 I take a closer look at the WKB method in the context of stochastic systems. It has previously been used to describe a variety of systems, including: the evolutionary dynamics of coexisting bacteria [53], predator–prey systems [54], epidemic models [55–59], and evolutionary games [60]. However, little attention is devoted to really understanding and explaining the method used. In this chapter I will discuss the origins of the WKB method, explain the terminology used throughout the literature, make connections with the related field of large deviations theory, and highlight what makes this method superior to other approaches under specific conditions. A central feature of the method is the construction of a landscape, or ‘quasi-potential’ for stochastic systems. Waddington’s epigenetic landscape [61] is a primitive example of such a construct, but similar pictures can provide intuition and quantitative predictions for a range of stochastic models. The power of this method will be illustrated by considering various applications, including the cancer initiation model discussed in Chap. 5.

Finally I will summarise the findings of this work and discuss the avenues of future research that this thesis has promoted in Chap. 7.

References

1. T.C. Schelling, Dynamic models of segregation. *J. Math. Sociol.* **1**, 143 (1971)
2. D. Helbing, B. Tilch, Generalized force model of traffic dynamics. *Phys. Rev. E* **58**, 133 (1998)
3. J.-P. Bouchaud, Power laws in economics and finance: some ideas from physics. *Quant. Finance* **1**, 105 (2001)
4. R.N. Mantegna, H.E. Stanley, *Introduction to Econophysics: Correlations and Complexity in Finance* (Cambridge University Press, Cambridge, 2007)
5. J.D. Murray, How the leopard gets its spots. *Sci. Am.* **258**, 80 (1988)
6. A.M. Turing, The chemical basis of morphogenesis. *Philos. Trans. R. Soc. B* **237**, 37 (1952)
7. M.C. Cross, H.S. Greenside, *Pattern Formation and Dynamics in Non-Equilibrium Systems* (Cambridge University Press, Cambridge, 2009)
8. J.D. Murray, A pre-pattern formation mechanism for animal coat markings. *J. Theor. Biol.* **88**, 161 (1981)
9. J.C. Maxwell, *Theory of Heat* (Longmans, London, 1871)
10. J.C. Maxwell, V. Illustrations of the dynamical theory of gases. Part I. On the motions and collisions of perfectly elastic spheres. *Philos. Mag.* **19**, 19 (1860)
11. L. Boltzmann, Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung respektive den Sätzen über das Wärmegleichgewicht. *Wien. Ber.* **76**, 373 (1877)
12. N.G. van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, Amsterdam, 2007)
13. D. Alonso, A.J. McKane, M. Pascual, Stochastic amplification in epidemics. *J. R. Soc. Interface* **4**, 575 (2007)
14. M.A. Nowak, *Evolutionary Dynamics* (Harvard University Press, Cambridge, 2006)
15. A.J. McKane, T.J. Newman, Predator-prey cycles from resonant amplification of demographic stochasticity. *Phys. Rev. Lett.* **94**, 218102 (2005)
16. R.A. Fisher, On the dominance ratio. *Proc. R. Soc. Edinb.* **42**, 321 (1922)
17. J.B.S. Haldane, A mathematical theory of natural and artificial selection. V. Selection and mutation. *Proc. Cambridge Philos. Soc.* **23**, 838 (1927)

18. S. Wright, Evolution in Mendelian populations. *Genetics* **16**, 97 (1931)
19. M. Kimura, On the probability of fixation of mutant genes in a population. *Genetics* **47**, 713 (1962)
20. T.L. Vincent, J.S. Brown, *Evolutionary Game Theory, Natural Selection, and Darwinian Dynamics* (Cambridge University Press, Cambridge, 2005)
21. Cancer Research UK. All cancers combined: Key facts. *Cancer Research UK*, (2014)
22. S. Mukherjee, *The Emperor of all Maladies: A Biography of Cancer* (Fourth Estate, London, 2011)
23. C. Nordling, A new theory on the cancer-inducing mechanism. *Br. J. Cancer* **7**, 68 (1953)
24. P. Armitage, R. Doll, The age distribution of cancer and a multi-stage theory of carcinogenesis. *Br. J. Cancer* **8**, 1 (1954)
25. J. Fisher, Multiple-mutation theory of carcinogenesis. *Nature* **181**, 651 (1958)
26. A.G. Knudson, Mutation and cancer: statistical study of retinoblastoma. *Proc. Natl. Acad. Sci. U.S.A.* **68**, 820 (1971)
27. S.H. Moolgavkar, The multistage theory of carcinogenesis and the age distribution of cancer in man. *J. Natl. Cancer Inst.* **61**, 49 (1978)
28. S.H. Moolgavkar, A.G. Knudson, Mutation and cancer: a model for human carcinogenesis. *J. Natl. Cancer Inst.* **66**, 1037 (1981)
29. I. Bozic, T. Antal, H. Ohtsuki, H. Carter, D. Kim, S. Chen, R. Karchin, K.W. Kinzler, B. Vogelstein, M.A. Nowak, Accumulation of driver and passenger mutations during tumor progression. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 18545 (2010)
30. I. Bozic, J.G. Reiter, B. Allen, T. Antal, K. Chatterjee, P. Shah, Y.S. Moon, A. Yaqubie, N. Kelly, D.T. Le, E.J. Lipson, P.B. Chapman, L.A. Diaz, Jr, B. Vogelstein, M.A. Nowak, Evolutionary dynamics of cancer in response to targeted combination therapy. *eLife* **2**, e00747 (2013)
31. J. Denes, D. Krewski, An exact representation for the generating function for the Moolgavkar-Venzon-Knudson two-stage model of carcinogenesis with stochastic stem cell growth. *Math. Biosci.* **131**, 185 (1996)
32. T. Antal, P. Krapivsky, Exact solution of a two-type branching process: models of tumor progression. *J. Stat. Mech.* **2011**, P08018 (2011)
33. M.A. Nowak, F. Michor, Y. Iwasa, The linear process of somatic evolution. *Proc. Natl. Acad. Sci. U.S.A.* **100**, 14966 (2003)
34. N.L. Komarova, A. Sengupta, M.A. Nowak, Mutation-selection networks of cancer initiation: tumor suppressor genes and chromosomal instability. *J. Theor. Biol.* **223**, 433 (2003)
35. A.M. Colman, *Game Theory and its Applications in the Social and Biological Sciences* (Butterworth-Heinemann, Oxford, 1995)
36. J.W. Weibull, *Evolutionary Game Theory* (MIT Press, Cambridge, 1995)
37. J. Hofbauer, K. Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge University Press, Cambridge, 1998)
38. H. Gintis, *Game Theory Evolving* (Princeton University Press, Princeton, 2009)
39. W.H. Sandholm, *Population Games and Evolutionary Dynamics* (MIT Press, Cambridge, 2010)
40. R.C. Maclean, I. Gudelj, Resource competition and social conflict in experimental populations of yeast. *Nature* **441**, 498 (2006)
41. J. Gore, H. Youk, A. van Oudenaarden, Snowdrift game dynamics and facultative cheating in yeast. *Nature* **459**, 253 (2009)
42. R.C. MacLean, A. Fuentes-Hernandez, D. Greig, L.D. Hurst, I. Gudelj, A mixture of “cheats” and “co-operators” can enable maximal group benefit. *PLoS Biol.* **8**, e1000486 (2010)
43. X.-Y. Li, C. Pietschke, S. Fraune, P.M. Altrock, T.C.G. Bosch, A. Traulsen, Which games are growing bacterial populations playing? *J. R. Soc. Interface* **12**, 20150121 (2015)
44. P. Ashcroft, P.M. Altrock, T. Galla, Fixation in finite populations evolving in fluctuating environments. *J. R. Soc. Interface* **11**, 20140663 (2014)
45. P. Ashcroft, A. Traulsen, T. Galla, When the mean is not enough: calculating fixation time distributions in birth-death processes. *Phys. Rev. E* **92**, 042154 (2015)
46. Y. Iwasa, F. Michor, M.A. Nowak, Stochastic tunnels in evolutionary dynamics. *Genetics* **166**, 1571 (2004)

47. M.A. Nowak, F. Michor, N.L. Komarova, Y. Iwasa, Evolutionary dynamics of tumor suppressor gene inactivation. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 10635 (2004)
48. Y. Iwasa, F. Michor, N.L. Komarova, M.A. Nowak, Population genetics of tumor suppressor genes. *J. Theor. Biol.* **233**, 15 (2005)
49. F. Michor, Y. Iwasa, Dynamics of metastasis suppressor gene inactivation. *J. Theor. Biol.* **241**, 676 (2006)
50. S.R. Proulx, The rate of multi-step evolution in Moran and Wright-Fisher populations. *Theor. Popul. Biol.* **80**, 197 (2011)
51. H. Haeno, Y.E. Maruvka, Y. Iwasa, F. Michor, Stochastic tunneling of two mutations in a population of cancer cells. *PLoS ONE* **8**, e65724 (2013)
52. M. Assaf, B. Meerson, Extinction of metastable stochastic populations. *Phys. Rev. E* **81**, 021116 (2010)
53. I. Lohmar, B. Meerson, Switching between phenotypes and population extinction. *Phys. Rev. E* **84**, 051901 (2011)
54. O. Gottesman, B. Meerson, Multiple extinction routes in stochastic population models. *Phys. Rev. E* **85**, 021140 (2012)
55. O.A. van Herwaarden, J. Grasman, Stochastic epidemics: major outbreaks and the duration of the endemic period. *J. Math. Biol.* **33**, 581 (1995)
56. A. Kamenev, B. Meerson, Extinction of an infectious disease: a large fluctuation in a nonequilibrium system. *Phys. Rev. E* **77**, 061107 (2008)
57. M.I. Dykman, I.B. Schwartz, A.S. Landsman, Disease extinction in the presence of random vaccination. *Phys. Rev. Lett.* **101**, 078101 (2008)
58. A.J. Black, A.J. McKane, WKB calculation of an epidemic outbreak distribution. *J. Stat. Mech.* **2011**, P12006 (2011)
59. L. Billings, L. Mier-Y-Teran-Romero, B. Lindley, I.B. Schwartz, Intervention-based stochastic disease eradication. *PLoS ONE* **8**, e70211 (2013)
60. A.J. Black, A. Traulsen, T. Galla, Mixing times in evolutionary games. *Phys. Rev. Lett.* **109**, 028101 (2012)
61. C.H. Waddington, *The Strategy of the Genes* (Allen and Unwin, London, 1957)

Chapter 2

Technical Background

In this chapter I will outline the mathematical and numerical techniques used throughout this thesis to analyse stochastic systems.

2.1 Definition: Stochastic Processes

A stochastic process is defined as a function of stochastic or ‘random’ variables and time [1]. In this thesis the function describes how the random variables jump between values as time progresses. We will focus on how the probability distributions of the stochastic variables evolve in time, and the quantities that can be obtained from this. As an example consider the stochastic variables $\mathbf{n}(t) = [n_1(t), n_2(t), \dots]$, where $n_i(t)$ is an integer that describes the number of particles of type i at time t . We can picture the classic probability exemplar of coloured balls in a bag, with time evolution described by balls being removed and/or added to the bag at random. The vector $\mathbf{n}(t)$ describes the state of the system. If the initial state of the system is known at time t_0 to be $\mathbf{n}_0$, then the state of the system at a future time will be described by a distribution for the probability that the state $\mathbf{n}$ will be observed at time t . We write this as $P(\mathbf{n}, t | \mathbf{n}_0, t_0)$, where the vertical line is to be read as “given that”. This is known as a conditional probability.

This thesis will focus on a subclass of stochastic processes in which the future behaviour of a system is determined only by the state of the system at the present time, and has no dependence on past states. These processes are known as Markov processes, after the Russian mathematician Andrey Markov (1856–1922). If the system was observed in the state $\mathbf{n}_j$ at time t_j for $j = 0, 1, \dots, k-1$ ($t_0 < t_1 < \dots < t_{k-1}$), then the Markov property can be written as

$$P(\mathbf{n}_k, t_k | \mathbf{n}_0, t_0; \mathbf{n}_1, t_1; \dots; \mathbf{n}_{k-1}, t_{k-1}) = P(\mathbf{n}_k, t_k | \mathbf{n}_{k-1}, t_{k-1}), \quad (2.1)$$

where $t_{k-1} < t_k$. The semicolons in the above expression are to be read as “and”.

2.2 Chapman–Kolmogorov and Master Equations

The probability of observing a specific ‘path’ through state-space is described by the joint probability distribution

$$P(\mathbf{n}_0, t_0; \mathbf{n}_1, t_1; \dots; \mathbf{n}_k, t_k). \quad (2.2)$$

This joint distribution is related to the conditional distribution through Bayes’¹ rule [1], which states

$$P(\mathbf{n}_{j+1}, t_{j+1}; \dots; \mathbf{n}_k, t_k | \mathbf{n}_0, t_0; \dots; \mathbf{n}_j, t_j) = \frac{P(\mathbf{n}_0, t_0; \dots; \mathbf{n}_k, t_k)}{P(\mathbf{n}_0, t_0; \dots; \mathbf{n}_j, t_j)}. \quad (2.3)$$

Together with the Markov assumption (2.1), we can write

$$\begin{aligned} P(\mathbf{n}_0, t_0; \dots; \mathbf{n}_k, t_k) &= P(\mathbf{n}_k, t_k | \mathbf{n}_{k-1}, t_{k-1}) P(\mathbf{n}_0, t_0; \dots; \mathbf{n}_{k-1}, t_{k-1}) \\ &= P(\mathbf{n}_k, t_k | \mathbf{n}_{k-1}, t_{k-1}) P(\mathbf{n}_{k-1}, t_{k-1} | \mathbf{n}_{k-2}, t_{k-2}) P(\mathbf{n}_0, t_0; \dots; \mathbf{n}_{k-2}, t_{k-2}) \\ &\vdots \\ &= \left[\prod_{j=1}^k P(\mathbf{n}_j, t_j | \mathbf{n}_{j-1}, t_{j-1}) \right] P(\mathbf{n}_0, t_0). \end{aligned} \quad (2.4)$$

From this equation we can recover two fundamental relations. First, taking $k = 1$ in Eq. (2.4) and then summing over $\mathbf{n}_0$ (known as marginalising the joint distribution), we recover

$$P(\mathbf{n}_1, t_1) = \sum_{\mathbf{n}_0} P(\mathbf{n}_1, t_1 | \mathbf{n}_0, t_0) P(\mathbf{n}_0, t_0). \quad (2.5)$$

Second, taking $k = 2$ in Eq. (2.4) and summing over $\mathbf{n}_1$ gives

$$P(\mathbf{n}_0, t_0; \mathbf{n}_2, t_2) = \sum_{\mathbf{n}_1} P(\mathbf{n}_2, t_2 | \mathbf{n}_1, t_1) P(\mathbf{n}_1, t_1 | \mathbf{n}_0, t_0) P(\mathbf{n}_0, t_0). \quad (2.6)$$

Dividing by $P(\mathbf{n}_0, t_0)$ and applying Bayes’ rule gives

$$P(\mathbf{n}_2, t_2 | \mathbf{n}_0, t_0) = \sum_{\mathbf{n}_1} P(\mathbf{n}_2, t_2 | \mathbf{n}_1, t_1) P(\mathbf{n}_1, t_1 | \mathbf{n}_0, t_0). \quad (2.7)$$

¹Rev. Thomas Bayes (1701–1761).

This is the Chapman–Kolmogorov equation, named after Sydney Chapman (1888–1970) and Andrey Kolmogorov (1903–1987). Analogous expressions for Eqs. (2.5) and (2.7) can be obtained for continuous state-spaces by replacing sums with the corresponding integrals.

To simplify the Chapman–Kolmogorov equation we will look at the evolution of the distribution in a small time-step, Δt , conditioned on an initial configuration. We write the Chapman–Kolmogorov equation as

$$P(\mathbf{n}, t + \Delta t | \mathbf{n}_0, t_0) = \sum_{\mathbf{n}'} P(\mathbf{n}, t + \Delta t | \mathbf{n}', t) P(\mathbf{n}', t | \mathbf{n}_0, t_0). \quad (2.8)$$

To evaluate the term $P(\mathbf{n}, t + \Delta t | \mathbf{n}', t)$, i.e. what is the probability for the system to be found in state $\mathbf{n}$ a short period of time after it was observed in state $\mathbf{n}'$, we introduce the transition rate $w_{\mathbf{n}, \mathbf{n}'}$. This is the rate per unit time at which the transition from $\mathbf{n}'$ to $\mathbf{n}$ occurs. All processes that will be considered in this chapter will have time-independent reaction rates. Such processes are referred to as ‘homogeneous’ Markov processes [1]. We can expand the jump probability as

$$P(\mathbf{n}, t + \Delta t | \mathbf{n}', t) = \delta_{\mathbf{n}, \mathbf{n}'} + w_{\mathbf{n}, \mathbf{n}'} \Delta t + \mathcal{O}(\Delta t^2). \quad (2.9)$$

As the system must be found somewhere at time $t + \Delta t$, we have the normalisation condition $\sum_{\mathbf{n}} P(\mathbf{n}, t + \Delta t | \mathbf{n}', t) = 1$ for all $\mathbf{n}'$. From this we find

$$P(\mathbf{n}, t + \Delta t | \mathbf{n}', t) = (1 - \delta_{\mathbf{n}, \mathbf{n}'}) w_{\mathbf{n}, \mathbf{n}'} \Delta t + \delta_{\mathbf{n}, \mathbf{n}'} (1 - w_{\mathbf{n}, \mathbf{n}} \Delta t) + \mathcal{O}(\Delta t^2), \quad (2.10)$$

where $w_{\mathbf{n}, \mathbf{n}} = \sum_{\mathbf{n}' \neq \mathbf{n}} w_{\mathbf{n}', \mathbf{n}}$. Substituting this relation into Eq. (2.8) gives

$$\begin{aligned} P(\mathbf{n}, t + \Delta t | \mathbf{n}_0, t_0) &= \sum_{\mathbf{n}' \neq \mathbf{n}} w_{\mathbf{n}, \mathbf{n}'} \Delta t P(\mathbf{n}', t | \mathbf{n}_0, t_0) \\ &\quad + \left[1 - \sum_{\mathbf{n}' \neq \mathbf{n}} w_{\mathbf{n}', \mathbf{n}} \Delta t \right] P(\mathbf{n}, t | \mathbf{n}_0, t_0) + \mathcal{O}(\Delta t^2). \end{aligned} \quad (2.11)$$

Taking the limit $\Delta t \rightarrow 0$ and rearranging gives the continuous-time master equation

$$\dot{P}(\mathbf{n}, t | \mathbf{n}_0, t_0) = \sum_{\mathbf{n}' \neq \mathbf{n}} [w_{\mathbf{n}, \mathbf{n}'} P(\mathbf{n}', t | \mathbf{n}_0, t_0) - w_{\mathbf{n}', \mathbf{n}} P(\mathbf{n}, t | \mathbf{n}_0, t_0)]. \quad (2.12)$$

This has the intuitive interpretation that the probability of being in state $\mathbf{n}$ increases due to transitions into the state (with rate $w_{\mathbf{n}, \mathbf{n}'}$) and decreases due to transitions out of the state (with rate $w_{\mathbf{n}', \mathbf{n}}$).

Sometimes it is more intuitive to label the transition rates not by their initial and final state, but by the initial state and the stoichiometric effect of that transition. We write $T_{\mathbf{n}}^{\mathbf{v}} = w_{\mathbf{n}+\mathbf{v}, \mathbf{n}}$ for the transition rate from state $\mathbf{n}$ to state $\mathbf{n} + \mathbf{v}$, where $\mathbf{v}$ is

the so-called stoichiometric coefficient [1]. The master equation (2.12) can then be rewritten as

$$\dot{P}_{\mathbf{n}}(t) = \sum_{\nu} [T_{\mathbf{n}-\nu}^{\nu} P_{\mathbf{n}-\nu}(t) - T_{\mathbf{n}}^{\nu} P_{\mathbf{n}}(t)], \quad (2.13)$$

where we have suppressed the initial condition notation for optical convenience. Throughout it will be assumed that all probabilities are conditioned on being in state $\mathbf{n}_0$ at time $t_0 = 0$, such that $P_{\mathbf{n}}(t) = P(\mathbf{n}, t | \mathbf{n}_0, 0)$. The sum in Eq. (2.13) runs over all possible reactions.

2.3 Birth–Death Process

An interesting case which will feature prominently in this thesis is the one-dimensional one-step process, often referred to as a birth–death process [2]. This process has been used to describe the proliferation of a disease, such as influenza, or the extinction of a colonising species [2]. We will use birth–death processes to describe the evolution of a population of two competing types of individual, labelled type *A* and type *B*. We assume that the total number of individuals, N , is constant, such that the state of the population can be described by a single number. We let i be the number of individuals of type *A* ($0 \leq i \leq N$), $N - i$ is then the number of individuals of type *B*. Only two possible reactions are allowed:

- (i) Birth events where the number of type-*A* individuals increases by one and the number of type-*B* individuals decreases by one. We will often use the notation $i \rightarrow i + 1$;
- (ii) Death events where the number of type-*A* individuals decreases by one and the number of type-*B* individuals increases by one, labelled as $i \rightarrow i - 1$.

This process is illustrated in Fig. 2.1. We can replace the transition rates T_i^{ν} in the master equation (2.13) with the birth rate b_i if $\nu = +1$ and the death rate d_i if $\nu = -1$. These transition rates are extensive; they scale linearly with the system size




Fig. 2.1 The birth–death process described by Eq. (2.14). Given the system is in state i , it can only jump to states $i + 1$ or to state $i - 1$ with rates b_i and d_i respectively. The rates b_0 and d_N determine the type of boundary. If $b_0 = d_N = 0$ then the boundaries are absorbing, otherwise they are reflecting. A quantity of interest in the absorbing boundary case is the fixation probability, ϕ_i . This is the probability that the system reaches the absorbing boundary at state N , given that it is initially in state i

N such that larger systems have more frequent reactions [1]. The master equation for the birth–death process is

$$\dot{P}_i(t) = b_{i-1}P_{i-1}(t) + d_{i+1}P_{i+1}(t) - (b_i + d_i)P_i(t). \quad (2.14)$$

The birth and death rates must satisfy the boundary conditions $b_N = d_0 = 0$, such that the number of individuals of a single type cannot leave the domain $0 \leq i \leq N$. The states $i = 0$ and $i = N$ may have a special form, which leads to further boundary conditions. This is discussed below.

Equation (2.14) can also be written in the more compact matrix representation as

$$\dot{\mathbf{P}}(t) = \mathbb{W} \cdot \mathbf{P}(t), \quad (2.15)$$

where the $(N + 1) \times (N + 1)$ tridiagonal matrix $\mathbb{W}$ has elements $w_{i+1,i} = b_i$, $w_{i-1,i} = d_i$, and $w_{i,i} = -(b_i + d_i)$. All other elements are zero. Direct integration of this equation gives the formal solution

$$\mathbf{P}(t) = e^{\mathbb{W}t} \cdot \mathbf{P}(0). \quad (2.16)$$

Despite its compact form, this solution is not very helpful as matrix exponentials are notoriously difficult to evaluate and the computational time to evaluate this function can scale exponentially with the size of the state-space. In Chap. 4 we will use (a truncated form of) Eq. (2.16) to compute some statistics of birth–death processes. It is worth noting that Eq. (2.15) and solution (2.16) are not just limited to birth–death processes, but can also be written down for the more general master equation (2.13) with appropriate choices of the matrix $\mathbb{W}$ and indexing of the state-space.

The birth–death process can also be described in discrete time (as can other processes). Substituting the birth and death rates into Eq. (2.11) gives the discrete-time master equation

$$P_i(t + \Delta t) = \Delta t b_{i-1} P_{i-1}(t) + \Delta t d_{i+1} P_{i+1}(t) + (1 - \Delta t b_i - \Delta t d_i) P_i(t). \quad (2.17)$$

A suitable choice for the time-step is $\Delta t = 1/N$ as one unit of time is then a generation with N possible birth and death events. Thus the probability of a birth or death event happening in a single time-step are b_i/N and d_i/N . These quantities are intensive. Any global rescaling of the reaction probabilities corresponds to a global rescaling of time. This allows us to absorb the time-step $\Delta t = 1/N$ into the transition probabilities b_i and d_i , provided that $b_i \geq 0$, $d_i \geq 0$, and $b_i + d_i \leq 1$ for all $0 \leq i \leq N$, and set $\Delta t = 1$ throughout such that the variable t simply counts the number of time-steps. This is the convention followed in numerous sources [3–5]. The number of generations, which is the unit of time used in the continuous-time framework, is then given by t/N .

Backward Equations

We can introduce a companion to the master equation which is given by the adjoint of Eq. (2.14) [1, 2]. For the birth–death process this ‘backward’ equation reads

$$\dot{Q}_{j;i}(t) = b_i Q_{j;i+1}(t) + d_i Q_{j;i-1}(t) - (b_i + d_i) Q_{j;i}(t), \quad (2.18)$$

where $Q_{j;i}(t)$ is the probability to be found in state j a period of time t after being found in state i . The forward master equation (2.14) describes the distribution at a future time given a fixed initial condition. The backward master equation (2.18), however, has a fixed final condition and variable initial state. This is analogous to the Schrödinger and Heisenberg pictures of quantum mechanics, where the time-dependence is moved between the state vectors and the operators [1, 6]. The use of this form of equation will become obvious in the next section. In matrix form the backward master equation can be written as

$$\dot{\mathbf{Q}}_j(t) = \mathbf{Q}_j(t) \cdot \mathbb{W}, \quad (2.19)$$

where $\mathbb{W}$ is the same matrix as appears in Eq. (2.15) and $\mathbf{Q}_j(t)$ is now a row vector.

We can also write the backward master equation (2.18) in discrete-time form as

$$Q_{j;i}(t + \Delta t) = \Delta t b_i Q_{j;i+1}(t) + \Delta t d_i Q_{j;i-1}(t) + (1 - \Delta t b_i - \Delta t d_i) Q_{j;i}(t), \quad (2.20)$$

and we will use this form in the next section to derive arrival-time statistics of birth–death processes with absorbing boundaries. As before we will set $\Delta t = 1$ throughout.

2.4 Fixation Probability and Mean Fixation Times

We continue with the example of birth–death processes. If the birth and death rates at the boundaries satisfy $b_0 = d_N = 0$, then the states 0 and N are absorbing; no further dynamics can occur once the population has reached one of these states. The state $i = 0$ corresponds to the extinction of type- A individuals from the population and the fixation of type- B individuals. Once extinct, no birth events can take place that reintroduce type- A individuals (except in the case where mutations can occur during reproduction events as discussed in the next section). In the state $i = N$, the population consists entirely of type- A individuals. Here type- A individuals have reached fixation and type- B individuals have become extinct from the population.

Explicit expressions for the fixation probabilities and mean fixation times for the birth–death process can be obtained from Eq. (2.20) [or equivalently Eq. (2.18) in continuous time]. The calculations can be found in numerous sources [1–5, 7, 8], but we will repeat them here for completeness. We will consider the case of a discrete-time process, however all results in this section apply to continuous-time processes

too. As described above we choose the time-step to be $\Delta t = 1$ and rescale the transition probabilities b_i and d_i accordingly.

Fixation Probability

To compute the probability that type- A individuals reach fixation, we set $j = N$ in Eq. (2.20) and take the limit $t \rightarrow \infty$. Writing $\phi_i(t) = \lim_{t \rightarrow \infty} Q_{N;i}(t)$, we obtain

$$\phi_i = b_i \phi_{i+1} + d_i \phi_{i-1} + (1 - b_i - d_i) \phi_i. \quad (2.21)$$

This has the intuitive interpretation that the probability of type- A individuals reaching fixation from state i is given by the probability of hopping to $i+1$ and reaching fixation from there, plus the probability of hopping to $i-1$ and reaching fixation from there, plus the probability of not hopping and reaching fixation from i . Equation (2.21) is subject to the boundary conditions $\phi_0 = 0$ and $\phi_N = 1$. It can be solved explicitly by introducing the difference variable $v_i = \phi_i - \phi_{i-1}$ to give $b_i v_{i+1} = d_i v_i$. The solution can be found recursively to give

$$v_i = \frac{d_{i-1}}{b_{i-1}} v_{i-1} = \frac{d_{i-1}}{b_{i-1}} \frac{d_{i-2}}{b_{i-2}} v_{i-2} = \left(\prod_{j=1}^{i-1} \frac{d_j}{b_j} \right) v_1 = \left(\prod_{j=1}^{i-1} \frac{d_j}{b_j} \right) \phi_1, \quad (2.22)$$

where we have used the boundary conditions to write $v_1 = \phi_1 - \phi_0 = \phi_1$. Taking the sum over the v_i yields $\sum_{k=1}^N v_k = \phi_N = 1$, and hence we can write

$$1 = \sum_{k=1}^N v_k = \phi_1 \sum_{k=1}^N \prod_{j=1}^{k-1} \frac{d_j}{b_j} = \phi_1 \left(1 + \sum_{k=1}^{N-1} \prod_{j=1}^k \frac{d_j}{b_j} \right). \quad (2.23)$$

Introducing the compact notation $\gamma_j = d_j/b_j$, we arrive at

$$\phi_1 = \frac{1}{1 + \sum_{k=1}^{N-1} \prod_{j=1}^k \gamma_j}, \quad (2.24a)$$

$$\phi_i = \sum_{k=1}^i v_k = \phi_1 \sum_{k=1}^i \prod_{j=1}^{k-1} \gamma_j = \frac{1 + \sum_{k=1}^{i-1} \prod_{j=1}^k \gamma_j}{1 + \sum_{k=1}^{N-1} \prod_{j=1}^k \gamma_j}. \quad (2.24b)$$

Mean Unconditional Fixation Time

A similar procedure is used to calculate the mean fixation times. We first consider the case of the unconditional fixation time, that is the number of time-steps taken until either type A or type B have reached fixation, or simply the number of time-steps taken to reach either of the absorbing boundaries. This quantity is itself a stochastic variable which follows a distribution. This distribution can be characterised by the infinite set of moments, the most informative of which is the mean. We introduce the variable $\vartheta_i(t) = Q_{0;i}(t) + Q_{N;i}(t)$, which is the cumulative probability to have

arrived at either absorbing state t time-steps after being in state i . The probability to arrive at an absorbing boundary at time t is then $\vartheta_i(t) - \vartheta_i(t-1)$. The mean unconditional fixation time, given the initial condition i , is then defined by

$$t_i = \frac{\sum_{t=0}^{\infty} t [\vartheta_i(t) - \vartheta_i(t-1)]}{\sum_{t=0}^{\infty} [\vartheta_i(t) - \vartheta_i(t-1)]} = \sum_{t=0}^{\infty} t [\vartheta_i(t) - \vartheta_i(t-1)], \quad (2.25)$$

where we have used $\vartheta_i(t \leq 0) = 0$ and $\lim_{t \rightarrow \infty} \vartheta_i(t) = 1$ to find $\sum_{t=0}^{\infty} [\vartheta_i(t) - \vartheta_i(t-1)] = 1$. From the backward master equation (2.20) it can be seen that $\vartheta_i(t)$ satisfies

$$\vartheta_i(t+1) = b_i \vartheta_{i+1}(t) + d_i \vartheta_{i-1}(t) + (1 - b_i - d_i) \vartheta_i(t). \quad (2.26)$$

Subtracting $\vartheta_i(t)$ from both sides, multiplying by t and summing yields

$$\begin{aligned} \sum_{t=0}^{\infty} t [\vartheta_i(t+1) - \vartheta_i(t)] &= b_i \sum_{t=0}^{\infty} t [\vartheta_{i+1}(t) - \vartheta_{i+1}(t-1)] \\ &\quad + d_i \sum_{t=0}^{\infty} t [\vartheta_{i-1}(t) - \vartheta_{i-1}(t-1)] \\ &\quad + (1 - b_i - d_i) \sum_{t=0}^{\infty} t [\vartheta_i(t) - \vartheta_i(t-1)]. \end{aligned} \quad (2.27)$$

The left-hand side of this equation needs some work to extract the mean fixation time t_i . We can write this as

$$\begin{aligned} \sum_{t=0}^{\infty} t [\vartheta_i(t+1) - \vartheta_i(t)] &= \sum_{t=1}^{\infty} (t-1) [\vartheta_i(t) - \vartheta_i(t-1)] \\ &= \sum_{t=1}^{\infty} t [\vartheta_i(t) - \vartheta_i(t-1)] - \sum_{t=1}^{\infty} [\vartheta_i(t) - \vartheta_i(t-1)] \\ &= \sum_{t=1}^{\infty} t [\vartheta_i(t) - \vartheta_i(t-1)] - \sum_{t=0}^{\infty} [\vartheta_i(t) - \vartheta_i(t-1)] \\ &= t_i - 1, \end{aligned} \quad (2.28)$$

where we have used $\vartheta_i(0) = 0$ and in the last step we have used $\lim_{t \rightarrow \infty} \vartheta_i(t) = 1$. Equation (2.27) can now be written as

$$t_i = b_i t_{i+1} + d_i t_{i-1} + (1 - b_i - d_i) t_i + 1. \quad (2.29)$$

This also has the intuitive interpretation that the time to reach fixation from state i is given by the probability of hopping to $i + 1$ and reaching fixation from there, plus the probability of hopping to $i - 1$ and reaching fixation from there, plus the probability of not hopping and reaching fixation from i , plus the time taken for this step. This equation has boundary conditions $t_0 = t_N = 0$. Using Eq. (2.29), the difference variable $v_i = t_i - t_{i-1}$ satisfies $v_i = \gamma_{i-1} v_{i-1} - 1/b_{i-1}$. The solution can be found recursively to give

$$v_i = t_1 \prod_{m=1}^{i-1} \gamma_m - \sum_{\ell=1}^{i-1} \frac{1}{b_\ell} \prod_{m=\ell+1}^{i-1} \gamma_m, \quad (2.30)$$

where we have used the boundary conditions to write $v_1 = t_1 - t_0 = t_1$. Taking the sum over the set $\{v_i\}$ yields $\sum_{k=i+1}^N v_k = -t_i$. In particular we have $\sum_{k=2}^N v_k = -t_1$. With this we can write

$$t_1 = -t_1 \sum_{k=2}^N \prod_{m=1}^{k-1} \gamma_m + \sum_{k=2}^N \sum_{\ell=1}^{k-1} \frac{1}{b_\ell} \prod_{m=\ell+1}^{k-1} \gamma_m, \quad (2.31)$$

and from this we can calculate

$$t_1 = \phi_1 \sum_{k=1}^{N-1} \sum_{\ell=1}^k \frac{1}{b_\ell} \prod_{m=\ell+1}^k \gamma_m, \quad (2.32a)$$

$$t_i = - \sum_{k=i+1}^N v_k = -t_1 \sum_{k=i}^{N-1} \prod_{m=1}^k \gamma_m + \sum_{k=i}^{N-1} \sum_{\ell=1}^k \frac{1}{b_\ell} \prod_{m=\ell+1}^k \gamma_m. \quad (2.32b)$$

These are expressions for the average number of time-steps taken to reach either absorbing boundary. The mean unconditional fixation time measured in units of generations is found by dividing Eq. (2.32) by the system size, N .

If we instead considered a continuous-time process with extensive birth and death rates b_i and d_i , Eq. (2.32) describe the mean unconditional fixation time in units of generations. The fraction $1/b_\ell$ in the Eq. (2.32) ensures that time is measured in units of $1/N$, i.e. in generations.

Mean Conditional Fixation Time

The conditional fixation time is determined in a very similar way. We here focus on the fixation of type A . We introduce the variable $\varphi_i(t) = Q_{N;i}(t)$, which is the cumulative probability to have arrived at the all-type- A state t time-steps after being in state i . The probability to arrive in state N at time t is $\varphi_i(t) - \varphi_i(t-1)$, and hence the mean conditional fixation time, given the initial condition i , is then defined by

$$t_{i|A} = \frac{\sum_{t=0}^{\infty} t [\varphi_i(t) - \varphi_i(t-1)]}{\sum_{t=0}^{\infty} [\varphi_i(t) - \varphi_i(t-1)]} = \frac{1}{\phi_i} \sum_{t=0}^{\infty} t [\varphi_i(t) - \varphi_i(t-1)], \quad (2.33)$$

where ϕ_i is the fixation probability defined in Eq. (2.24) and $\varphi_i(t \leq 0) = 0$.

From Eq. (2.20) it can be seen that $\varphi_i(t)$ satisfies

$$\varphi_i(t+1) = b_i \varphi_{i+1}(t) + d_i \varphi_{i-1}(t) + (1 - b_i - d_i) \varphi_i(t). \quad (2.34)$$

Subtracting $\varphi_i(t)$ from both sides, multiplying by t and summing yields

$$\begin{aligned} \sum_{t=0}^{\infty} t [\varphi_i(t+1) - \varphi_i(t)] &= b_i \sum_{t=0}^{\infty} t [\varphi_{i+1}(t) - \varphi_{i+1}(t-1)] \\ &\quad + d_i \sum_{t=0}^{\infty} t [\varphi_{i-1}(t) - \varphi_{i-1}(t-1)] \\ &\quad + (1 - b_i - d_i) \sum_{t=0}^{\infty} t [\varphi_i(t) - \varphi_i(t-1)]. \end{aligned} \quad (2.35)$$

Following the procedure above, and introducing the variable $\theta_i = \phi_i t_{i|A}$, we can arrive at the expression

$$\theta_i = b_i \theta_{i+1} + d_i \theta_{i-1} + (1 - b_i - d_i) \theta_i + \phi_i. \quad (2.36)$$

We note that Eqs. (2.29) and (2.36) are very similar, but the difference is more than just a global pre-factor ϕ_i . This equation has boundary conditions $\theta_0 = \theta_N = 0$. Using Eq. (2.36), the difference variable $\eta_i = \theta_i - \theta_{i-1}$ satisfies $\eta_i = \gamma_{i-1} \eta_{i-1} - \phi_{i-1}/b_{i-1}$. This can again be solved recursively to give

$$\eta_i = \theta_1 \prod_{m=1}^{i-1} \gamma_m - \sum_{\ell=1}^{i-1} \frac{\phi_\ell}{b_\ell} \prod_{m=\ell+1}^{i-1} \gamma_m, \quad (2.37)$$

where we have used the boundary conditions to write $\eta_1 = \theta_1 - \theta_0 = \theta_1$. Taking the sum over the set $\{\eta_i\}$ yields $\sum_{k=i+1}^N \eta_k = -\theta_i$. In particular we have $\sum_{k=2}^N \eta_k = -\theta_1$. With this we can write

$$\theta_1 = -\theta_1 \sum_{k=2}^N \prod_{m=1}^{k-1} \gamma_m + \sum_{k=2}^N \sum_{\ell=1}^{k-1} \frac{\phi_\ell}{b_\ell} \prod_{m=\ell+1}^{k-1} \gamma_m, \quad (2.38)$$

and from this we can calculate

$$\theta_1 = \phi_1 \sum_{k=1}^{N-1} \sum_{\ell=1}^k \frac{\phi_\ell}{b_\ell} \prod_{m=\ell+1}^k \gamma_m, \quad (2.39a)$$

$$\theta_i = - \sum_{k=i+1}^N \eta_k = -\theta_1 \sum_{k=i}^{N-1} \prod_{m=1}^k \gamma_m + \sum_{k=i}^{N-1} \sum_{\ell=1}^k \frac{\phi_\ell}{b_\ell} \prod_{m=\ell+1}^k \gamma_m. \quad (2.39b)$$

The average number of time-steps until type- A individuals take over the population is then given by $t_{i|A} = \theta_i / \phi_i$. Again these expressions hold for continuous-time processes with extensive birth and death rates, where time is intrinsically measured in units of generations.

Higher Moments

Although the mean fixation time can provide useful information, higher moments are often sought after to provide a more complete picture. For example, the variance of the random variable X ,

$$\text{var}(X) = \langle X^2 \rangle - \langle X \rangle^2, \quad (2.40)$$

which describes the spread of the values of X requires the computation of the second moment, $\langle X^2 \rangle$. We use the notation $\langle \cdot \rangle$ to denote the average value. In general the r -th moment of the unconditional fixation time is defined by

$$\tau_i^{(r)} = \frac{\sum_{t=0}^{\infty} t^r [\vartheta_i(t) - \vartheta_i(t-1)]}{\sum_{t=0}^{\infty} [\vartheta_i(t) - \vartheta_i(t-1)]} = \sum_{t=0}^{\infty} t^r [\vartheta_i(t) - \vartheta_i(t-1)], \quad (2.41)$$

where again $\vartheta_i(t) - \vartheta_i(t-1)$ is the probability to reach either of the absorbing boundaries at time t . In analogy with Eq. (2.27) we can write

$$\begin{aligned} \sum_{t=0}^{\infty} t^r [\vartheta_i(t+1) - \vartheta_i(t)] &= b_i \sum_{t=0}^{\infty} t^r [\vartheta_{i+1}(t) - \vartheta_{i+1}(t-1)] \\ &\quad + d_i \sum_{t=0}^{\infty} t^r [\vartheta_{i-1}(t) - \vartheta_{i-1}(t-1)] \\ &\quad + (1 - b_i - d_i) \sum_{t=0}^{\infty} t^r [\vartheta_i(t) - \vartheta_i(t-1)]. \end{aligned} \quad (2.42)$$

For the left-hand side we write

$$\begin{aligned}
 \sum_{t=0}^{\infty} t^r [\vartheta_i(t+1) - \vartheta_i(t)] &= \sum_{t=1}^{\infty} (t-1)^r [\vartheta_i(t) - \vartheta_i(t-1)] \\
 &= \sum_{t=1}^{\infty} \sum_{k=0}^r \binom{r}{k} (-1)^{r-k} t^k [\vartheta_i(t) - \vartheta_i(t-1)], \\
 &= \tau_i^{(r)} + \sum_{k=0}^{r-1} \binom{r}{k} (-1)^{r-k} \tau_i^{(k)}, \tag{2.43}
 \end{aligned}$$

where $\tau_i^{(0)} = 1$. Thus the r -th moment of the unconditional fixation time is dependent on all lower moments. It satisfies the recursive equation [2, 4, 9, 10]

$$\tau_i^{(r)} = b_i \tau_{i+1}^{(r)} + d_i \tau_{i-1}^{(r)} + (1 - b_i - d_i) \tau_i^{(r)} - s_i^{(r-1)}, \tag{2.44}$$

where $s_i^{(r-1)} = \sum_{k=0}^{r-1} \binom{r}{k} (-1)^{r-k} \tau_i^{(k)}$. The moments also satisfy the boundary conditions $\tau_0^{(r)} = \tau_N^{(r)} = 0$.

Equation (2.44) is of the same form as Eq. (2.36). As a result we can immediately write down the closed-form expressions

$$\tau_1^{(r)} = \phi_1 \sum_{k=1}^{N-1} \sum_{\ell=1}^k \frac{-s_{\ell}^{(r-1)}}{b_{\ell}} \prod_{m=\ell+1}^k \gamma_m, \tag{2.45a}$$

$$\tau_i^{(r)} = -\tau_1^{(r)} \sum_{k=i}^{N-1} \prod_{m=1}^k \gamma_m + \sum_{k=i}^{N-1} \sum_{\ell=1}^k \frac{-s_{\ell}^{(r-1)}}{b_{\ell}} \prod_{m=\ell+1}^k \gamma_m. \tag{2.45b}$$

The derivation of conditional fixation time moments follows the same calculation. Writing $\theta_i^{(r)} = \phi_i \tau_{i|A}^{(r)}$, where $\tau_{i|A}^{(r)}$ is the r -th moment of the conditional fixation time distribution, and $s_{i|A}^{(r-1)} = \sum_{k=0}^{r-1} \binom{r}{k} (-1)^{r-k} \tau_{i|A}^{(k)}$, this quantity satisfies

$$\theta_1^{(r)} = \phi_1 \sum_{k=1}^{N-1} \sum_{\ell=1}^k \frac{-\phi_{\ell} s_{\ell|A}^{(r-1)}}{b_{\ell}} \prod_{m=\ell+1}^k \gamma_m, \tag{2.46a}$$

$$\theta_i^{(r)} = -\theta_1^{(r)} \sum_{k=i}^{N-1} \prod_{m=1}^k \gamma_m + \sum_{k=i}^{N-1} \sum_{\ell=1}^k \frac{-\phi_{\ell} s_{\ell|A}^{(r-1)}}{b_{\ell}} \prod_{m=\ell+1}^k \gamma_m. \tag{2.46b}$$

The conditional fixation time moments are then found via $\tau_{i|A}^{(r)} = \theta_i^{(r)} / \phi_i$.

Once the moments have been calculated, the distribution can be recovered by considering the moment generating function. This is defined as the Laplace² transform of the arrival time density [2] and is denoted as $F_i(z)$, where z is a continuous (and, in general, complex) variable. For the unconditional arrival time we have

$$\begin{aligned}
 F_i(z) &= \sum_{t=0}^{\infty} e^{-zt} [\vartheta_i(t) - \vartheta_i(t-1)] \\
 &= \sum_{t=0}^{\infty} \sum_{j=0}^{\infty} \frac{(-zt)^j}{j!} [\vartheta_i(t) - \vartheta_i(t-1)] \\
 &= \sum_{j=0}^{\infty} \frac{(-z)^j}{j!} \sum_{t=0}^{\infty} t^j [\vartheta_i(t) - \vartheta_i(t-1)] \\
 &= \sum_{j=0}^{\infty} \frac{(-z)^j}{j!} \tau_i^{(j)}.
 \end{aligned} \tag{2.47}$$

The arrival time density is then recovered by performing the inverse Laplace transform of Eq. (2.47) [1, 2],

$$\vartheta_i(t) - \vartheta_i(t-1) = \frac{1}{2\pi i} \int e^{zt} F_i(z) dz. \tag{2.48}$$

In practice one computes a finite set of moments (say the first one hundred) and sums these to obtain $F_i(z)$, which can then be inverted. In Chap. 4 we go beyond this approach and compute closed-form expressions for the exact arrival-time distributions.

2.5 Equilibration

Sometimes the problem we are studying will not have absorbing boundaries. This may be the case if mutations can happen during reproduction events, or due to migration from an external population. In these scenarios, an extinct type of individual may be spontaneously reintroduced. In birth–death processes, this corresponds to having a non-zero birth rate from state $i = 0$ and non-zero death rate from state $i = N$, i.e. $b_0 > 0$ and $d_N > 0$. We retain the conditions $d_0 = 0$ and $b_N = 0$ such that the state-space is still restricted to $0 \leq i \leq N$. Thus if the population is in state $i = 0$, the only possible transition is to state $i = 1$ with rate b_0 . Hence, this is called a reflecting boundary [2]. The same is true at the opposite boundary, where $N \rightarrow N-1$ with rate d_N is the only transition from state $i = N$.

²Pierre-Simon Laplace (1749–1827).

Such problems are no longer characterised by the fixation probability, but by the stationary distribution of the population state. The timescale of this problem is the time taken to approach this distribution, and is known as the mixing time [11].

Stationary Distribution

If $P_i(t)$ is the probability to find the system in state i at time t , then the stationary distribution is given by $P_i^{\text{st}} = \lim_{t \rightarrow \infty} P_i(t)$. We can take the infinite-time limit of the forward master equation (2.14) [or equivalently the discrete-time equation (2.17)] to give

$$0 = b_{i-1}P_{i-1}^{\text{st}} + d_{i+1}P_{i+1}^{\text{st}} - (b_i + d_i)P_i^{\text{st}} \text{ for } 0 \leq i \leq N, \quad (2.49)$$

where we have used $\lim_{t \rightarrow \infty} \dot{P}_i(t) = 0$. Equation (2.49) can be written in the form [1, 2]

$$b_{i-1}P_{i-1}^{\text{st}} - d_iP_i^{\text{st}} = b_iP_i^{\text{st}} - d_{i+1}P_{i+1}^{\text{st}}, \quad (2.50)$$

which states that the net flow of probability from state $i - 1$ to state i is the same for all i . Considering the boundary terms, this net flow must be zero, i.e.

$$b_{i-1}P_{i-1}^{\text{st}} - d_iP_i^{\text{st}} = 0 \text{ for } 1 \leq i \leq N. \quad (2.51)$$

This equation can be recursively solved to find P_i^{st} , which satisfies

$$P_i^{\text{st}} = \frac{b_{i-1}}{d_i} P_{i-1}^{\text{st}} = \left(\prod_{j=1}^i \frac{b_{j-1}}{d_j} \right) P_0^{\text{st}}. \quad (2.52)$$

The value of P_0^{st} is then determined by the normalisation $\sum_{j=0}^N P_j^{\text{st}} = 1$. This gives

$$P_0^{\text{st}} = \left(\sum_{j=0}^N \prod_{k=1}^j \frac{b_{k-1}}{d_k} \right)^{-1}. \quad (2.53)$$

Combining this with Eq. (2.52) gives

$$P_i^{\text{st}} = \frac{\Gamma_i}{\sum_{j=0}^N \Gamma_j}, \quad \Gamma_i = \prod_{j=1}^i \frac{b_{j-1}}{d_j}. \quad (2.54)$$

Mixing Time

The timescale of the dynamics is characterised by the so-called mixing time, t_{mix} [11, 12]. This is the time taken for the probability distribution, $\mathbf{P}(t)$, to come within a specified distance of the stationary distribution $\mathbf{P}^{\text{st}}$. That is to say t_{mix} is the first time at which $d[\mathbf{P}(t_{\text{mix}}), \mathbf{P}^{\text{st}}] = \varepsilon$. The distance between two distributions $\mathbf{P}$ and $\mathbf{Q}$ commonly used in this context is [11, 12]

$$d[\mathbf{P}, \mathbf{Q}] = \frac{1}{2} \sum_{i=0}^N |P_i - Q_i|. \quad (2.55)$$

The factor $1/2$ is used so that distributions which are maximally different have a distance of one. For example, if $P_i = \delta_{i,j}$ and $Q_i = \delta_{i,k}$, where $\delta_{\alpha,\beta}$ is the Kronecker delta and $j \neq k$, then

$$d[\mathbf{P}, \mathbf{Q}] = \frac{1}{2} \sum_{i=0}^N |P_i - Q_i| = \frac{1}{2} \sum_{i=0}^N |\delta_{i,j} - \delta_{i,k}| = \frac{1}{2} (|1| + |-1|) = 1. \quad (2.56)$$

Example

In Fig. 2.2 we show the approach to equilibrium for a simple process: A population of size N consists of two types of individual, A and B , with i individuals of type A . An individual is randomly chosen in the population. If that individual is of type- A (chosen with probability i/N), it switches to type- B with rate μN . If the chosen individual is of type- B (probability $(N - i)/N$), it can switch to type- A with rate λN . Hence the birth and death reaction rates are

$$b_i = \lambda \times (N - i), \quad (2.57a)$$

$$d_i = \mu \times i. \quad (2.57b)$$

These reaction rates are extensive, and here we are considering a continuous-time setup.

As shown in Fig. 2.2a, soon after the dynamics has started the distribution broadens and resembles a normal distribution. It approaches the stationary distribution (labelled $t = \infty$) at a decreasing rate. This is highlighted by considering the distance




Fig. 2.2 **a** The stationary distribution of the birth–death process described by Eq. (2.57). Distributions are found from numeric integration of the master equation (2.14). The stationary distribution ($t = \infty$) is given by Eq. (2.54). **b** The distance between $\mathbf{P}(t)$ and the stationary distribution (2.54) calculated using Eq. (2.55). Here t_{mix} is illustrated for $\varepsilon = 0.1$. The parameters are $\lambda = 0.2$, $\mu = 0.1$, and $N = 100$, and the process was initialised from $P_i(0) = \delta_{i,10}$

to the stationary distribution using Eq. (2.55), and is shown in Fig. 2.2b. The slope of this approach corresponds to the slowest eigenvalue of the master equation.³ To specify the time-to-equilibration we must impose a threshold, ε , as described above. This is illustrated in Fig. 2.2b.

2.6 Deterministic Dynamics and Stability Analysis

In the above example the mean value of i at time t is given approximately by the location of the peak of the distribution at that time. This mean evolves according to an ordinary differential equation (ODE). We define the mean value of i at time t as $\langle i \rangle = \sum_{i=0}^N i P_i(t)$, where the angle brackets indicate an average over many realisations of the stochastic process up to time t . The ODE for the evolution of the mean is found by multiplying the master equation (2.14) by i and summing over all allowed values, such that

$$\sum_{i=0}^N i \dot{P}_i(t) = \sum_{i=0}^N [i b_{i-1} P_{i-1}(t) + i d_{i+1} P_{i+1}(t) - i(b_i + d_i) P_i(t)]. \quad (2.58)$$

Using the boundary conditions $d_0 = b_N = 0$, and the definition of $\langle i \rangle$, we can write Eq. (2.58) as

$$\frac{d\langle i \rangle}{dt} = \sum_{i=0}^N (b_i - d_i) P_i(t) = \langle b_i \rangle - \langle d_i \rangle. \quad (2.59)$$

We can now introduce the continuous variable $x = \lim_{N \rightarrow \infty} i/N$, such that $0 \leq x \leq 1$. This is the deterministic limit; the distribution of x approaches a delta-function centred at $\lim_{N \rightarrow \infty} \langle i \rangle / N$, and hence the evolution of x is completely specified by the ODE for the mean. This ODE is found by dividing Eq. (2.59) by N and taking the limit $N \rightarrow \infty$, which gives

$$\dot{x} = f_+(x) - f_-(x), \quad (2.60)$$

where we have introduced the functions


$$f_+(x) = \lim_{N \rightarrow \infty} \frac{\langle b_i \rangle}{N} = \lim_{N \rightarrow \infty} \frac{b_{\langle i \rangle}}{N} = \lim_{N \rightarrow \infty} \frac{b_{Nx}}{N} \quad (2.61a)$$

$$f_-(x) = \lim_{N \rightarrow \infty} \frac{d_{Nx}}{N}. \quad (2.61b)$$

The step $\langle b_i \rangle = b_{\langle i \rangle}$ does not hold in general as b_i could be a non-linear function of i ; only in the deterministic limit can we make the assumption $\langle i^m \rangle = \langle i \rangle^m$.

³Specifically of the matrix $\mathbb{W}$ shown in Eq. (2.15).

Fig. 2.3 The phase portrait of Eq. (2.62) which describes the time evolution of the mean number of type-A individuals in the birth-death process specified by reactions (2.57). As in Fig. 2.2, the parameters are $\lambda = 0.2$ and $\mu = 0.1$



For the reaction scheme in Eq. (2.57), the deterministic equation is simply

$$\dot{x} = \lambda(1 - x) - \mu x = \lambda - (\lambda + \mu)x. \quad (2.62)$$

Plotting the phase portrait of this equation reveals the flow of the system, as shown in Fig. 2.3. When $\dot{x} > 0$, the flow is to the right, and when $\dot{x} < 0$ the flow is to the left. There exists a point at which $\dot{x} = 0$ (solid circle in Fig. 2.3). Setting $\dot{x} = 0$ in Eq. (2.62) and solving for x recovers $x^* = \lambda/(\lambda + \mu)$. This is a fixed point [13]. In this example it is stable as the flow from both sides is towards it. This stability can be further illustrated by direct integration of Eq. (2.62) which gives

$$x(t) = x(0)e^{-(\lambda+\mu)t} + \frac{\lambda}{\lambda + \mu} [1 - e^{-(\lambda+\mu)t}]. \quad (2.63)$$


From this we can see $\lim_{t \rightarrow \infty} x(t) = x^*$.

For more general problems with non-linear reaction rates and in higher dimensions, determining the stability of fixed points is not quite as simple. In such problems the approach is to linearise the problem about the fixed point. Hence this method is known as linear stability analysis [13].

The deterministic equations of motion obtained from the master equation (2.13) can be recovered by following an analogous procedure to the one described above. Multiplying Eq. (2.13) by $\mathbf{n}$ and summing over $\mathbf{n}$ gives

$$\begin{aligned} \sum_{\mathbf{n}} \mathbf{n} \dot{P}_{\mathbf{n}} &= \sum_{\mathbf{n}} \sum_{\mathbf{v}} \mathbf{n} [T_{\mathbf{n}-\mathbf{v}}^{\mathbf{v}} P_{\mathbf{n}-\mathbf{v}} - T_{\mathbf{n}}^{\mathbf{v}} P_{\mathbf{n}}] \\ &= \sum_{\mathbf{v}} \left[\sum_{\mathbf{n}} (\mathbf{n} + \mathbf{v}) T_{\mathbf{n}}^{\mathbf{v}} P_{\mathbf{n}} - \sum_{\mathbf{n}} \mathbf{n} T_{\mathbf{n}}^{\mathbf{v}} P_{\mathbf{n}} \right] \\ &= \sum_{\mathbf{v}} \mathbf{v} \left[\sum_{\mathbf{n}} T_{\mathbf{n}}^{\mathbf{v}} P_{\mathbf{n}} \right] \\ \Rightarrow \frac{d \langle \mathbf{n} \rangle}{dt} &= \sum_{\mathbf{v}} \mathbf{v} \langle T_{\mathbf{n}}^{\mathbf{v}} \rangle. \end{aligned} \quad (2.64)$$

Fig. 2.4 The stability of fixed points in two dimensions can be classified in terms of the determinant [$\Delta = \det(\mathbb{J})$], trace [$\tau = \text{Tr}(\mathbb{J})$], and discriminant ($\tau^2 - 4\Delta$) of the Jacobian $\mathbb{J}$. Figure is adapted from Ref. [13]



Dividing by the system size N and taking the limit $N \rightarrow \infty$ gives

$$\dot{\mathbf{x}} = \sum_v v f_v(\mathbf{x}) = \mathbf{A}(\mathbf{x}), \quad (2.65)$$

where $\mathbf{x} = \lim_{N \rightarrow \infty} \mathbf{n}/N$ and $f_v(\mathbf{x}) = \lim_{N \rightarrow \infty} T_{N\mathbf{x}}^v/N$.

Fixed points are found by solving $\mathbf{A}(\mathbf{x}^*) = \mathbf{0}$. By expanding Eq. (2.65) about $\mathbf{x}^*$ to first order in $\mathbf{u} = \mathbf{x} - \mathbf{x}^*$, we arrive at

$$\dot{\mathbf{u}} = \mathbb{J} \cdot \mathbf{u}, \quad (2.66)$$

where $\mathbb{J}_{i,j} = \partial A_i / \partial x_j$ is evaluated at $\mathbf{x} = \mathbf{x}^*$ and is known as the Jacobian.⁴ This equation is linear in $\mathbf{u}$ and can be solved by expressing $\mathbf{u}$ as a linear combination of the eigenvectors of $\mathbb{J}$, such that

$$\mathbf{u}(t) = \sum_i c_i \mathbf{v}^{(i)} e^{\lambda_i t}, \quad (2.67)$$

where λ_i are the eigenvalues of $\mathbb{J}$ with corresponding right-eigenvectors $\mathbf{v}^{(i)}$. The coefficients c_i are determined by the initial condition. If all eigenvalues have negative real parts, then $\lim_{t \rightarrow \infty} \mathbf{u}(t) = \mathbf{0}$, and $\mathbf{x}$ will approach $\mathbf{x}^*$, which is a stable fixed point, or attractor. If all eigenvalues have positive real parts then $\mathbf{u}$ will diverge and $\mathbf{x}^*$ is an unstable fixed point. If the set of eigenvalues contains a mixture of positive and negative real values, then there will be divergence in the eigendirections associated with the positive eigenvalues and $\mathbf{x}^*$ is a saddle point. Non-zero imaginary parts of the eigenvalues are associated with cycles and spirals, but these do not feature in this thesis [13].

The case of a two-dimensional system will feature in Chap. 5. In this case we can infer the eigenvalue properties from the determinant (Δ), trace (τ), and discriminant ($\tau^2 - 4\Delta$) of the Jacobian $\mathbb{J}$, as characterised in Fig. 2.4.

⁴After Carl Gustav Jacob Jacobi (1804–1851), who will reappear in Chaps. 5 and 6.

2.7 Continuous State-Space Approximations

Exact solutions of the master equation are usually non-existent; for most cases it is analytically intractable. However, there exists a collection of approximation schemes that allow us to make further analytical progress. These schemes involve replacing the discrete variables $\mathbf{n}$ by their continuous analogues $\mathbf{x} = \mathbf{n}/N$. This allows the master equation to be approximated by a partial differential equation (PDE), namely the Fokker–Planck equation.⁵

Fokker–Planck Equation

The simplest way to arrive at a PDE that describes the probability density is to directly substitute $\mathbf{x} = \mathbf{n}/N$ into the master equation (2.13). We can introduce the function $f_v(\mathbf{x}) = T_{N\mathbf{x}}^v/N$, which is reminiscent of the function introduced in Eq. (2.65) without imposing the infinite system-size limit. We also introduce the continuous probability density $\rho(\mathbf{x}, t)$ to replace the discrete distribution $P_{\mathbf{n}}(t)$. With this we can rewrite the master equation (2.13) as

$$\dot{\rho}(\mathbf{x}, t) \approx \sum_v \left[N f_v \left(\mathbf{x} - \frac{\mathbf{v}}{N} \right) \rho \left(\mathbf{x} - \frac{\mathbf{v}}{N}, t \right) - N f_v(\mathbf{x}) \rho(\mathbf{x}, t) \right]. \quad (2.68)$$

If we now consider N to be a large (but not infinite) parameter, then we can expand the above equation as a Taylor series about $\mathbf{x}$. Truncating the expression up to terms $\mathcal{O}(N^{-1})$, we can write

$$\begin{aligned} \dot{\rho}(\mathbf{x}, t) &= \sum_v \left\{ - \sum_i \frac{\partial}{\partial x_i} [v_i f_v(\mathbf{x}) \rho(\mathbf{x}, t)] + \frac{1}{2N} \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} [v_i v_j f_v(\mathbf{x}) \rho(\mathbf{x}, t)] \right\} \\ &= - \sum_i \frac{\partial}{\partial x_i} [A_i(\mathbf{x}) \rho(\mathbf{x}, t)] + \frac{1}{2N} \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} [\mathbb{B}_{i,j}(\mathbf{x}) \rho(\mathbf{x}, t)], \end{aligned} \quad (2.69)$$

where $\mathbf{A}(\mathbf{x}) = \sum_v \mathbf{v} f_v(\mathbf{x})$ as described in Eq. (2.65), and where the diffusion matrix $\mathbb{B}(\mathbf{x})$ has elements

$$\mathbb{B}_{i,j} = \sum_v v_i v_j f_v(\mathbf{x}). \quad (2.70)$$

Equation (2.69) is called the Fokker–Planck equation. The first term is representative of the deterministic *drift*, and the second *diffusion* term represents the spread of the probability density due to the stochasticity [14]. There is a complication that needs to be considered with this derivation; we need to ensure that the equation maintains a positive probability density. Including a higher number of terms in Eq. (2.69) [$\mathcal{O}(N^{k \geq 2})$] would violate this condition unless the infinite series is considered. This is the Pawula theorem [15, 16].

⁵Adriaan Fokker (1877–1972) and Max Planck (1858–1947).

This derivation is closely related to two rigorous approximation schemes. These are the Kramers–Moyal⁶ expansion [17, 18] which considers a truncated series of jump moments [14], and the van Kampen⁷ system-size expansion which explicitly considers a large parameter to recover an equation which describes the fluctuations about the deterministic trajectory (2.65) [1]. The Fokker–Planck approximation of the master equation does not feature prominently in this thesis, but for further discussion about its derivation from the master equation see Refs. [1, 14].

Stochastic Differential Equation

The process described by the master equation (2.13) is a discrete jump process. The Fokker–Planck equation (2.69) describes the evolution of a continuous variable that is affected by some noise [1]. We write this as

$$\dot{\mathbf{x}} = \mathbf{A}(\mathbf{x}) + \frac{1}{\sqrt{N}} \mathbb{G}(\mathbf{x}) \cdot \boldsymbol{\eta}(t), \quad (2.71)$$

where $\mathbb{B} = \mathbb{G} \cdot \mathbb{G}^T$, and the $\eta_i(t)$ random variables drawn from a Gaussian distribution with zero mean and correlation function

$$\langle \eta_i(t) \eta_j(t') \rangle = \delta_{i,j} \delta(t - t'). \quad (2.72)$$

Equation (2.71) is known as a stochastic differential equation (SDE) or a Langevin equation, after Paul Langevin (1872–1946). It is an extension to the deterministic equation (2.65), which can be recovered by taking the limit $N \rightarrow \infty$ in Eq. (2.71).

Again there are complications with this derivation, including the famous Itô–Stratanovich⁸ dilemma [1, 14, 16]. This states that Eq. (2.71) alone is not well-specified if $\mathbb{G}$ is dependent on the state of the system $\mathbf{x}$. This is termed multiplicative noise, and if this is the case it must be stated whether Eq. (2.71) is to be interpreted in the Itô or Stratanovich sense. Further discussion of this issue here would only pose as a distraction and would be of no value for the remainder of this thesis. More details can be found in Refs. [1, 14, 16]. SDEs only appear in this thesis in Chap. 6, and we will always use the Itô interpretation. Equation (2.71) should also be interpreted as an Itô SDE.

2.8 Evolutionary Game Theory

The interactions between individuals in a population can sometimes be more complicated than A “dies with constant rate μ ”. Instead, this interaction rate may have a non-linear dependence on the state of the population. These interactions can be

⁶Hendrik Kramers (1894–1952) and José Enrique Moyal (1910–1998).

⁷Nico van Kampen (1921–2013).

⁸Kiyoshi Itô (1915–2008) and Ruslan Stratanovich (1930–1997).

formalised in an evolutionary game [8, 19–22]. Such games can be used to describe conflict over food or territory, cheating in resource allocation, as well as interactions between variants of a gene [10, 23–26]. In an evolutionary normal-form game each individual can be associated with one out of a finite set of strategies. A payoff matrix quantifies the reward received by a given individual when it interacts with another individual [21].

The dynamics of populations interacting in such a game are often described by deterministic replicator equations (discussed at the end of this section) or similar differential equations [20, 22, 27]. While deterministic dynamics are useful to understand the action of selection, a stochastic approach is required to understand the impact of fluctuations in finite populations [28, 29]. Deterministic approaches fail to capture effects such as fixation and extinction, or the convergence to a stationary distribution in systems with mutation [3, 12, 30–32].

The interaction between two individuals in a two-strategy evolutionary game is characterised by the payoff matrix,

$$\begin{array}{c|cc} & A & B \\ \hline A & R & S \\ B & T & P \end{array} \quad (2.73)$$

A type-*A* individual encountering another of its kind receives the reward *R*, and it receives the sucker's payoff *S* when interacting with a type-*B* individual. In turn, an individual of type *B* interacting with an individual of type *A* obtains the temptation payoff *T*, and *P* is the punishment payoff for each individual if they are both of type *B*.

If there are *i* individuals of type *A* in the population and *N* − *i* individuals of type *B*, the expected payoffs for each type of player are

$$\pi_A(i) = \frac{i-1}{N-1}R + \frac{N-i}{N-1}S, \quad (2.74a)$$

$$\pi_B(i) = \frac{i}{N-1}T + \frac{N-i-1}{N-1}P. \quad (2.74b)$$

There exist three general types of two-player two-strategy evolutionary games which can be described by the payoff matrix (2.73) and payoffs (2.74):

Dominance: If $R > T$ and $S > P$, then playing strategy *A* will return a higher payoff irrespective of the composition of the population. This type is then always favoured. Likewise, if the inequalities are reversed such that $T > R$ and $P > S$, then playing strategy *B* will always return a higher payoff. This latter scenario captures the well-studied Prisoner's dilemma [8]. This also leads to the concept of Nash equilibria, after John Nash (1928–2015). A Nash equilibrium exists if, with all players playing the same strategy, no single deviation from that strategy can lead to an increase in payoff [8, 33]. In the Prisoner's dilemma, the temptation to cheat (strategy *B*) is greater than the reward to cooperate with another cooperator

(strategy A), and the punishment for mutual defection is a better option than receiving the sucker's payoff by cooperating with a defector. Hence defection, or strategy B , is the Nash strategy in this game.

Coexistence: If $T > R$ and $S > P$, then choosing the opposite strategy to the other player yields the highest payoff. Hence the payoff is maximised in a heterogeneous population. There exists no pure-strategy Nash equilibrium in this setup.

Coordination: If $R > T$ and $P > S$, then payoff is maximised when both players adopt the same strategy. In this case there is a bi-stability and both strategies are Nash equilibria [8].

The rate at which strategies spread through the population is a function of the pay-offs in the evolutionary game. A constant parameter $\beta > 0$, the so-called intensity of selection, is introduced to control the how much the game affects the population dynamics. It can be thought of as the inverse temperature. For $\beta \rightarrow \infty$ (low temperature), the dynamics is controlled by the evolutionary game. For $\beta \rightarrow 0$ (high temperature), the dynamics is dominated by stochasticity. The regime of weak selection is interesting as it allows for perturbative treatment to obtain analytic results [34, 35]. Many functional forms of the mapping between payoffs and population dynamics are possible, but the update rules can be divided into two distinct classes [36]:

Pairwise-comparison process: Here two individuals are randomly chosen from the population, and the second adopts the strategy of the first with a probability which is a function of the payoff *difference*. For example, if a type- A individual is chosen first, and a type- B chosen second, then the type- B adopts strategy A with probability determined by $g[\beta \Delta \pi(i)]$, where $\Delta \pi(i) = \pi_A(i) - \pi_B(i)$. The reverse interaction (A adopting strategy B) occurs with probability $g[-\beta \Delta \pi(i)]$. The exact functional dependence has to satisfy the basic properties that $0 \leq g \leq 1$ and that it is a monotonically increasing function of payoff. Common forms for this include the linear mapping [12]

$$g[\pm \beta \Delta \pi(i)] = \frac{1}{2} [1 \pm \beta \Delta \pi(i)], \quad (2.75)$$

or the Fermi⁹ mapping [34]

$$g[\pm \beta \Delta \pi(i)] = \frac{1}{1 + e^{\mp \beta \Delta \pi(i)}}. \quad (2.76)$$

This process leads to birth and death rates (extensive) of the form

$$b_i = \frac{i(N-i)}{N} g[\beta \Delta \pi(i)], \quad (2.77a)$$

$$d_i = \frac{i(N-i)}{N} g[-\beta \Delta \pi(i)], \quad (2.77b)$$

⁹Enrico Fermi (1901–1954).




Fig. 2.5 **a** and **b** show the birth rates and death rates from Eqs. (2.77) and (2.79). **c** shows the total exit rate from state i , which is $b_i + d_i$. **d** shows the ratio of b_i and d_i . For this example, which represents a coexistence game, the payoff matrix parameters are $R = 1.0$, $S = 1.5$, $T = 1.9$, and $P = 1.0$, the intensity of selection is $\beta = 0.5$ and the system size is $N = 20$

which are similar for both the linear and the Fermi mappings as shown in Fig. 2.5a, b. It is easy to show that $g[+\beta\Delta\pi(i)] + g[-\beta\Delta\pi(i)] = 1$ for both Eqs. (2.75) and (2.76). Thus the total rate at which the population exits state i is simply $i(N-i)/i$, as shown in Fig. 2.5c.

Fitness-based process: Here individuals reproduce with a rate determined by the individuals' *own* expected payoff. There is no comparison in this process. After a reproduction event a random individual is removed from the population to conserve the fixed population size. Again the functional form of the payoff dependence must be monotonically increasing. An exponential mapping is a frequent choice, such that an individual's reproductive fitness is given by [37, 38]

$$f_A(i) = e^{\beta\pi_A(i)}, \quad (2.78a)$$

$$f_B(i) = e^{\beta\pi_B(i)}. \quad (2.78b)$$

The birth and death rates are then given by

$$b_i = \frac{i(N-i)}{N} \frac{f_A(i)}{\bar{f}(i)}, \quad (2.79a)$$

$$d_i = \frac{i(N-i)}{N} \frac{f_B(i)}{\bar{f}(i)}, \quad (2.79b)$$

where the fitnesses are scaled by the mean population fitness $\bar{f}(i) = [ie^{\beta\pi_A(i)} + (N-i)e^{\beta\pi_B(i)}]/N$. This ensures events occur on the same timescale as the above pairwise comparison process, i.e. Eq. (2.79) is of similar magnitude to Eq. (2.77), as shown in Fig. 2.5a–c. Interestingly the ratio of birth and death rates is identical for the Fermi pairwise-comparison process and the exponential fitness process, as shown in Fig. 2.5d [36].

The reaction rates in Eqs. (2.77) and (2.79) are the update rules of the Moran¹⁰ process [4, 39] with frequency-dependent selection. These rules have been widely used in evolutionary game theory [5, 8, 40]. The Moran process represents a birth–death process in which the population size remains constant, and by construction it has absorbing states at $i = 0$ and $i = N$. The process can be represented by the reaction scheme



Finally, the above interaction schemes can easily be generalised for larger strategy spaces [8, 21]. However only 2×2 games will appear in this thesis.

Replicator Dynamics

By considering the deterministic dynamics of the process specified by the reactions (2.80), we can write

$$\dot{x} = x(1-x)[\mathcal{F}_A(x) - \mathcal{F}_B(x)], \quad (2.81)$$

where $\mathcal{F}_A(x)$ and $\mathcal{F}_B(x)$ represent the frequency-dependent reaction rates, e.g. $\mathcal{F}_A(x) = g[\beta\Delta\pi(Nx)]$. Expanding this equation gives

$$\begin{aligned} \dot{x} &= x[\mathcal{F}_A(x) - x\mathcal{F}_A(x) - (1-x)\mathcal{F}_B(x)] \\ &= x[\mathcal{F}_A(x) - \bar{\mathcal{F}}(x)]. \end{aligned} \quad (2.82)$$

This is called the replicator equation as the fitness of A approaches the mean population fitness $\bar{\mathcal{F}} = x\mathcal{F}_A(x) + (1-x)\mathcal{F}_B(x)$ [20, 22, 27].

2.9 Numerical Simulation Methods

Numerical simulations of a stochastic process can serve many purposes. They can be used to confirm the correctness of analytic work or verify if assumptions made in the analysis are valid. They can also present some physical intuition about how a system evolves, and this can in turn guide the analytic procedure. The aim is to

¹⁰Pat Moran (1917–1988).

generate random numbers that are distributed according to the solution of the master equation.

Monte-Carlo Methods

For a discrete-time system that is described by the master equation (2.11), the corresponding process is easily described: if the system is in state $\mathbf{n}$ at time t , then at time $t + \Delta t$ the system will be in state $\mathbf{n}' \neq \mathbf{n}$ with probability $w_{\mathbf{n}',\mathbf{n}}\Delta t$, or will remain in state $\mathbf{n}$ with probability $1 - \sum_{\mathbf{n}' \neq \mathbf{n}} w_{\mathbf{n}',\mathbf{n}}\Delta t$. The simulation of this process follows directly from this. Namely, we set $\mathbf{n}(t) = \mathbf{n}$ and evaluate the values of $w_{\mathbf{n}',\mathbf{n}}\Delta t$ and $1 - \sum_{\mathbf{n}' \neq \mathbf{n}} w_{\mathbf{n}',\mathbf{n}}\Delta t$. These values are the weights associated with choosing the value of $\mathbf{n}(t + \Delta t)$, such that only a single random number needs to be drawn.

This discrete-time method, however, suffers from two fundamental flaws. Firstly, only a single reaction can occur in a single time-step, and secondly, a large number of time-steps may pass by without the system changing state. Choosing a very small time-step to cure the first problem amplifies the second, and choosing a long time-step to cure the second problem amplifies the first flaw. To fix these problems we will consider a continuous-time setup where the length of the time-step is a random number that is dependent on the state of the system.

Gillespie Algorithm

Trying to deduce the process that is described by the continuous-time master equation (2.13) is significantly more difficult than the discrete-time case. We now need to answer two questions: which is the next reaction to occur, and at what time does this happen? As both of these quantities are random variables, we need to compute the joint probability density $p(\mathbf{v}, \tau)$, which is the probability density that the next reaction to occur from $\mathbf{n}$ is to $\mathbf{n} + \mathbf{v}$, and this happens after a period of time τ [41]. If we consider the reaction to occur in the time interval $[t + \tau, t + \tau + d\tau]$, then we can write

$$p(\mathbf{v}, \tau)d\tau = P_0(\mathbf{n}, t + \tau | \mathbf{n}, t) T_{\mathbf{n}}^{\mathbf{v}} d\tau, \quad (2.83)$$

where $P_0(\mathbf{n}, t + \tau | \mathbf{n}, t)$ is the probability that no reactions have occurred in the time interval $[t, t + \tau]$. This quantity is found by considering only transitions out of state $\mathbf{n}$ in the master equation (2.13). By relabelling the initial condition in Eq. (2.12), we can write,

$$\frac{dP_0(\mathbf{n}, t + \tau | \mathbf{n}, t)}{d\tau} = - \sum_{\mathbf{v}} T_{\mathbf{n}}^{\mathbf{v}} P_0(\mathbf{n}, t + \tau | \mathbf{n}, t). \quad (2.84)$$

This can be directly integrated to show that $P_0(\mathbf{n}, t + \tau | \mathbf{n}, t)$ decreases exponentially with τ ,

$$P_0(\mathbf{n}, t + \tau | \mathbf{n}, t) = \exp \left[- \left(\sum_{\mathbf{v}} T_{\mathbf{n}}^{\mathbf{v}} \right) \tau \right]. \quad (2.85)$$

The value $\sum_{\mathbf{v}} T_{\mathbf{n}}^{\mathbf{v}} = a_0(\mathbf{n})$ is the total rate at which the system leaves state $\mathbf{n}$. Substituting this into Eq. (2.83) we can write

$$p(\mathbf{v}, \tau) = \frac{T_{\mathbf{n}}^{\mathbf{v}}}{a_0(\mathbf{n})} \times a_0(\mathbf{n}) \exp[-a_0(\mathbf{n})\tau]. \quad (2.86)$$

Hence the probability density $p(\mathbf{v}, \tau)$ can be decomposed into two independent densities: the first describes the choice of reaction with weight $T_{\mathbf{n}}^{\mathbf{v}}/a_0(\mathbf{n})$, and the second describes the time-step which is exponentially distributed. Hence by drawing two independent random numbers for each time-step, we can simulate a process that is distributed exactly as described by the solution of the continuous-time master equation (2.12). This procedure was popularised in Ref. [41], and is known as the Gillespie algorithm.

References

1. N.G. van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, Amsterdam, 2007)
2. N. Goel, N. Richter-Dyn, *Stochastic Models in Biology* (Academic Press, New York, 1974)
3. T. Antal, I. Scheuring, Fixation of strategies for an evolutionary game in finite populations. *Bull. Math. Biol.* **68**, 1923 (2006)
4. W.J. Ewens, *Mathematical Population Genetics. I. Theoretical Introduction* (Springer, New York, 2004)
5. A. Traulsen, C. Hauert, Stochastic evolutionary game dynamics, in *Reviews of Nonlinear Dynamics and Complexity*, vol. II, ed. by H.G. Schuster (Wiley-VCH, Weinheim, 2009)
6. R. Shankar, *Principles of Quantum Mechanics* (Springer, New York, 1994)
7. S. Karlin, H.M. Taylor, *A Second Course in Stochastic Processes* (Academic Press, New York, 1981)
8. M.A. Nowak, *Evolutionary Dynamics* (Harvard University Press, Cambridge, MA, 2006)
9. S. Redner, *A Guide to First-Passage Processes* (Cambridge University Press, Cambridge, UK, 2001)
10. P.M. Altrock, A. Traulsen, F.A. Reed, Stability properties of underdominance in finite subdivided populations. *PLoS Comput. Biol.* **7**, e1002260 (2011)
11. D.A. Levin, Y. Peres, E.L. Wilmer, *Markov Chains and Mixing Times* (AMS Publishing, Providence, RI, 2009)
12. A.J. Black, A. Traulsen, T. Galla, Mixing times in evolutionary games. *Phys. Rev. Lett.* **109**, 028101 (2012)
13. S. Strogatz, *Nonlinear Dynamics and Chaos* (Westview Press, Boulder, CO, 2000)
14. C.W. Gardiner, *Handbook of Stochastic Methods* (Springer, New York, 2009)
15. R. Pawula, Approximation of the linear Boltzmann equation by the Fokker-Planck equation. *Phys. Rev.* **162**, 186 (1967)
16. H. Risken, *The Fokker-Planck Equation* (Springer, Berlin, 1989)
17. H.A. Kramers, Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica* **7**, 284 (1940)
18. J. Moyal, Stochastic processes and statistical physics. *J. R. Stat. Soc.* **11**, 151 (1949)
19. J.W. Weibull, *Evolutionary Game Theory* (MIT Press, Cambridge, MA, 1995)
20. J. Hofbauer, K. Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge University Press, Cambridge, UK, 1998)
21. H. Gintis, *Game Theory Evolving* (Princeton University Press, Princeton, NJ, 2009)
22. W.H. Sandholm, *Population Games and Evolutionary Dynamics* (MIT Press, Cambridge, MA, 2010)
23. J. Maynard Smith, G.R. Price, The logic of animal conflict. *Nature* **246**, 15 (1973)
24. M. Milinski, Tit For Tat in sticklebacks and the evolution of cooperation. *Nature* **325**, 433 (1987)

25. R.C. MacLean, A. Fuentes-Hernandez, D. Greig, L.D. Hurst, I. Gudelj, A mixture of “cheats” and “co-operators” can enable maximal group benefit. *PLoS Biol.* **8**, e1000486 (2010)
26. A. Traulsen, F.A. Reed, From genes to games: cooperation and cyclic dominance in meiotic drive. *J. Theor. Biol.* **299**, 120 (2012)
27. P.D. Taylor, L. Jonker, Evolutionarily stable strategies and game dynamics. *Math. Biosci.* **40**, 145 (1978)
28. T. Lenormand, D. Roze, F. Rousset, Stochasticity in evolution. *Trends Ecol. Evol.* **24**, 157 (2009)
29. A.J. Black, A.J. McKane, Stochastic formulation of ecological models and their applications. *Trends Ecol. Evol.* **27**, 337 (2012)
30. P.M. Altrock, A. Traulsen, T. Galla, The mechanics of stochastic slowdown in evolutionary games. *J. Theor. Biol.* **311**, 94 (2012)
31. H. Arnoldt, M. Timme, S. Grosskinsky, Frequency-dependent fitness induces multistability in coevolutionary dynamics. *J. R. Soc. Interface* **9**, 3387 (2012)
32. J. Du, B. Wu, P.M. Altrock, L. Wang, Aspiration dynamics of multi-player games in finite populations. *J. R. Soc. Interface* **11**, 20140077 (2014)
33. J.F. Nash, Equilibrium points in n -person games. *Proc. Natl. Acad. Sci. USA* **36**, 48 (1950)
34. P.M. Altrock, A. Traulsen, Fixation times in evolutionary games under weak selection. *New J. Phys.* **11**, 013012 (2009)
35. B. Wu, P.M. Altrock, L. Wang, A. Traulsen, Universality of weak selection. *Phys. Rev. E* **82**, 046106 (2010)
36. B. Wu, B. Bauer, T. Galla, A. Traulsen, Fitness-based models and pairwise-comparison models of evolutionary games are typically different even in unstructured populations. *New J. Phys.* **17**, 023043 (2015)
37. A. Traulsen, N. Shresh, M.A. Nowak, Analytical results for individual and group selection of any intensity. *Bull. Math. Biol.* **70**, 1410 (2008)
38. P.M. Altrock, A. Traulsen, Deterministic evolutionary game dynamics in finite populations. *Phys. Rev. E* **80**, 011909 (2009)
39. P.A.P. Moran, *The Statistical Processes of Evolutionary Theory* (Clarendon Press, Oxford, UK, 1962)
40. C. Taylor, D. Fudenberg, A. Sasaki, M.A. Nowak, Evolutionary game dynamics in finite populations. *Bull. Math. Biol.* **66**, 1621 (2004)
41. D.T. Gillespie, Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* **81**, 2340 (1977)

Chapter 3

Finite Populations in Switching Environments

3.1 Introduction

As discussed in the previous section, selection acts on different phenotypes, such as ‘resident’ and ‘mutant’, and changes the population composition. The environment in which a population evolves determines the direction of selection. Changes in the state of the environment can alter these selective pressures such that different phenotypes are selected for as time progresses. Characterising the evolutionary dynamics of the population in such a system is non-trivial, but we will shed light on this issue in this chapter. This work originally appeared in [1], although some of the notation has been changed to fit in with the rest of this thesis.

Time-varying environments are relevant in the evolution of bacterial populations subject to environment modulations by a host [2, 3], or varying antibiotic stress [4]. An illustrative example is the evolution of normal ‘sensitive’ cells and resistant ‘persister’ cells. A stochastic model of this was examined by Kussell et al. [5], where periods of antibiotics were turned on and off. During times of antibiotic stress the growth rate of normal cells was reduced, but the resistant cells sustain population levels. Stochastic switching between the phenotypes ensured that any lost phenotypes could be reintroduced. This can be seen from the simulation trajectory of this process shown in Fig. 3.1.

Further mathematical analysis was carried out in the large-population limit in the absence of intrinsic stochasticity. The population dynamics are then controlled by the differential equation [5]

$$\frac{d\mathbf{n}}{dt} = \mathbb{A}_{\varepsilon(t)} \cdot \mathbf{n}(t), \quad (3.1)$$

where $\mathbf{n}(t)$ is the population vector that describes the number of individuals of each species at time t , and $\varepsilon(t)$ is the state of the environment which takes discrete values. The environment can be periodic as shown in Fig. 3.1, or more generally can follow a continuous-time stochastic process. The matrix $\mathbb{A}_{\varepsilon(t)}$ describes the growth of the




Fig. 3.1 A stochastic trajectory of a population subjected to periodic antibiotic stress. Gray horizontal bars indicate the absence of the drug, and white spaces indicate periods of antibiotics. Dark blue lines represent the number of wild-type cells which are sensitive to the drug (and hence their numbers decrease during antibiotics). The pale blue line represents the number of resistant or persister phenotypes, which are able to rescue the population from extinction. Orange and green lines should be ignored in this instance. Figure is taken from Ref. [5]

number of individuals of each species, including changes from one species to another, when the environment is in state $\varepsilon(t)$.

The model (3.1) was supported by Acar et al. [6] who provided experimental evidence by modifying the environment in which two strains of yeast grow. More complicated studies of dynamics in switching environments rely on cells ‘sensing’ the environment [7] and on the history or information of the environment that is collected during a cell’s lifetime [8, 9]. These examples illustrate that the assumption of an interaction structure independent of time is not always realistic. At the same time it is largely an open question how complex interactions between phenotypes together with spontaneous changes in the environment influence the evolutionary dynamics.

While deterministic dynamics are useful to understand the action of selection, a stochastic approach is necessary to capture effects such as fixation and extinction. Environmental variability in stochastic systems has been investigated in predator–prey models [10]. It has also been studied in the context of evolutionary games, where (continuous) extrinsic noise is added to model parameters [11].

Rather than selecting a specific form for the dynamics, we will use the generic birth–death framework such that our results apply to a wide class of population dynamics. For the environmental dynamics, we will follow the work of Kussell et al. [5] and consider an environment that stochastically switches between a discrete set

of states. It is convenient to describe the dynamics in discrete time, as will be seen below. The model we use is described in detail in Sect. 3.2. In Sect. 3.3 the theory is developed that allows us to calculate fixation probabilities and mean fixation times of a rare mutant under fluctuating environmental conditions. We then expand on the two-environment scenario where further analytical progress can be made. To illustrate our theoretical results we study the fixation properties in an evolutionary game that stochastically switches between a coexistence game and a coordination game in Sect. 3.4. We determine environmental conditions under which the success of a rare invading mutant is maximal. This is seen to occur at a non-trivial combination of switching rates.

If we introduce mutations into the dynamics, the possibility of fixation and extinction is removed as a lost species can be reintroduced. Instead we seek to describe the stationary distribution of a population evolving in a switching environment in Sect. 3.5. We derive approximations for this stationary distribution, which we show are valid for a large range of switching rates in the two-environment scenario.

3.2 Model

We consider populations consisting of a fixed number, N , of individuals. Each individual can be of one of two types, A or B , which we refer to as ‘mutant’ and ‘resident’, respectively. The population is well mixed; every individual can interact with any other individual. The state of the population (without considering the environmental dynamics) is fully characterised by the number, i , of individuals of type A . The remaining $N - i$ individuals are then of type B . We furthermore assume that at any one time the environment can be in one of Ω discrete states, labelled $\sigma \in \Lambda$, where Λ is the space of states of the environment ($|\Lambda| = \Omega$). Hence the state of the entire system at any time is given by the pair (i, σ) .

The discrete-time birth–death dynamics of the population for a given environment, σ , is specified by the transition probabilities $b_i^{(\sigma)}$ and $d_i^{(\sigma)}$ of a one-step process. This is analogous to the birth and death rates discussed in Sect. 2.3. Specifically, if the system is in state (i, σ) the population transitions to state $i + 1$ in the next time-step with probability $b_i^{(\sigma)}$. Similarly the state of the population in the next time step is $i - 1$ with probability $d_i^{(\sigma)}$. These transitions are shown as short curved arrows in Fig. 3.2. With probability $1 - b_i^{(\sigma)} - d_i^{(\sigma)}$ the population remains in state i . We always assume that $b_i^{(\sigma)} \geq 0$, $d_i^{(\sigma)} \geq 0$, and $b_i^{(\sigma)} + d_i^{(\sigma)} \leq 1$ for all (i, σ) . Throughout this chapter we will use the time-step $\Delta t = 1$, as described in Sect. 2.3. The quantity t then counts the number of time-steps. All times shown in figures will be expressed in generations, i.e. t/N .

For now we will assume that the states $i = 0$ (all-type- B) and $i = N$ (all-type- A) are absorbing, i.e. $b_0^{(\sigma)} = 0$ and $d_N^{(\sigma)} = 0$ for all $\sigma \in \Lambda$. In the absence of further mutation events a type, once absent, can never be re-introduced. If mutations occur




Fig. 3.2 A population undergoes a one-step birth–death process, such that given the population is in state i , in one time-step it may transition to $i - 1$ or $i + 1$, or remain at i . The states $i = 0$ and $i = N$ are absorbing in all environments (no arrows out of these states). The transition probabilities are dependent on the state of the environment, indicated by *solid* versus *dashed* arrows in environments σ and σ' , respectively. The environment switches from state σ to σ' with probability $\mu_{\sigma', \sigma}$ in any one time-step. The quantity $\phi_i^{(\sigma)}$ represents the probability of fixation, as discussed in Sect. 3.3

during the dynamics, then the states $i = 0$ and $i = N$ are no longer absorbing and the system converges to a unique, non-trivial stationary state. We consider this case in Sect. 3.5.

In our approach the environment evolves from one state to another *independently* of the state of the population. This is a simplification that allows the following analysis to take place, but this model still captures a wide array of natural scenarios. In this discrete-time setup we take the dynamics of the environment to be a simple Markov chain, described by the transition matrix $\mathbb{M}$ of size $\Omega \times \Omega$. The entry $\mu_{\sigma', \sigma}$ of the matrix $\mathbb{M}$ represents the probability that the environment changes to state σ' in the next time-step, if it is currently in state σ , as shown in Fig. 3.2. The matrix $\mathbb{M}$ is a stochastic matrix, $\sum_{\sigma'} \mu_{\sigma', \sigma} = 1$ for all $\sigma \in \Lambda$. In other words all columns sum to one. To ensure this we set $\mu_{\sigma, \sigma} = 1 - \sum_{\sigma' \neq \sigma} \mu_{\sigma', \sigma}$.

If the system is in state (i, σ) at a given time, it may transition to 3Ω possible states in any one time-step. These are given by (i, σ') , $(i + 1, \sigma')$ and $(i - 1, \sigma')$, where $\sigma' \in \Lambda$ can be any of the Ω states of the environment. If we write $R_{(j, \sigma'), (i, \sigma)}$ for the probability of a transition from (i, σ) to (j, σ') , we have

$$R_{(i+1, \sigma'), (i, \sigma)} = \mu_{\sigma', \sigma} b_i^{(\sigma)}, \quad (3.2a)$$

$$R_{(i-1, \sigma'), (i, \sigma)} = \mu_{\sigma', \sigma} d_i^{(\sigma)}, \quad (3.2b)$$

$$R_{(i, \sigma'), (i, \sigma)} = \mu_{\sigma', \sigma} \left(1 - b_i^{(\sigma)} - d_i^{(\sigma)} \right). \quad (3.2c)$$

No transitions from (i, σ) to (j, σ') can occur when $|i - j| > 1$. In this setup the birth and death probabilities are determined by the state of the environment at the *beginning* of the discrete time-step.

In analogy with Eq. (2.17), the master equation for the probability to be in state (i, σ) is given by

$$\begin{aligned}
 P_{(i,\sigma)}(t+1) &= \sum_{\sigma' \in \Lambda} [R_{(i,\sigma),(i-1,\sigma')} P_{(i-1,\sigma')}(t) + R_{(i,\sigma),(i+1,\sigma')} P_{(i+1,\sigma')}(t) + R_{(i,\sigma),(i,\sigma')} P_{(i,\sigma')}(t)] \\
 &= \sum_{\sigma' \in \Lambda} \mu_{\sigma,\sigma'} \left[b_{i-1}^{(\sigma')} P_{(i-1,\sigma')}(t) + d_{i+1}^{(\sigma')} P_{(i+1,\sigma')}(t) + \left(1 - b_i^{(\sigma')} - d_i^{(\sigma')}\right) P_{(i,\sigma')}(t) \right].
 \end{aligned} \tag{3.3}$$

The backward equation, in analogy with Eq. (2.20), is

$$\begin{aligned}
 Q_{(j,\sigma''); (i,\sigma)}(t+1) &= \sum_{\sigma' \in \Lambda} \left[R_{(i+1,\sigma'), (i,\sigma)} Q_{(j,\sigma''); (i+1,\sigma')}(t) + R_{(i-1,\sigma'), (i,\sigma)} Q_{(j,\sigma''); (i-1,\sigma')}(t) \right. \\
 &\quad \left. + R_{(i,\sigma'), (i,\sigma)} Q_{(j,\sigma''); (i,\sigma')}(t) \right] \\
 &= \sum_{\sigma' \in \Lambda} \mu_{\sigma',\sigma} \left[b_i^{(\sigma')} Q_{(j,\sigma''); (i+1,\sigma')}(t) + d_i^{(\sigma')} Q_{(j,\sigma''); (i-1,\sigma')}(t) \right. \\
 &\quad \left. + \left(1 - b_i^{(\sigma')} - d_i^{(\sigma')}\right) Q_{(j,\sigma''); (i,\sigma')}(t) \right].
 \end{aligned} \tag{3.4}$$

3.3 Mathematical Framework

We now demonstrate how to calculate the fixation probability and mean fixation times in birth–death processes with an arbitrary number of environmental states. This framework is based on the discrete-time process. In the associated publication, Ref. [1], we present similar methods for analysing continuous-time processes alongside the discrete-time formulation.

Fixation Probability

The fixation probability, $\phi_i^{(\sigma)}$, is the probability that the system ends up in the absorbing state with N individuals of type A , conditioned on the initial state (i, σ) . The probability of fixation of a single mutant, $\phi_1^{(\sigma)}$, is of particular interest; rare mutations can introduce a previously absent strategy into the population, and typically there is only one individual of this novel type initially. To obtain an expression for the fixation probability, we follow Sect. 2.4 and introduce the $\phi_i^{(\sigma)} = \lim_{t \rightarrow \infty} \sum_{\sigma''} Q_{(N,\sigma''); (i,\sigma)}(t)$. We sum over all final environmental states (σ'') as we are only interested in the probability of type A taking over the population, not the environment in which fixation is reached. From Eq. (3.4) we obtain

$$\phi_i^{(\sigma)} = \sum_{\sigma' \in \Lambda} \mu_{\sigma',\sigma} \left[b_i^{(\sigma')} \phi_{i+1}^{(\sigma')} + d_i^{(\sigma')} \phi_{i-1}^{(\sigma')} + \left(1 - b_i^{(\sigma')} - d_i^{(\sigma')}\right) \phi_i^{(\sigma')} \right]. \tag{3.5}$$

This is to be solved along with the boundary conditions $\phi_0^{(\sigma)} = 0$ and $\phi_N^{(\sigma)} = 1$ for all $\sigma \in \Lambda$.

To obtain a formal solution, we introduce $\psi_i^{(\sigma)} = \sum_{\sigma'} \mu_{\sigma', \sigma} \phi_i^{(\sigma')}$. The boundary conditions $\phi_0^{(\sigma)} = 0$ and $\phi_N^{(\sigma)} = 1$ translate into $\psi_0^{(\sigma)} = 0$ and $\psi_N^{(\sigma)} = 1$ for all $\sigma \in \Lambda$, where the second expression follows from $\mathbb{M}$ being a stochastic matrix. With this notation we have

$$\phi_i^{(\sigma)} = b_i^{(\sigma)} \left(\psi_{i+1}^{(\sigma)} - \psi_i^{(\sigma)} \right) - d_i^{(\sigma)} \left(\psi_i^{(\sigma)} - \psi_{i-1}^{(\sigma)} \right) + \psi_i^{(\sigma)}. \quad (3.6)$$

In matrix form we can write $\boldsymbol{\psi}_i = \boldsymbol{\phi}_i \cdot \mathbb{M}$, where $\boldsymbol{\psi}_i$ and $\boldsymbol{\phi}_i$ are row vectors with Ω components.¹ Using $\boldsymbol{\phi}_i = \boldsymbol{\psi}_i \cdot \mathbb{M}^{-1}$, we obtain

$$\left(\psi_{i+1}^{(\sigma)} - \psi_i^{(\sigma)} \right) = \gamma_i^{(\sigma)} \left(\psi_i^{(\sigma)} - \psi_{i-1}^{(\sigma)} \right) + \frac{1}{b_i^{(\sigma)}} \left[(\boldsymbol{\psi}_i \cdot \mathbb{M}^{-1})^{(\sigma)} - \psi_i^{(\sigma)} \right], \quad (3.7)$$

where $\gamma_i^{(\sigma)} = d_i^{(\sigma)} / b_i^{(\sigma)}$. This formalism requires the matrix $\mathbb{M}$ to be invertible. However, it will be shown that there is no anomalous behaviour when $\det(\mathbb{M}) = 0$.

To keep the notation compact we define the difference variable $v_i^{(\sigma)} = \psi_i^{(\sigma)} - \psi_{i-1}^{(\sigma)}$. Using $\psi_0^{(\sigma)} = 0$, we have $\sum_{j=1}^i v_j^{(\sigma)} = \psi_i^{(\sigma)}$. With this notation we can write Eq. (3.7) in the following form

$$v_{i+1}^{(\sigma)} = \gamma_i^{(\sigma)} v_i^{(\sigma)} + \frac{1}{b_i^{(\sigma)}} \left[\sum_{j=1}^i \mathbf{v}_j \cdot (\mathbb{M}^{-1} - \mathbb{I}) \right]^{(\sigma)}, \quad (3.8)$$

where $\mathbb{I}$ is the $\Omega \times \Omega$ identity matrix. This relation expresses the vector $\mathbf{v}_{i+1}$ in terms of the vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_i$. We can therefore express all vectors $\mathbf{v}_i$ ($i = 2, \dots, N$) in terms of $\mathbf{v}_1$. The constraint $\sum_{i=1}^N \mathbf{v}_i = \boldsymbol{\psi}_N = (1, \dots, 1)$ then determines $\mathbf{v}_1$ self-consistently. We note that the resulting set of equations is linear in the set $\{v_1^{(\sigma)}\}$. Hence a solution can be obtained in closed form, in principle. In practice one inverts the linear system using one of the standard algebraic manipulation packages. Once $\mathbf{v}_1$ has been found, the other components $\mathbf{v}_i$, with $i = 2, \dots, N$, can be computed via Eq. (3.8). One then uses $\boldsymbol{\phi}_i = \sum_{j=1}^i \mathbf{v}_j \cdot \mathbb{M}^{-1}$ to find the fixation probabilities starting with i individuals of type A in environment σ , $\phi_i^{(\sigma)}$.

We note here that algebraically inverting the linear system (3.8) when N is large is difficult due to the increasing number of terms in the corresponding expressions. Thus, at present, this theory is limited computationally to relatively small system-sizes. We have shown it is accurate up to $N = \mathcal{O}(100)$.

In the case of a single environment, $\Omega = 1$, the matrix $\mathbb{M}$ is simply the 1×1 identity matrix, and the solution to Eq. (3.8) simplifies to the well-known result for birth–death processes, Eq. (2.24) [12–15].

¹As the backward equation is the adjoint of the forward equation, the column vectors of probability are transposed to give row vectors.

Mean Unconditional Fixation Time

We write $t_i^{(\sigma)}$ for the expected number of time-steps taken to reach any one of the two absorbing states, given that the system is started in state (i, σ) . Just like we did in Eq. (2.25), we define $\vartheta_i^{(\sigma)}(t) = \sum_{\sigma''} [\mathcal{Q}_{(0,\sigma''); (i,\sigma)}(t) + \mathcal{Q}_{(N,\sigma''); (i,\sigma)}(t)]$ and

$$t_i^{(\sigma)} = \sum_{t=0}^{\infty} t \left[\vartheta_i^{(\sigma)}(t) - \vartheta_i^{(\sigma)}(t-1) \right]. \quad (3.9)$$

The mean fixation times fulfill the boundary conditions $t_0^{(\sigma)} = t_N^{(\sigma)} = 0$. Proceeding as described in Sect. 2.4, from the backward master equation (3.4) we can express the mean unconditional fixation time as

$$t_i^{(\sigma)} = \sum_{\sigma' \in \Lambda} \mu_{\sigma', \sigma} \left[b_i^{(\sigma)} t_{i+1}^{(\sigma')} + d_i^{(\sigma)} t_{i-1}^{(\sigma')} + \left(1 - b_i^{(\sigma)} - d_i^{(\sigma)}\right) t_i^{(\sigma')} \right] + 1. \quad (3.10)$$

Introducing the variable $\xi_i^{(\sigma)} = \sum_{\sigma'} \mu_{\sigma', \sigma} t_i^{(\sigma')}$, we have

$$t_i^{(\sigma)} = b_i^{(\sigma)} \left(\xi_{i+1}^{(\sigma)} - \xi_i^{(\sigma)} \right) - d_i^{(\sigma)} \left(\xi_i^{(\sigma)} - \xi_{i-1}^{(\sigma)} \right) + \xi_i^{(\sigma)} + 1. \quad (3.11)$$

With the notation $v_i^{(\sigma)} = \xi_i^{(\sigma)} - \xi_{i-1}^{(\sigma)}$, and using $\sum_{j=1}^i v_j^{(\sigma)} = \xi_i^{(\sigma)}$, we arrive at

$$v_{i+1}^{(\sigma)} = \gamma_i^{(\sigma)} v_i^{(\sigma)} + \frac{1}{b_i^{(\sigma)}} \left[\sum_{j=1}^i \mathbf{v}_j \cdot (\mathbb{M}^{-1} - \mathbb{I}) \right]^{(\sigma)} - \frac{1}{b_i^{(\sigma)}}. \quad (3.12)$$

This relation allows one to express all vectors $\mathbf{v}_i$ ($i = 2, \dots, N$) in terms of $\mathbf{v}_1$. The constraint $\sum_{i=1}^N \mathbf{v}_i = (0, \dots, 0)$ from the boundary condition then determines $\mathbf{v}_1$, and the mean unconditional fixation times are computed using $\mathbf{t}_i = \sum_{j=1}^i \mathbf{v}_j \cdot \mathbb{M}^{-1}$. As described previously, the $\mathbf{t}_i$ are expressed in units of elementary time-steps. In the results presented below, times are expressed in generations, i.e. we plot t/N on the time axes.

Mean Conditional Fixation Time

We write $t_{i|A}^{(\sigma)}$ for the mean fixation time conditioned on absorption in the all-A state, given that the system is initially in state (i, σ) . To find this conditional fixation time, we proceed along similar lines as per the derivation of Eq. (2.39). We introduce the variable $\varphi_i^{(\sigma)}(t) = \sum_{\sigma''} \mathcal{Q}_{(N,\sigma''); (i,\sigma)}(t)$, and define the mean conditional fixation time as

$$t_{i|A}^{(\sigma)} = \frac{1}{\phi_i^{(\sigma)}} \sum_{t=0}^{\infty} t \left[\varphi_i^{(\sigma)}(t) - \varphi_i^{(\sigma)}(t-1) \right]. \quad (3.13)$$

From the backward master equation (3.4), it can be seen that $t_{i|A}^{(\sigma)}$ satisfies

$$\phi_i^{(\sigma)} t_{i|A}^{(\sigma)} = \sum_{\sigma' \in \Lambda} \mu_{\sigma', \sigma} \left[b_i^{(\sigma)} \phi_{i+1}^{(\sigma')} t_{i+1|A}^{(\sigma')} + d_i^{(\sigma)} \phi_{i-1}^{(\sigma')} t_{i-1|A}^{(\sigma')} + \left(1 - b_i^{(\sigma)} - d_i^{(\sigma)} \right) \phi_i^{(\sigma')} t_i^{(\sigma')} \right] + \phi_i^{(\sigma)}. \quad (3.14)$$

Introducing the variable $\theta_i^{(\sigma)} = \phi_i^{(\sigma)} t_{i|A}^{(\sigma)}$, which has boundary conditions $\theta_0^{(\sigma)} = \theta_N^{(\sigma)} = 0$, and $\zeta_i^{(\sigma)} = \sum_{\sigma'} \mu_{\sigma', \sigma} \theta_i^{(\sigma')}$, we have

$$\theta_i^{(\sigma)} = b_i^{(\sigma)} \left(\zeta_{i+1}^{(\sigma)} - \zeta_i^{(\sigma)} \right) - d_i^{(\sigma)} \left(\zeta_i^{(\sigma)} - \zeta_{i-1}^{(\sigma)} \right) + \zeta_i^{(\sigma)} + \phi_i^{(\sigma)}. \quad (3.15)$$

Now introducing the difference variable $\eta_i^{(\sigma)} = \zeta_i^{(\sigma)} - \zeta_{i-1}^{(\sigma)}$ and noting that $\sum_{j=1}^i \eta_j^{(\sigma)} = \zeta_i^{(\sigma)}$, we arrive at

$$\eta_{i+1}^{(\sigma)} = \gamma_i^{(\sigma)} \eta_i^{(\sigma)} + \frac{1}{b_i^{(\sigma)}} \left[\sum_{j=1}^i \eta_j \cdot (\mathbb{M}^{-1} - \mathbb{I}) \right]^{(\sigma)} - \frac{1}{b_i^{(\sigma)}} \phi_i^{(\sigma)}. \quad (3.16)$$

The set $\{\theta_i^{(\sigma)}\}$ can then be found using $\theta_i = \sum_{j=1}^i \eta_j \cdot \mathbb{M}^{-1}$. Results for the mean conditional fixation time can then be obtained using $t_{i|A}^{(\sigma)} = \theta_i^{(\sigma)} / \phi_i^{(\sigma)}$.

Switching Between Two Environments

We now focus on the case of environments which can be in one of two possible states, i.e. $\Omega = 2$. We label the two states as $\sigma = \pm 1$ ($\Lambda = \{+1, -1\}$). The matrix $\mathbb{M}$ can then be written as

$$\mathbb{M} = \begin{pmatrix} 1 - p_+ & p_- \\ p_+ & 1 - p_- \end{pmatrix}, \quad (3.17)$$

where the quantity p_σ is the probability that in a given time-step the environment switches *from* state σ to $-\sigma$, i.e. if the environment is in state $+1$, it will switch to state -1 in the next time-step with probability p_+ .

We recall that our theoretical results require the inversion of $\mathbb{M}$. Excluding the case when $\Delta = \det(\mathbb{M}) = 1 - p_+ - p_-$ vanishes, this inversion can be carried out straightforwardly,

$$\mathbb{M}^{-1} = \frac{1}{\Delta} \begin{pmatrix} 1 - p_- & -p_- \\ -p_+ & 1 - p_+ \end{pmatrix}. \quad (3.18)$$

For the case $\Delta = 0$, i.e. $p_+ + p_- = 1$, we have verified that there is no anomalous behaviour of simulation results, see below (Fig. 3.6).

To find the fixation probability in this two-environment system, we substitute the inverse matrix (3.18) into the general result (3.8). This reduces to the recursion

$$v_{i+1}^{(\sigma)} = \gamma_i^{(\sigma)} v_i^{(\sigma)} + \frac{1}{b_i^{(\sigma)}} \frac{p_\sigma}{\Delta} \sum_{j=1}^i \left(v_j^{(\sigma)} - v_j^{(-\sigma)} \right), \quad (3.19)$$

along with the boundary condition $\sum_{j=1}^N \mathbf{v}_j = (1, 1)$. The fixation probability is obtained via

$$\phi_i^{(\sigma)} = \left[\sum_{j=1}^i \mathbf{v}_j \cdot \mathbb{M}^{-1} \right]^{(\sigma)} = \frac{1}{\Delta} \sum_{j=1}^i \left[(1 - p_{-\sigma}) v_j^{(\sigma)} - p_{\sigma} v_j^{(-\sigma)} \right]. \quad (3.20)$$

Similarly, Eq. (3.12) for the mean unconditional fixation time reduces to

$$v_{i+1}^{(\sigma)} = \gamma_i^{(\sigma)} v_i^{(\sigma)} + \frac{1}{b_i^{(\sigma)}} \frac{p_{\sigma}}{\Delta} \sum_{j=1}^i \left(v_j^{(\sigma)} - v_j^{(-\sigma)} \right) - \frac{1}{b_i^{(\sigma)}}, \quad (3.21)$$

along with the boundary condition $\sum_{i=1}^N \mathbf{v}_j = (0, 0)$. The mean unconditional fixation times are then found as

$$t_i^{(\sigma)} = \left[\sum_{j=1}^i \mathbf{v}_j \cdot \mathbb{M}^{-1} \right]^{(\sigma)} = \frac{1}{\Delta} \sum_{j=1}^i \left[(1 - p_{-\sigma}) v_j^{(\sigma)} - p_{\sigma} v_j^{(-\sigma)} \right]. \quad (3.22)$$

Finally, Eq. (3.16) for the mean conditional fixation time reduces to

$$\eta_{i+1}^{(\sigma)} = \gamma_i^{(\sigma)} \eta_i^{(\sigma)} + \frac{1}{b_i^{(\sigma)}} \frac{p_{\sigma}}{\Delta} \sum_{j=1}^i \left(\eta_j^{(\sigma)} - \eta_j^{(-\sigma)} \right) - \frac{1}{b_i^{(\sigma)}} \phi_i^{(\sigma)}, \quad (3.23)$$

along with the boundary condition $\sum_{i=1}^N \eta_j = (0, 0)$. The mean conditional fixation times are then found as

$$t_{i|A}^{(\sigma)} = \frac{1}{\phi_i^{(\sigma)}} \left[\sum_{j=1}^i \eta_j \cdot \mathbb{M}^{-1} \right]^{(\sigma)} = \frac{1}{\phi_i^{(\sigma)}} \frac{1}{\Delta} \sum_{j=1}^i \left[(1 - p_{-\sigma}) \eta_j^{(\sigma)} - p_{\sigma} \eta_j^{(-\sigma)} \right]. \quad (3.24)$$

Effective Description for Fast Switching

We say the environment is in the ‘fast switching’ regime if the lifetime of the environmental states are much shorter than the mean fixation time in either environment, i.e. in the fast switching regime we expect the state of the environment to change many times before fixation is reached. If this is the case we expect the population dynamics to be controlled by a set of *effective* transition probabilities, which are weighted averages of the original transition probabilities in the different environmental states. The weights are given by the fraction of time spent in each environmental state.

In the two-environment scenario, the dynamics of σ follow a so-called telegraph process [16]. Writing $P_{\sigma}(t)$ for the probability for the environment to be found in state σ at time t , the evolution of σ is described by the forward master equation

$P_\sigma(t+1) = \sum_{\sigma'} \mu_{\sigma,\sigma'} P_{\sigma'}(t)$. In matrix form this is simply $\mathbf{P}(t+1) = \mathbb{M} \cdot \mathbf{P}(t)$. By diagonalising the 2×2 matrix $\mathbb{M}$ (Eq. (3.17)), we can express $\mathbf{P}(t)$ in terms of the eigenvectors ($\mathbf{v}_1$ and $\mathbf{v}_2$) and eigenvalues ($\lambda_1 = 1$ and $\lambda_2 = 1 - p_+ - p_-$) of $\mathbb{M}$, which gives

$$\mathbf{P}(t) = c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2 \lambda_2^t, \quad (3.25)$$

where c_1 and c_2 are determined by the initial condition. Taking the long-time limit of this expression reveals the fraction of time spent in the state σ is $p_{-\sigma}/(p_\sigma + p_{-\sigma})$ for $\sigma \in \{-1, +1\}$. Hence the time spent in state σ decreases with increasing p_σ if $p_{-\sigma}$ is held fixed. Using this long-time limit, the effective transition probabilities are given by

$$b_i^{\text{eff}} = \frac{p_-}{p_+ + p_-} b_i^{(+)} + \frac{p_+}{p_+ + p_-} b_i^{(-)}, \quad (3.26)$$

and likewise for d_i^{eff} .

In this approximation the dynamics of the population are mapped onto the familiar birth–death process on the set $i \in \{0, 1, \dots, N\}$ with absorbing states $i = 0$ and $i = N$, as described in Sect. 2.3. Using the expressions from Sect. 2.4 for the fixation probabilities and mean fixation times, we can describe the fixation properties of the population in the fast-switching regime. For the fixation probability of a single mutant, we have

$$\phi_1^{\text{eff}} = \frac{1}{1 + \sum_{k=1}^{N-1} \prod_{j=1}^k \gamma_j^{\text{eff}}}. \quad (3.27)$$

We have here written $\gamma_i^{\text{eff}} = d_i^{\text{eff}}/b_i^{\text{eff}}$. The corresponding approximations for the mean unconditional and conditional fixation times of a single mutant are

$$t_1^{\text{eff}} = \phi_1^{\text{eff}} \sum_{k=1}^{N-1} \sum_{\ell=1}^k \frac{1}{b_\ell^{\text{eff}}} \prod_{m=\ell+1}^k \gamma_m^{\text{eff}}, \quad (3.28)$$

$$t_{1|A}^{\text{eff}} = \sum_{k=1}^{N-1} \sum_{\ell=1}^k \frac{1}{b_\ell^{\text{eff}}} \phi_\ell^{\text{eff}} \prod_{m=\ell+1}^k \gamma_m^{\text{eff}}, \quad (3.29)$$

respectively. These expressions describe exactly the fixation properties of a birth–death system with the effective transition probabilities; the nature of our approximation is to assume that the birth–death process in quickly changing environments can be described by the effective transition probabilities in Eq. (3.26).

Finally we note that this theory is *independent* of the invertibility of the switching matrix $\mathbb{M}$ (also shown in Fig. 3.6).

3.4 Switching Between Two Games

As a direct application of the general theory we have developed, we now consider evolutionary game dynamics in well-mixed, finite populations. We modify the dynamics described in Sect. 2.8 so that at any point in time the environment is in one of two discrete states ($\sigma \in \{+1, -1\}$). In this model the state of the environment affects the payoff structure. The interaction between individuals is characterised by the payoff matrix

$$\begin{array}{c|cc} & A & B \\ \hline A & R^{(\sigma)} & S^{(\sigma)} \\ \hline B & T^{(\sigma)} & P^{(\sigma)} \end{array} \quad (3.30)$$

If the environment is in state σ , and if there are i individuals of type A in the population and $N - i$ individuals of type B , the expected payoffs for each type of player are

$$\pi_A^{(\sigma)}(i) = \frac{i-1}{N-1} R^{(\sigma)} + \frac{N-i}{N-1} S^{(\sigma)}, \quad (3.31a)$$

$$\pi_B^{(\sigma)}(i) = \frac{i}{N-1} T^{(\sigma)} + \frac{N-i-1}{N-1} P^{(\sigma)}, \quad (3.31b)$$

as described in Eq. (2.74). For this example we choose an exponential mapping between expected payoff and fitness as described in Eq. (2.78). In environment σ the fitnesses are

$$f_A^{(\sigma)}(i) = \exp \left[\beta \pi_A^{(\sigma)}(i) \right], \quad (3.32a)$$

$$f_B^{(\sigma)}(i) = \exp \left[\beta \pi_B^{(\sigma)}(i) \right], \quad (3.32b)$$

where $\beta > 0$ is the so-called intensity of selection. Following the update rules of the Moran process as described in Eq. (2.79), the intensive transition probabilities are given by

$$b_i^{(\sigma)} = \frac{i(N-i)}{N^2} \frac{f_A^{(\sigma)}(i)}{\bar{f}^{(\sigma)}(i)} \quad (3.33a)$$

$$d_i^{(\sigma)} = \frac{i(N-i)}{N^2} \frac{f_B^{(\sigma)}(i)}{\bar{f}^{(\sigma)}(i)}, \quad (3.33b)$$

where $\bar{f}^{(\sigma)}(i) = [i f_A^{(\sigma)}(i) + (N-i) f_B^{(\sigma)}(i)]/N$ is the average fitness of the population.

To reduce the parameter space of our model, we consider $R^{(\sigma)} = P^{(\sigma)} = 1$ in the payoff matrix (3.30). The type of game is then determined by the off-diagonal terms. We choose $S^{(\sigma)} = 1 + \sigma q$ and $T^{(\sigma)} = 1 + \sigma r$, where q and r are real-valued

parameters. Thus we have the parametrised payoff matrix

$$\begin{array}{c|cc} & A & B \\ \hline A & 1 & 1 + \sigma q \\ B & 1 + \sigma r & 1. \end{array} \quad (3.34)$$

This parametrisation does not span the entire space of all 2×2 games, but it covers the three general types discussed in Sect. 2.8:

Dominance: $1 + \sigma q > 1$ and $1 + \sigma r < 1$ (or $1 + \sigma q < 1$ and $1 + \sigma r > 1$), type A (or type B) always has the higher fitness irrespective of the composition of the population. This type is then always favoured by selection.

Coexistence: $1 + \sigma q > 1$ and $1 + \sigma r > 1$, selection drives the population away from the absorbing boundaries.

Coordination: $1 + \sigma q < 1$ and $1 + \sigma r < 1$, the population exhibits bi-stability and selection drives the population towards the monomorphic states.

In the last two cases there exists an internal point in frequency space for which the direction of selection changes its sign, i.e. at which the gradient of selection, $b_i^{(\sigma)} - d_i^{(\sigma)}$, is zero. This is referred to as the selection-balance point. This point can be calculated by solving $b_{i^*}^{(\sigma)} = d_{i^*}^{(\sigma)}$ for i^* . From the transition probabilities and fitnesses in Eqs. (3.32) and (3.33), we see i^* satisfies $\pi_A^{(\sigma)}(i^*) = \pi_B^{(\sigma)}(i^*)$. From Eq. (3.31) we find $i^*/N = q/(q+r)$ for both $\sigma = \pm 1$. In the dominance game there are no such turning points. For the remainder of this chapter we focus on switching between coexistence and coordination games. More precisely we choose $q > 0$ and $r > 0$ in Eq. (3.34). The coexistence game corresponds to $\sigma = +1$ and the coordination game to $\sigma = -1$.

Results

In Fig. 3.3a we show a sample trajectory of a simulation in which a single mutant reaches fixation. The gradient of selection, $b_i^{(\sigma)} - d_i^{(\sigma)}$, for the two fixed environments is shown in Fig. 3.3b, c. During periods when the environment is in the coexistence state (light background; $\sigma = +1$) the population fluctuates about the selection-balance point (dashed line). During periods when the environment is in the coordination state (shaded background; $\sigma = -1$) the population is driven away from the selection-balance point. In the final period in the coordination state the mutant is driven to fixation.

In Fig. 3.4 we show the variation of fixation probability and mean conditional fixation time with the switching parameters p_+ and p_- obtained from our theoretical framework. The fixation probability in this example depends non-trivially on the environmental switching parameters; we find a combination, $p_+ \simeq p_-$, for which fixation of a single mutant is most likely, as shown in Fig. 3.4a, b. The initial state of the environment has very little effect on the fixation probability for $p_{\pm} \gtrsim 0.1$. In this region the switching process of the environment is too fast for the initial condition to have any significant effect on the population dynamics, and it is here that we expect the effective description to approximate the system well. For $p_{\pm} \lesssim$




Fig. 3.3 **a** A sample trajectory (time series) of the fraction of individuals of type A. White background corresponds to the environment being in the $\sigma = +1$ coexistence state, while the shaded background corresponds to the $\sigma = -1$ coordination state. Dashed line is the location of the point at which selection balances, which is the same in both states of the environment. **b** Gradient of selection in the $\sigma = +1$ coexistence state, $b_i^{(+)} - d_i^{(+)}$. Solid circle shows location of the point of selection balance, and arrows indicate the direction and magnitude of flow towards this point. **c** Gradient of selection in the $\sigma = -1$ coordination state, $b_i^{(-)} - d_i^{(-)}$. Empty circle shows location of the point of selection balance, and arrows indicate the direction and magnitude of flow away from this point. For the realisation in panel **a** and the selection bias shown in **b** and **c**, the payoff matrix parameters are $q = 0.5$ and $r = 0.9$, the selection intensity is $\beta = 1$, the system size is $N = 100$, and the switching probabilities are $p_+ = 10^{-3}$ and $p_- = 10^{-4}$. Time is measured in generations

0.1 the fixation probability is affected by the initial condition. This effect can be understood by considering the deterministic gradient of selection of the two games, which are qualitatively the same as those shown in Fig. 3.3b, c.² When starting in the coordination game, selection pushes the mutant towards extinction. Hence fixation is more likely if the initial state is the coexistence game.

The mean conditional fixation times shown in Fig. 3.4c, d show very little dependence on the initial state of the environment. Systems started in the coordination environment will tend to reach extinction relatively quickly due to initial adverse selection, unless the environment switches to the coexistence state early on. Thus the sample of runs that reach fixation started in the coordination game will be dominated by runs in which the environment switches soon after the start of the run. Then we expect that the value of the mean conditional fixation time is close to the one obtained when starting in the coexistence game.

The fixation time is small for $p_+ \gg p_-$ when the environment is found mostly in the coordination game, and large when the environment is mostly in the coexistence state ($p_+ \ll p_-$). If fixation happens, it will generally be quicker in the coordination game than in the coexistence game [17, 18]. This is due to the adverse selection bias in the coordination game at low mutant numbers, as shown in Fig. 3.3c. The more time the system spends in this region of adverse selection the less likely it is for the mutant to reach fixation. Thus if fixation happens in a coordination game then it happens fast. In the coexistence game on the other hand the direction of selection is towards the balance point, as shown in Fig. 3.3b. The system can ‘afford’ to spend significant

²Payoff parameters are the same in Figs. 3.3 and 3.4. Although system-size and selection intensity are different, these parameters do not affect the qualitative features of the gradient of selection.




Fig. 3.4 Theoretical predictions of the fixation probability and mean conditional fixation time of a single mutant as a function of the switching parameters p_+ and p_- . Panel **a** shows the fixation probability (Eq. 3.20) when starting from the $\sigma = +1$ coexistence game, and panel **b** shows the same when starting from the $\sigma = -1$ coordination game. Panel **c** shows the mean conditional fixation time (Eq. 3.24) (time measured in generations) when starting from the $\sigma = +1$ coexistence game, and panel **d** shows the same when starting from the $\sigma = -1$ coordination game. Horizontal lines correspond to the data shown in Fig. 3.5. The payoff matrix parameters are $q = 0.5$ and $r = 0.9$, the selection intensity is $\beta = 0.5$, and the system size is $N = 50$

time in the region of small mutant numbers and still reach fixation eventually even after repeated excursions through frequency space. Thus there is no need for fixation to occur quickly, and conditional fixation times can be long. These observations make it plausible that the mean conditional fixation time will generally decrease when less time is spent in the coexistence game, which is exactly what we find in Fig. 3.4c, d. Other choices of the parameters q and r for which the two games are a coexistence game and a coordination game, reveal that the behaviour of the mean conditional fixation times is robust under such changes.

To compare how the analytic predictions compare with simulation results we take a cross-section of the data, shown by the horizontal lines in Fig. 3.4a–d. This data is shown in Fig. 3.5. The theoretical predictions of Eqs. (3.20) and (3.24), indicated by solid lines, are in convincing agreement with simulation data (symbols). The predictions of the effective theory, Eqs. (3.27) and (3.29), agree well with simulation results in the fast-switching region. The effective theory qualitatively matches the data, but unsurprisingly there are systematic deviations when switching is slow.




Fig. 3.5 **a** Fixation probability of a single mutant at fixed p_- computed from simulation results (symbols; crosses correspond to $\sigma(0) = +1$ and circles to $\sigma(0) = -1$), along with the theoretical results (solid lines) from Eq. (3.20), and the fast-switching approximation result (dashed line) of Eq. (3.27). **b** Mean conditional fixation time (in generations) of a single mutant at fixed p_- computed from simulation results as described above, along with the exact theoretical results (solid lines) of Eq. (3.24), and the fast-switching approximation result (dashed line) of Eq. (3.29). The parameters are as in Fig. 3.4, and $p_- = 0.01$

In Fig. 3.6, to show that there is no anomalous behaviour when the switching matrix $\mathbb{M}$ is singular, we keep $p_- = 0.5$ fixed and sweep p_+ across 0.5, which is the point at which $\Delta = 0$. The simulation data does not indicate any singularity or anomalous behaviour, and this is also confirmed by the effective theory which does not require $\mathbb{M}$ to be invertible. The apparent issue is hence not a fundamental problem, but merely an artificial singularity that arises from the procedure that we use to analyse this system.




Fig. 3.6 **a** Fixation probability of a single mutant at fixed $p_- = 0.5$ computed from simulation results (symbols; crosses correspond to $\sigma(0) = +1$ and circles to $\sigma(0) = -1$), along with the theoretical results (solid lines) from Eq. (3.20), and the fast-switching approximation result (dashed line) of Eq. (3.27). **b** Mean conditional fixation time (in generations) of a single mutant at fixed $p_- = 0.5$ computed from simulation results as described above, along with the exact theoretical results (solid lines) of Eq. (3.24), and the fast-switching approximation result (dashed line) of Eq. (3.29). The parameters are as in Fig. 3.4. Simulation data and the effective theory reveal no singular or anomalous behaviour at $p_+ = 0.5$




Fig. 3.7 **a** Fixation probability of a single mutant at fixed p_- started in the $\sigma(0) = +1$ coexistence game for various system sizes. Simulation results (symbols) are shown along with the theoretical results (solid lines) from Eq. (3.20). **b** Mean conditional fixation time (in generations) of a single mutant. The main panel shows simulation results, as described above along with the exact theoretical results (solid lines) Eq. (3.24). Inset panel shows the theoretical results scaled by the system size. Open symbols are averages over 10^4 simulation runs, filled symbols correspond to averaging over 10^6 runs. The parameters are as in Fig. 3.4, and $p_- = 0.01$

The features observed in Fig. 3.5, i.e. the peak in the fixation probability and shape of the mean conditional fixation time as a function of p_+ , are found to be robust when the system size is increased, as shown in Fig. 3.7. Fixation probabilities generally decrease with system size, but the observed peak becomes sharper. This effect is highlighted by the logarithmic axis in Fig. 3.7a.

The mean conditional fixation time scales exponentially in N when the system spends most of its time in the coexistence state. This scaling is in-line with existing results [17, 19]. If the majority of time is spent in the coordination game, then the mean conditional fixation time increases sub-linearly with N , again in agreement with existing literature [17]. As can be seen in the inset of Fig. 3.7b, there is a ‘critical’ switching rate $p_+ \simeq p_- = 0.01$ at which the scaling of the fixation time (in generations) is linear in N . This is in agreement with the neutral theory result [17].

Finally we comment on how varying the payoff matrix parameters affects the evolutionary outcome. By considering the deterministic gradient of selection, we have shown that the selection-balance point is located at $i^*/N = q/(q + r)$ in both the coexistence and coordination games. If the environment is fixed to the coexistence-game state, fixation of the mutant is more likely if the selection-balance point is close to the fixated state as there is a greater region in which selection favours the mutant. This is shown in Fig. 3.8a. In a fixed coordination-game environment the reverse is the case. The range of adverse selection is to the left of the balance point, and so fixation is less likely the closer the point of selection balance is to the fixated state, again shown in Fig. 3.8a. These views are validated through computation of the fixation probability in the fixed games using Eq. (2.24) as a function of the fixed point location, shown in Fig. 3.8b.

Using these ideas we can infer how the system with a switching environment will behave. For $q \ll r$, i.e. a selection-balance point close to $i = 0$, we expect that the fixation probability will increase the more time is spent in the coordination-game




Fig. 3.8 **a** Illustration of selection bias in the two environments for different locations of the balance point; **b** Fixation probability of a single mutant in the fixed coexistence and coordination games calculated using Eq. (2.24). Here $q + r = 1$ such that the selection-balance point is located at $i^*/N = q$. The dotted line is the neutral theory result $\phi_1 = 1/N$, found by setting $\gamma_i = d_i/b_i = 1$ in Eq. (2.4). Remaining parameters are $\beta = 0.5$ and $N = 50$

environment, i.e. $\phi_1^{(\sigma)}$ is an increasing function of the probability p_+ with which the system leaves the $\sigma = +1$ state (coexistence game). For $q \gg r$, i.e. i^* close to $i = N$, the reverse is the case. Fixation is more likely in the coexistence game ($\sigma = +1$), and the fixation probability is hence a decreasing function of p_+ at fixed p_- .

For $q \approx r$ the situation is less clear. The fixation probability will be comparable in both games if the environment is frozen, as shown in Fig. 3.8b. Two effects here conspire to produce a non-trivial outcome:

- (i) Consider the case in which the system is mostly in the coordination-game state, i.e. $p_+ \gg p_-$. It is plausible that an occasional switch to a coexistence game will make fixation more likely than in a constant coordination game. This is because the coexistence-game environment pushes the system away from extinction at low mutant numbers. In the regime of $p_+ \gg p_-$ we thus expect the fixation probability to increase as p_+ is lowered. In other words, $\phi_1^{(\sigma)}$ is a decreasing function at large p_+ .
- (ii) Similarly, if the system is mostly in the coexistence-game environment ($p_+ \ll p_-$), short periods of time in the coordination game can make fixation more likely. This is because selection at large mutant numbers is directed towards fixation in the coordination game. At $p_+ \ll p_-$ we expect $\phi_1^{(\sigma)}$ to be an increasing function of p_+ .

These two effects taken together generate a maximum of the fixation probability at intermediate values of $p_+ \approx p_-$, which is exactly what we find in Fig. 3.4.

To confirm our picture we find the value of p_+ that maximises fixation probability as a function of q and r in Fig. 3.9 (for a given $p_- = 0.01$). The point of selection balance is $1/(1+r/q)$. The presence of diagonal structures in Fig. 3.9 shows that the behaviour of the fixation probability is determined by the location of the selection-balance point. If this point is close to the fixation state $i = N$ ($q \gg r$, bottom-right




Fig. 3.9 The value of p_+ at which $\phi_1^{(+)}$ is maximal given $p_- = 0.01$ as a function of q and r . Fixation probabilities found using Eq. (3.20). Diagonal structures indicate only the ratio between r and q , i.e. the location of the selection-balance point, determines the qualitative behaviour. Remaining parameters are $\beta = 0.5$ and $N = 50$

in Fig. 3.9), then the fixation probability is maximal for vanishing p_+ . If this point is close to the extinction state ($q \ll r$, top-left in Fig. 3.9), then the fixation probability is maximal for large p_+ . For intermediate locations of the selection-balance point ($q \approx r$) fixation is maximised at a non-trivial combination of environment states. Starting the environment in the $\sigma = -1$ coordination game produces an almost identical picture to Fig. 3.9.

3.5 Mutation–Selection Equilibria

We now consider systems with mutations occurring during the dynamics. This removes the possibility of fixation and extinction. The combination of mutation, selection, and noise can lead to non-trivial stationary states. We introduce mutation by modifying the discrete-time transition probabilities of Eq. (3.33) and now use

$$b_i^{(\sigma)} = (1 - u) \frac{i(N - i)}{N^2} \frac{f_A^{(\sigma)}(i)}{\bar{f}^{(\sigma)}(i)} + u \frac{(N - i)^2}{N^2}, \quad (3.35a)$$

$$d_i^{(\sigma)} = (1 - u) \frac{i(N - i)}{N^2} \frac{f_B^{(\sigma)}(i)}{\bar{f}^{(\sigma)}(i)} + u \frac{i^2}{N^2}, \quad (3.35b)$$

where $u \ll 1$ is the mutation rate. The transition probabilities $b_0^{(\sigma)} = d_N^{(\sigma)} = u$ are now non-zero, and so the states $i = 0$ and $i = N$ are no longer absorbing.

The stationary probability $\rho_{(i,\sigma)}$ of finding the system in state (i, σ) ($i = 0, 1, \dots, N$; $\sigma \in \Lambda$) is obtained, as described in Sect. 2.5, by taking the infinite-time limit of the forward master equation (3.3). This gives

$$\rho_{(i,\sigma)} = \sum_{\sigma' \in \Lambda} \mu_{\sigma,\sigma'} \left[b_{i-1}^{(\sigma')} \rho_{(i-1,\sigma')} + d_{i+1}^{(\sigma')} \rho_{(i+1,\sigma')} + \left(1 - b_i^{(\sigma')} - d_i^{(\sigma')}\right) \rho_{(i,\sigma')} \right]. \quad (3.36)$$

This equation is of the form $\rho_{(i,\sigma)} = \sum_{\sigma'} \sum_j R_{(i,\sigma),(j,\sigma')} \rho_{(j,\sigma')}$, and it is solved by finding the eigenvector of the linear operator R corresponding to the eigenvalue $\lambda = 1$. The stationary distribution for the state of the population is found by summing over all states of the environment, $\rho_i = \sum_{\sigma} \rho_{(i,\sigma)}$. This solution is the exact stationary distribution of the population.

If the switching probabilities are large, we are in the fast-switching limit described by Eq. (3.26). In this regime one might expect the stationary distribution of the population to be approximated by the distribution of a system controlled by the effective transition rates, b_i^{eff} and d_i^{eff} . The resulting effective stationary distribution of the population, ρ_i^{eff} , is given by Eq. (2.54), i.e.

$$\rho_i^{\text{eff}} = \frac{\Gamma_i^{\text{eff}}}{\sum_{j=0}^N \Gamma_j^{\text{eff}}}, \quad \Gamma_i^{\text{eff}} = \prod_{j=1}^i \frac{b_{j-1}^{\text{eff}}}{d_j^{\text{eff}}}. \quad (3.37)$$

If the switching probabilities are small, then the environment states are long-lived. In this regime the population will relax to the stationary state of the current environment before the next switching event. With this, one might expect that the stationary distribution of the population is given by the weighted average of the stationary distributions one would obtain in the respective single environments. The stationary distribution in a single fixed environment, $\rho_i^{(\sigma)}$, can again be read off from Eq. (2.54) as

$$\rho_i^{(\sigma)} = \frac{\Gamma_i^{(\sigma)}}{\sum_{j=0}^N \Gamma_j^{(\sigma)}}, \quad \Gamma_i^{(\sigma)} = \prod_{j=1}^i \frac{b_{j-1}^{(\sigma)}}{d_j^{(\sigma)}}. \quad (3.38)$$

This can also be derived from Eq. (3.36) by assuming that the transition matrix of the environment, $\mathbb{M}$, is the identity matrix. In fact $\lim_{p_{\sigma} \rightarrow 0} \mathbb{M} = \mathbb{I}$. The average stationary distribution over many slow-switching environments can then be written as

$$\bar{\rho}_i = \sum_{\sigma' \in \Lambda} \rho_{\sigma'} \rho_i^{(\sigma')}, \quad (3.39)$$

where ρ_{σ} is the (stationary) probability that the environment is in state σ .

Results

For the environmental dynamics we will again consider the two-state scenario discussed in Sect. 3.4, where the environment switches between a coexistence game and a coordination game. Figure 3.10a shows the stationary distributions of the population




Fig. 3.10 **a** The stationary distributions in the single-environment coexistence game $\rho_{i,+}^{(1)}$ (dotted line) and coordination game $\rho_{i,-}^{(1)}$ (dashed lines) calculated from Eq. (3.38), along with the ‘average’ $\bar{\rho}_i$ (Eq. 3.39) and effective $\rho_{i,\text{eff}}$ (Eq. 3.37) approximate stationary solutions (solid lines) for $p_+ = p_-$. **b** The exact solution ρ_i (Eq. (3.36); solid lines and symbols for identification) at different combinations of magnitudes of $p_+ = p_-$. The payoff matrix parameters are $q = 0.5$ and $r = 0.9$, the system size is $N = 50$, the selection intensity is $\beta = 0.5$, and the mutation probability is $u = 0.02$

for the fixed environments (calculated using Eq. (3.38)), and the approximate stationary distributions ρ_i^{eff} and $\bar{\rho}_i$ for identical switching parameters. In a constant coexistence game ($\sigma = +1$) the stationary distribution is peaked about the point at which the gradient of selection changes sign, and in a fixed coordination game ($\sigma = -1$) we find a distribution which is strongly peaked about the $i \simeq N$ state. The asymmetry is due to the imbalanced payoff matrix used, such that the basin of attraction for the $i \simeq N$ state is much larger than for the $i \simeq 0$ state. For the parameters chosen in the Fig. 3.10, the selection-balance point is at $i^* \approx 18$. For equal switching rates, $p_+ = p_-$, the averaged stationary distribution $\bar{\rho}_i$ lies exactly in between the two single-environment distributions. The effective distribution ρ_i^{eff} is approximately uniform in the centre of the domain, with a lower probability of being found close to the domain boundaries. This reflects the fact that for equal switching probabilities the effective game is close to neutral, but frequent mutations push the population to the interior.

The exact solution (Eq. (3.36)) is plotted in Fig. 3.10b for a range of magnitudes of $p_+ = p_-$. For large p_σ we are in the fast-switching regime, and hence the exact solution closely matches the effective solution ρ_i^{eff} . For small p_σ we are in the slow-switching regime and the exact solution approaches the averaged solution $\bar{\rho}_i$. For $p_\sigma \lesssim 10^{-2}$, the exact solution matches the features of the single-environment distributions, with a peak at $i \simeq N$ and at the coexistence point i^* . Interestingly, this solution also predicts a (small) peak at the $i \simeq 0$ state, a feature which is not seen in the single-environment distributions, or in the approximate distributions.

The distributions ρ_i^{eff} and $\bar{\rho}_i$ are both approximations. To evaluate the accuracy of these distributions we compute the distance from the exact eigenvector solution given in Eq. (3.36). For the distances we will use the measure




Fig. 3.11 **a** The distance (Eq. 3.40a) between the effective distribution (Eq. 3.37) and the exact stationary distributions (Eq. 3.36). **b** The distance (Eq. 3.40b) between the averaged stationary distribution (Eq. (3.39)) and the exact stationary distribution. The remaining parameters are as in Fig. 3.10


$$d^{\text{eff}} = \frac{1}{2} \sum_{i=0}^N |\rho_i^{\text{eff}} - \rho_i|, \quad (3.40a)$$

$$\bar{d} = \frac{1}{2} \sum_{i=0}^N |\bar{\rho}_i - \rho_i|, \quad (3.40b)$$

which was introduced in Sect. 2.5. These distances are plotted in Fig. 3.11 as a function of switching parameters p_+ and p_- . The approach based on effective transition rates (Fig. 3.11a) is found to be accurate over a large range of switching probabilities away from the slow-switching $p_\sigma \rightarrow 0$ limit. Conversely, the weighted-average distribution (Fig. 3.11b) is inaccurate for a large range of p_σ , but it becomes increasingly accurate if the dynamics of the environment is slow ($p_\sigma \rightarrow 0$). Both approximate distributions accurately predict the exact stationary distribution when the two switching rates are very disparate, i.e. $p_+ \ll p_-$ or vice versa (top-left and bottom-right corners of the two insets). In these regions the environment spends most of the time in one state, so that the model effectively reduces to the single-environment case. All approaches then collapse to the same result, which is the stationary distribution obtained in a single fixed environment.

To verify the accuracy of the analysis, we compare the stationary distribution and its approximations against some numerical data. This data can be obtained in numerous ways. Firstly the master equation (3.3) could be numerically integrated for a long time. Instead, we will compare with the stationary distribution obtained from an ensemble of simulations of the exact stochastic process described by reaction rates (3.35). This distribution is obtained by sampling stochastic simulations at multiple time points to create a set of distributions for the variable i , and then averaging over the set of distributions to obtain a closer approximation, Q_i , to the stationary

Fig. 3.12 The distance (Eq. (3.41)) between the analytic distributions (3.36), (3.37), and (3.39) with the stationary distributions obtained from simulation, Q . The remaining parameters are as in Fig. 3.10



distribution. These distributions are then compared with our analytical distributions using the distance measure

$$d[\rho, Q] = \frac{1}{2} \sum_{i=0}^N |\rho_i - Q_i|, \quad (3.41)$$

where ρ can be the exact, effective or averaged distribution. These distances are shown in Fig. 3.12. The accuracy of the exact solution is confirmed by simulations across many orders of magnitude of switching probabilities. Any deviations can be attributed to incomplete equilibration. For large switching probabilities the effective stationary distribution, ρ^{eff} , approximates the simulation results well. As expected the effective theory becomes inaccurate for slow switching, roughly below $p_\sigma \simeq 10^{-2}$, in our example. The averaged stationary distribution, $\bar{\rho}$, shows the opposite behaviour. It is in reasonable agreement with simulations for slow switching, but shows systematic deviations when the switching process is too fast for the population to react adiabatically.

An alternative approach to obtain the distance from simulation distributions would be to use the time-averaged distance. That is the distance between the sampled simulation distribution and the analytic distributions is taken at each time point, and then the set of distances is averaged over time. This approach was used in the publication associated with this chapter, Ref. [1], and it produces an almost identical picture to Fig. 3.12.

3.6 Summary

The dynamics of a population evolving under changing environmental conditions is an important concept in the study of bacterial populations. Some previous works have focused on deterministic analyses [5], or an environment following a continuous stochastic process [11]. Here we have taken a different route, and assumed that the environment switches between discrete states whilst retaining the demographic

stochasticity of the population. We have developed the mathematical formalism to describe fixation properties in a general birth–death process in an environment fluctuating between an arbitrary number of discrete states. The main results of this investigation are self-consistent expressions for the fixation probability of a mutant in a fixed-size population, as well as for the mean unconditional and conditional fixation times. For short-lived environments we put forward an approximation based on effective transition probabilities.

As a specific application we discuss the fixation properties in the context of an evolutionary game in a two-world scenario. The two states of the environment then correspond to two different payoff matrices of the underlying games. Simulations confirm our analytic solutions over a wide range of switching probabilities. The approximation based on effective transition probabilities is seen to reproduce simulation data in the limit of fast switching.

Focussing on the case of switching between a coexistence game and a coordination game, we find unexpected non-trivial behaviour of the fixation probability of a single mutant. We observe in our analytical results and in simulations that fixation can be more likely in a scenario in which the environment switches between the two games than in either of the constant environments. We provide an intuitive explanation for this effect, and we have investigated in detail the circumstances under which this phenomenon can occur.

Adding mutations to the dynamics removes the possibility of fixation, but introduces non-trivial stationary states. We develop a method for calculating this distribution, along with approximations for both long-lived and short-lived environmental states. These approximations are shown to agree well with simulations in their respective limits.

The general theory developed here now allows further investigation of evolutionary dynamics in time-varying environments. It provides a first mathematical characterisation of the effects one may expect in such systems. The closed-form self-consistent solutions will help to speed up future studies, and they may remove the need for extensive computer simulations.

While our work is mainly mathematical, we think that our theory can be used to interpret existing experimental studies such as those studied by Acar et al. [6]. For some biological systems it may be more appropriate to use constant selection in each environment, as opposed to frequency-dependent selection. Our example of switching between coexistence and coordination games was chosen to illustrate the theory and to show the rich dynamical behaviour that can be observed in these models. We note that both types of game have been observed in systems of experimental evolution [20–22]. We hope the formalism we have developed will be useful to analyse models closer to other biological applications, and potentially to guide future experiments on evolutionary systems in time-dependent environments.

References

1. P. Ashcroft, P.M. Altrock, T. Galla, Fixation in finite populations evolving in fluctuating environments. *J. R. Soc. Interface* **11**, 20140663 (2014)
2. S. Franzenburg, S. Fraune, P. Altrock, S. Kuenzel, J. Baines, A. Traulsen, T.C.G. Bosch, Bacterial colonization of hydra hatchlings follows a robust temporal pattern. *ISME J.* **7**, 781 (2013)
3. M. McFall-Ngai, M.G. Hadfield, T.C. Bosch, H.V. Carey, T. Domazet-Lošo, A.E. Douglas, N. Dubilier, G. Eberl, T. Fukami, S.F. Gilbert et al., Animals in a bacterial world, a new imperative for the life sciences. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 3229 (2013)
4. C.P. Goulart, M. Mahmudi, K.A. Crona, S.D. Jacobs, M. Kallmann, B.G. Hall, D.C. Greene, M. Barlow, Designing antibiotic cycling strategies by determining and understanding local adaptive landscapes. *PLoS ONE* **8**, e56040 (2013)
5. E. Kussell, R. Kishony, N.Q. Balaban, S. Leibler, Bacterial persistence: a model of survival in changing environments. *Genetics* **169**, 1807 (2005)
6. M. Acar, J.T. Mettetal, A. van Oudenaarden, Stochastic switching as a survival strategy in fluctuating environments. *Nat. Genet.* **40**, 471 (2008)
7. E. Kussel, S. Leibler, Phenotypic diversity, population growth, and information in fluctuating environments. *Science* **309**, 2075 (2005)
8. S. Leibler, E. Kussell, Individual histories and selection in heterogeneous populations. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 13183 (2010)
9. O. Rivoire, S. Leibler, The value of information for populations in varying environments. *J. Stat. Phys.* **142**, 1124 (2011)
10. U. Dobramysl, U.C. Täuber, Environmental versus demographic variability in two-species predator-prey models. *Phys. Rev. Lett.* **110**, 048105 (2013)
11. M. Assaf, M. Mobilia, E. Roberts, Cooperation dilemma in finite populations under fluctuating environments. *Phys. Rev. Lett.* **111**, 238101 (2013)
12. N. Goel, N. Richter-Dyn, *Stochastic Models in Biology* (Academic Press, New York, 1974)
13. M.A. Nowak, *Evolutionary Dynamics* (Harvard University Press, Cambridge MA, 2006)
14. A. Traulsen, C. Hauert, Stochastic evolutionary game dynamics, in *Reviews of Nonlinear Dynamics and Complexity*, vol. II, ed. by H.G. Schuster (Wiley-VCH, Weinheim, 2009)
15. S. Karlin, H.M. Taylor, *A Second Course in Stochastic Processes* (Academic Press, New York, 1981)
16. N.G. van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, Amsterdam, 2007)
17. T. Antal, I. Scheuring, Fixation of strategies for an evolutionary game in finite populations. *Bull. Math. Biol.* **68**, 1923 (2006)
18. P.M. Altrock, A. Traulsen, Fixation times in evolutionary games under weak selection. *New J. Phys.* **11**, 013012 (2009)
19. M. Mobilia, Oscillatory dynamics in rock-paper-scissors games with mutations. *J. Theor. Biol.* **264**, 1 (2010)
20. R.C. MacLean, A. Fuentes-Hernandez, D. Greig, L.D. Hurst, I. Gudelj, A mixture of "cheats" and "co-operators" can enable maximal group benefit. *PLoS Biol.* **8**, e1000486 (2010)
21. J. Gore, H. Youk, A. van Oudenaarden, Snowdrift game dynamics and facultative cheating in yeast. *Nature* **459**, 253 (2009)
22. R.C. Maclean, I. Gudelj, Resource competition and social conflict in experimental populations of yeast. *Nature* **441**, 498 (2006)

Chapter 4

Fixation Time Distributions in Birth–Death Processes

4.1 Introduction

As discussed in the previous chapter, the time that it takes for a mutation to reach fixation in a population is one of the fundamental quantities that is predicted by evolutionary dynamics. However, the fixation time is itself a random variable, and while the first moment can provide a good indication of the outcome in some circumstances, this approach can be insufficient when the distribution of fixation times is broad [1, 2]. To provide a complete answer to the question ‘how long does it take for a mutation to fixate in a population?’, we must compute the complete arrival time distribution. In this chapter we will focus on the original birth–death process described in Sect. 2.3. This work has been described in Ref. [3].

Although the master equation (2.14) describing the birth–death dynamics is linear, calculating fixation time distributions is more intricate than one may initially think. As described in Sect. 2.4, nested expressions for all moments of fixation times are known [2, 4, 5]. These are given by Eqs. (2.45) and (2.46). From these the distributions can, in principle, be constructed recursively up to arbitrary precision. However, this approach does not provide a simple closed-form solution or a means of efficiently sampling from the arrival time distribution.

An alternative approach is to diagonalise the linear operator of the master equation and to carry out the analysis in eigenspace. This method was used by Karlin and McGregor¹ to calculate the arrival time distribution at a state $N > 0$ of a birth–death process that has a reflecting boundary at state 0 [6]. Provided that the process is initialised in state 0, Karlin and McGregor showed that the distribution of arrival times can be expressed as the convolution of exponential distributions parametrised by each of the non-zero eigenvalues of the master equation [6, 7]. This is equivalent to stating that the arrival times are given by the sum of independent, exponentially distributed random variables. We will refer to this result as the Karlin–McGregor theorem. It has been examined in numerous sources in the probability theory

¹Samuel Karlin (1924–2007) and James McGregor (1921–1988).

literature [7–15]. However, the discussion of these matters is usually very terse, and not easily accessible to physicists or researchers in adjacent disciplines. Researchers in the theoretical biosciences are only recently beginning to use these ideas for the purpose of model reduction [16, 17]. The existing results are limited to specific initial conditions and types of birth–death chains, and a clear understanding of the analysis in eigenspace is lacking.

In this chapter we consider birth–death processes with two absorbing states and a general initial condition, describing the invasion (or extinction) of a number of mutants in a population of wild-type individuals. The model is described in detail in Sect. 4.2. In Sect. 4.3 we calculate closed-form expressions for the fixation time distributions in terms of the eigenvalues of the master equation, before we turn to the physical interpretation of these and the relation to the Karlin–McGregor results in Sect. 4.4. We illustrate these results by applying this framework to evolutionary games in Sect. 4.5. Here we demonstrate the inadequacy of the mean whilst highlighting the accuracy of our method. We then use our results to relate fixation processes to the equilibration dynamics of evolutionary systems with mutation (and hence with no absorbing states) in Sect. 4.6. In these latter systems the timescale is defined by the mixing time, as described in Sect. 2.5. In the limit of small mutation rates, we identify the relation between the timescales of equilibration and fixation. Finally, we explore the efficiency of our method in Sect. 4.7. We demonstrate that our method is a very effective model-reduction tool, which can generate samples from the arrival time distributions much faster than direct simulation of the original birth–death process. We also show that our calculation of the arrival time distribution in terms of the spectrum of the process is faster than straightforward numerical integration of the master equation.

4.2 Model

We study a continuous-time birth–death process with states $0 \leq i \leq N$, representing a population of constant size N with i individuals of the mutant type and $N - i$ of the resident wild-type. As described in Sect. 2.3, this process is characterised by the birth and death rates b_i and d_i ($0 \leq i \leq N$). Throughout the next section we will not specify a form for these birth and death rates, but in general they are non-linear functions of i . At the boundaries the birth and death rates satisfy $b_N = d_0 = 0$. Initially we will consider the boundaries to be absorbing such that we also have $b_0 = d_N = 0$, as illustrated in Fig. 4.1.

This process is described by the master equation for the probability, $P_i(t)$, to be found in state i at time t , given that the system was started in state i_0 at time $t = 0$. As in Sect. 2.3, we initially suppress the notation for the initial condition to maintain readability. We write the master equation in the matrix form

$$\dot{\mathbf{P}}(t) = \mathbb{W} \cdot \mathbf{P}(t), \quad (4.1)$$




Fig. 4.1 One-step birth–death process in a population of N individuals. The variable i denotes the number of invading mutants. The states $i = 0$ (extinction) and $i = N$ (fixation) are absorbing. Birth rates are labelled b_i and death rates d_i

where $\mathbf{P} = (P_0, P_1, \dots, P_N)^T$, and the $(N + 1) \times (N + 1)$ matrix $\mathbb{W}$ has elements $w_{i,i} = -(b_i + d_i)$, $w_{i+1,i} = b_i$ and $w_{i-1,i} = d_i$. Our objective in the next section is to calculate exactly the arrival time densities at the absorbing states, i.e. $\dot{P}_0(t)$ and $\dot{P}_N(t)$.

4.3 Mathematical Framework

To derive the arrival time densities it is convenient to focus only on the interior states, $1 \leq i \leq N - 1$, of the birth–death process shown in Fig. 4.1. We introduce the lower-case notation $\mathbf{p} = (p_1, \dots, p_{N-1})^T$ for the interior, where $p_i(t) = P_i(t)$ for $1 \leq i \leq N - 1$. We cannot call this quantity a probability distribution as, in general, it is not normalised; there is a constant ‘leaking’ from the interior to the absorbing states. The quantity $\mathbf{p}$ satisfies the equation

$$\dot{\mathbf{p}}(t) = \mathbb{A} \cdot \mathbf{p}(t) \quad (4.2)$$

where the matrix $\mathbb{A}$ is equal to the matrix $\mathbb{W}$ in Eq. (4.1) with the first and last rows and columns removed. We represent this removal by shading the elements to be removed, such that the matrix $\mathbb{A}$ is given by

$$\mathbb{A} = \begin{pmatrix} 0 & d_1 & 0 & & & & & & \\ 0 & -(b_1 + d_1) & d_2 & 0 & & & & & \\ 0 & b_1 & -(b_2 + d_2) & d_3 & & & & & \\ & \ddots & \ddots & \ddots & \ddots & & & & \\ & & 0 & b_{N-3} & -(b_{N-2} + d_{N-2}) & d_{N-1} & 0 & & \\ & & & 0 & b_{N-2} & -(b_{N-1} + d_{N-1}) & 0 & & \\ & & & & 0 & b_{N-1} & 0 & & \end{pmatrix}. \quad (4.3)$$

Ignoring the shading, this is the matrix $\mathbb{W}$. We can represent the $(N + 1) \times (N + 1)$ matrix $\mathbb{W}$ in terms of $(N - 1) \times (N - 1)$ matrix $\mathbb{A}$ in the following way:

$$\mathbb{W} = \begin{pmatrix} 0 & d_1 & 0 & \cdots & 0 \\ \vdots & \left[\begin{array}{c} \mathbb{A} \end{array} \right] & \vdots \\ 0 & \cdots & 0 & b_{N-1} & 0 \end{pmatrix}. \quad (4.4)$$

With this representation it is easy to show the important result that the eigenvalues of $\mathbb{A}$ coincide with the eigenvalues of $\mathbb{W}$. Considering $\det(\mathbb{W} - \lambda \mathbb{I}) = 0$ gives

$$0 = \begin{vmatrix} -\lambda & d_1 & 0 & \cdots & 0 \\ \vdots & \left[\begin{array}{c} \mathbb{A} - \lambda \mathbb{I} \end{array} \right] & \vdots \\ 0 & \cdots & 0 & b_{N-1} & -\lambda \end{vmatrix} = \lambda^2 \det(\mathbb{A} - \lambda \mathbb{I}). \quad (4.5)$$

Thus the eigenvalues of $\mathbb{W}$ are $\lambda = 0$ (with multiplicity 2) and the eigenvalues of $\mathbb{A}$.

We can discern some properties about the eigenvalues of $\mathbb{A}$. As the probability leaks from the interior states to the absorbing states, the truncated master equation (4.2) does not conserve probability and the elements of $\mathbf{p}(t)$ must approach zero at $t \rightarrow \infty$. Hence, all eigenvalues of the operator $\mathbb{A}$ in Eq. (4.2) must have a (strictly) negative real part. Furthermore, as the sub- and super-diagonal elements of $\mathbb{A}$ are all positive, the matrix is sign-symmetric and can be shown to be similar to a symmetric matrix. Hence all the eigenvalues of $\mathbb{A}$ are real. With this we can say that the eigenvalues of $\mathbb{A}$ are (strictly) negative. To avoid numerous minus signs and having to specify absolute values, we will work with the eigenvalues of $-\mathbb{A}$, which are positive. We label these as λ_i for $1 \leq i \leq N - 1$. At this stage we do not specify an ordering of the eigenvalues.

The formal solution to Eq. (4.2) is

$$\mathbf{p}(t) = \exp(\mathbb{A}t) \cdot \mathbf{p}(0), \quad (4.6)$$

where the initial condition is $p_i(0) = \delta_{i,i_0}$ ($1 \leq i_0 \leq N - 1$). To proceed analytically we need to make the matrix exponential in Eq. (4.6) more tractable. To do this we transform Eq. (4.6) from the time domain to the complex frequency domain by taking the Laplace transform. We define the Laplace transform as

$$\widehat{f}(s) = \mathcal{L}[f(t)] = \int_{0^-}^{\infty} f(t) e^{-st} dt. \quad (4.7)$$

The use of 0^- in the lower integration limit allows us to evaluate the Laplace transform of a delta function, $\delta(t)$, to unity. This will be useful later in this derivation.

Writing $\widehat{\mathbf{p}}(s) = \mathcal{L}[\mathbf{p}(t)]$, the Laplace transform of Eq. (4.6) is

$$\begin{aligned}
\widehat{\mathbf{p}}(s) &= \int_{0^-}^{\infty} \mathbf{p}(t) e^{-st} dt \\
&= \int_{0^-}^{\infty} \exp[-(s\mathbb{I} - \mathbb{A})t] \cdot \mathbf{p}(0) dt \\
&= -(s\mathbb{I} - \mathbb{A})^{-1} \cdot \left[e^{-(s\mathbb{I} - \mathbb{A})t} \right]_{0^-}^{\infty} \cdot \mathbf{p}(0) \\
&= (s\mathbb{I} - \mathbb{A})^{-1} \cdot \mathbf{p}(0).
\end{aligned} \tag{4.8}$$

As we know that $\lim_{t \rightarrow \infty} \mathbf{p}(t) = \mathbf{0}$, this integral is convergent at least for all $\text{Re}(s) > 0$.

Our strategy is to compute $\widehat{p}_1(s)$ and $\widehat{p}_{N-1}(s)$, and from these $\widehat{P}_0(s) = d_1 \widehat{p}_1(s)$ and $\widehat{P}_N(s) = b_{N-1} \widehat{p}_{N-1}(s)$. Then by performing the inverse transform, we can recover the arrival time densities $\dot{P}_0(t)$ and $\dot{P}_N(t)$.

Substituting $p_i(0) = \delta_{i,i_0}$ for the initial condition in Eq. (4.8), and reinstating the explicit notation, the quantities we want to evaluate are

$$\widehat{p}_{1|i_0}(s) = [(s\mathbb{I} - \mathbb{A})^{-1}]_{1,i_0}, \tag{4.9a}$$

$$\widehat{p}_{N-1|i_0}(s) = [(s\mathbb{I} - \mathbb{A})^{-1}]_{N-1,i_0}. \tag{4.9b}$$

To proceed, we recall that the (i, j) -th element of the inverse of any invertible matrix $\mathbb{B}$ is given by $[\mathbb{B}^{-1}]_{i,j} = C_{j,i} / \det \mathbb{B}$, where $C_{j,i}$ is the (j, i) -th co-factor of $\mathbb{B}$. Thus we can write

$$\widehat{p}_{1|i_0}(s) = [(s\mathbb{I} - \mathbb{A})^{-1}]_{1,i_0} = \frac{1}{\det(s\mathbb{I} - \mathbb{A})} C_{i_0,1}, \tag{4.10}$$

and likewise for the $(N-1, i_0)$ -th element.

To calculate the co-factor $C_{i_0,1}$, we remove row i_0 and column 1 from $s\mathbb{I} - \mathbb{A}$ and evaluate the determinant. Again we indicate removal of elements by shading, such that the co-factor is given by

$$C_{i_0,1} = (-1)^{i_0+1} \begin{vmatrix} s+a_1 & -d_2 & 0 & & & \\ -b_1 & s+a_2 & -d_3 & & & \\ \vdots & \ddots & \ddots & \ddots & \ddots & \\ 0 & -b_{i_0-1} & s+a_{i_0} & -d_{i_0+1} & 0 & \\ & \ddots & \ddots & \ddots & \ddots & \\ & & 0 & -b_{N-3} & s+a_{N-2} & -d_{N-1} \\ & & & 0 & -b_{N-2} & s+a_{N-1} \end{vmatrix}, \tag{4.11}$$

where we have introduced the notation $a_i = b_i + d_i$ for compactness. Using Laplace's formula, the co-factor can be expressed as

$$\begin{aligned}
C_{i_0,1} &= (-1)^{i_0+1} \left(\prod_{i=2}^{i_0} -d_i \right) \begin{vmatrix} s + a_{i_0+1} & -d_{i_0+2} & 0 & \\ -b_{i_0+1} & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & -d_{N-1} \\ & 0 & -b_{N-2} & s + a_{N-1} \end{vmatrix} \\
&= \left(\prod_{i=2}^{i_0} d_i \right) \det(s\mathbb{I} - \mathbb{A}_{(N-i_0-1)}) \\
&= \prod_{i=2}^{i_0} d_i \prod_{\alpha=1}^{N-i_0-1} (s + x_\alpha). \tag{4.12}
\end{aligned}$$

The matrix $\mathbb{A}_{(N-i_0-1)}$ consists of the rows and columns $i_0 + 1, \dots, N - 1$ of the matrix $\mathbb{A}$, i.e. it is the bottom right $(N - i_0 - 1) \times (N - i_0 - 1)$ sub-matrix of $\mathbb{A}$. The matrix $-\mathbb{A}_{(N-i_0-1)}$ has eigenvalues $x_\alpha > 0$ ($1 \leq \alpha \leq N - i_0 - 1$) and determinant

$$\det(-\mathbb{A}_{(N-i_0-1)}) = \prod_{\alpha=1}^{N-i_0-1} x_\alpha = \chi_{i_0}. \tag{4.13}$$

The $(i_0, N - 1)$ -th cofactor is required for the computation of $\widehat{p}_{N-1|i_0}(s)$. Thus we remove row i_0 and column $N - 1$ from $s\mathbb{I} - \mathbb{A}$ (again indicated by shading) and evaluate the determinant, such that

$$C_{i_0,N-1} = (-1)^{i_0+N-1} \begin{vmatrix} s + a_1 & -d_2 & 0 & & & \\ -b_1 & s + a_2 & -d_3 & 0 & & \\ \ddots & \ddots & \ddots & \ddots & \ddots & \\ 0 & -b_{i_0-1} & s + a_{i_0} & -d_{i_0+1} & 0 & \\ & \ddots & \ddots & \ddots & \ddots & \\ & & 0 & -b_{N-3} & s + a_{N-2} & -d_{N-1} \\ & & & 0 & -b_{N-2} & s + a_{N-1} \end{vmatrix}. \tag{4.14}$$

Again using Laplace's formula we arrive at

$$\begin{aligned}
C_{i_0,N-1} &= \left(\prod_{i=i_0}^{N-2} b_i \right) \det(s\mathbb{I} - \mathbb{A}^{(i_0-1)}) \\
&= \prod_{i=i_0}^{N-2} b_i \prod_{\alpha=1}^{i_0-1} (s + y_\alpha). \tag{4.15}
\end{aligned}$$

The matrix $\mathbb{A}^{(i_0-1)}$ consists of rows and columns $1, \dots, i_0 - 1$ of the matrix $\mathbb{A}$, i.e. it is the top-left $(i_0 - 1) \times (i_0 - 1)$ sub-matrix of $\mathbb{A}$. The matrix $-\mathbb{A}^{(i_0-1)}$ has eigenvalues $y_\alpha > 0$ ($1 \leq \alpha \leq i_0 - 1$) and determinant

$$\det(-\mathbb{A}^{(i_0-1)}) = \prod_{\alpha=1}^{i_0-1} y_\alpha = \psi_{i_0}. \quad (4.16)$$

Putting things together, and writing $\det(s\mathbb{I} - \mathbb{A}) = \prod_{\gamma=1}^{N-1} (s + \lambda_\gamma)$, we have the compact expressions

$$\widehat{p}_{1|i_0}(s) = \prod_{i=2}^{i_0} d_i \prod_{\alpha=1}^{N-i_0-1} (s + x_\alpha) \prod_{\gamma=1}^{N-1} \frac{1}{s + \lambda_\gamma}, \quad (4.17a)$$

$$\widehat{p}_{N-1|i_0}(s) = \prod_{i=i_0}^{N-2} b_i \prod_{\alpha=1}^{i_0-1} (s + y_\alpha) \prod_{\gamma=1}^{N-1} \frac{1}{s + \lambda_\gamma}. \quad (4.17b)$$

Using $\widehat{P}_{0|i_0}(s) = d_1 \widehat{p}_{1|i_0}(s)$ and $\widehat{P}_{N|i_0}(s) = b_{N-1} \widehat{p}_{N-1|i_0}(s)$, we obtain the Laplace transforms of the arrival time densities at sites $i = 0$ and $i = N$, respectively. They are given by

$$\widehat{P}_{0|i_0}(s) = D_{i_0} \prod_{\alpha=1}^{N-i_0-1} (s + x_\alpha) \prod_{\gamma=1}^{N-1} \frac{1}{s + \lambda_\gamma}, \quad (4.18a)$$

$$\widehat{P}_{N|i_0}(s) = B_{i_0} \prod_{\alpha=1}^{i_0-1} (s + y_\alpha) \prod_{\gamma=1}^{N-1} \frac{1}{s + \lambda_\gamma}, \quad (4.18b)$$

where $D_{i_0} = \prod_{i=1}^{i_0} d_i$ and $B_{i_0} = \prod_{i=i_0}^{N-1} b_i$.

To return to the time domain, we need to perform the inverse Laplace transform. To do this we will make use of the convolution theorem. For two functions $\mathcal{F}(t)$ and $\mathcal{G}(t)$ defined in the domain $t \geq 0$, the convolution of the two functions is defined as

$$[\mathcal{F} * \mathcal{G}](t) = \int_0^\infty \mathcal{F}(t - \tau) \mathcal{G}(\tau) d\tau, \quad (4.19)$$

where ‘ $*$ ’ is the convolution operator. The convolution theorem then states

$$\mathcal{L}^{-1}[\widehat{\mathcal{F}}(s) \cdot \widehat{\mathcal{G}}(s)] = [\mathcal{F} * \mathcal{G}](t). \quad (4.20)$$

Therefore, if the terms $s + x$ and $(s + \lambda)^{-1}$ in Eq. (4.18) can be expressed as Laplace transforms of specific functions, then the arrival time densities $\widehat{P}_{0|i_0}(t)$ and $\widehat{P}_{N|i_0}(t)$ can be simply expressed as the convolution of those functions.

It is easy to show that $(s + \lambda)^{-1}$ is proportional to the Laplace transform of an exponential distribution with parameter $\lambda > 0$: The exponential distribution is defined by

$$\mathcal{E}^{(\lambda)}(t) = \lambda e^{-\lambda t} \quad \text{for } t \geq 0. \quad (4.21)$$

The Laplace transform is obtained as follows

$$\begin{aligned}\mathcal{L}[\mathcal{E}^{(\lambda)}(t)] &= \int_{0-}^{\infty} \lambda e^{-\lambda t} e^{-st} dt \\ &= \lambda \int_{0-}^{\infty} e^{-(s+\lambda)t} dt.\end{aligned}\quad (4.22)$$

This integral is convergent in the region $\text{Re}(s) > -\lambda$. Within this region we have

$$\mathcal{L}[\mathcal{E}^{(\lambda)}(t)] = \frac{\lambda}{s + \lambda}, \quad (4.23)$$

and hence

$$\mathcal{L}^{-1}[(s + \lambda)^{-1}] = \frac{1}{\lambda} \mathcal{E}^{(\lambda)}(t). \quad (4.24)$$

Deriving the function of time that transforms into $s + z$ (where $z = x_\alpha$ or y_α) is a little more complicated, but we can show that $\delta(t) + z^{-1}\delta'(t)$ satisfies this condition (up to multiplication by a constant). The object $\delta'(t)$ is the derivative of the Dirac-delta distribution $\delta(t)$, which can be defined conveniently by its Fourier transform [18]. It has the form

$$\delta'(t) = \int_{-\infty}^{\infty} (i\omega) e^{i\omega t} d\omega. \quad (4.25)$$

Evaluating the Laplace transform of $\delta(t) + z^{-1}\delta'(t)$ gives

$$\begin{aligned}\mathcal{L}[\delta(t) + z^{-1}\delta'(t)] &= \int_{0-}^{\infty} [\delta(t) + z^{-1}\delta'(t)] e^{-st} dt \\ &= \int_{0-}^{\infty} e^{-st} \delta(t) dt + z^{-1} [e^{-st} \delta(t)]_{0-}^{\infty} + z^{-1} s \int_{0-}^{\infty} e^{-st} \delta(t) dt \\ &= 1 + z^{-1}s,\end{aligned}\quad (4.26)$$

where we have used $\lim_{t \rightarrow 0-} \delta(t) = 0$. This expression has no singularities, and thus the region of convergence in terms of s is the entire complex plane. Hence the inverse Laplace transform of $s + z$ is given by

$$\mathcal{L}^{-1}[s + z] = z [\delta(t) + z^{-1}\delta'(t)]. \quad (4.27)$$

It is useful here to define some further properties of the object $\delta'(t)$. When it is convolved with a test function $\mathcal{F}(t)$ with infinite support, one obtains (after integration by parts)

$$\int_{-\infty}^{\infty} \delta'(t - \tau) \mathcal{F}(\tau) d\tau = \mathcal{F}'(t). \quad (4.28)$$

If a test function $\mathcal{G}(t)$ has finite support, say $t \geq 0$, then one finds

$$\int_0^\infty \delta'(t - \tau) \mathcal{G}(\tau) d\tau = \mathcal{G}'(t) + \mathcal{G}(0) \delta(t). \quad (4.29)$$

Using the convolution theorem (4.19), and Eqs. (4.24) and (4.27), the inverse Laplace transform of Eq. (4.18) is

$$\dot{P}_{0|i_0}(t) = \frac{D_{i_0} \chi_{i_0}}{\Lambda} \mathcal{E}^{(\lambda_1)} * \dots * \mathcal{E}^{(\lambda_{N-1})} * (\delta + x_1^{-1} \delta') * \dots * (\delta + x_{N-i_0-1}^{-1} \delta'), \quad (4.30a)$$

$$\dot{P}_{N|i_0}(t) = \frac{B_{i_0} \psi_{i_0}}{\Lambda} \mathcal{E}^{(\lambda_1)} * \dots * \mathcal{E}^{(\lambda_{N-1})} * (\delta + y_1^{-1} \delta') * \dots * (\delta + y_{i_0-1}^{-1} \delta'), \quad (4.30b)$$

where $\Lambda = \det(-\mathbb{A}) = \prod_{\alpha=1}^{N-1} \lambda_\alpha$. For compactness we introduce the notation

$$E_\ell = \mathcal{E}^{(\lambda_1)} * \dots * \mathcal{E}^{(\lambda_\ell)}, \quad (4.31a)$$

$$R_\ell = (\delta + y_1^{-1} \delta') * \dots * (\delta + y_\ell^{-1} \delta'), \quad (4.31b)$$

such that we can write Eq. (4.30b) as

$$\dot{P}_{N|i_0}(t) = \frac{B_{i_0} \psi_{i_0}}{\Lambda} E_{N-1} * R_{i_0-1}. \quad (4.32)$$

The compact version of Eq. (4.30a) follows analogously.


Before turning to the physical interpretation of Eq. (4.32), it is useful to evaluate the convolution of an exponential distribution $\mathcal{E}^{(\lambda)}$ with an object of the form $\delta + y^{-1} \delta'$. Using the result of Eq. (4.29), we arrive at

$$\begin{aligned} \mathcal{E}^{(\lambda)} * (\delta + y^{-1} \delta') &= \int_0^\infty \lambda e^{-\lambda \tau} [\delta(t - \tau) + y^{-1} \delta'(t - \tau)] d\tau \\ &= \lambda e^{\lambda t} - \lambda^2 y^{-1} e^{\lambda t} + \lambda y^{-1} \delta(t) \\ &= \frac{\lambda}{y} \delta(t) + \left(1 - \frac{\lambda}{y}\right) \mathcal{E}^{(\lambda)}(t). \end{aligned} \quad (4.33)$$

Assuming $\lambda/y < 1$ (which will be the case throughout our analysis) this describes a convex combination of a point-mass at zero and an exponential distribution. To obtain samples from this hybrid distribution, one chooses $t = 0$ with probability λ/y , otherwise t is drawn from $\mathcal{E}^{(\lambda)}(t)$. An example of this distribution is illustrated in Fig. 4.2.

The result (4.33) has an important implication for the interpretation of Eq. (4.32). As Eq. (4.33) describes a distribution, the quantity $E_{N-1} * R_{i_0-1}$ in Eq. (4.32) is also a distribution. Thus we have

Fig. 4.2 A comparison between the distribution (4.33) and the exponential distribution (4.21). Parameters are $\lambda = 5$, $y = 25$, and the histograms are generated from 10^5 samples of each distribution



$$\begin{aligned} \int_0^\infty \dot{P}_{N|i_0}(t) dt &= \frac{B_{i_0} \psi_{i_0}}{\Lambda} \int_0^\infty E_{N-1} * R_{i_0-1} dt = \frac{B_{i_0} \psi_{i_0}}{\Lambda} \\ \Rightarrow \phi_{N|i_0} &= \frac{B_{i_0} \psi_{i_0}}{\Lambda}, \end{aligned} \quad (4.34)$$

where $\phi_{N|i_0}$ is the fixation probability of i_0 mutants.

Equations (4.30) have further implications. Choosing the initial conditions $i_0 = N - 1$ in Eq. (4.30a) and $i_0 = 1$ in Eq. (4.30b) means there are no objects of the form $\delta + z^{-1}\delta'$. Hence Eqs. (4.30) reduce to

$$\dot{P}_{0|N-1}(t) = \mathcal{E}^{(\lambda_1)} * \dots * \mathcal{E}^{(\lambda_{N-1})} = \phi_{0|N-1} E_{N-1}, \quad (4.35a)$$

$$\dot{P}_{N|1}(t) = \frac{B_1}{\Lambda} \mathcal{E}^{(\lambda_1)} * \dots * \mathcal{E}^{(\lambda_{N-1})} = \phi_{N|1} E_{N-1}. \quad (4.35b)$$

From this we conclude

$$\frac{\dot{P}_{0|N-1}(t)}{\phi_{0|N-1}} = \frac{\dot{P}_{N|1}(t)}{\phi_{N|1}}, \quad (4.36)$$

that is the conditional arrival time distribution at state $i = 0$ given $i_0 = N - 1$ is equal to the conditional arrival time distribution at state $i = N$ given $i_0 = 1$. This symmetry has been known for the mean fixation time [19, 20], and it was recently shown that the correspondence holds for the full distribution [17]. Our approach offers an alternative way to obtain this intriguing result.

4.4 Physical Interpretation

We now discuss two possible interpretations of Eq. (4.32) [or equivalently Eqs. (4.30)]. These different representations arise because the convolution operator (4.19) is commutative, such that we can order the convolutions in Eq. (4.32) in multiple ways. We proceed to analyse the different cases separately. In this section we only focus on arrival at state N . Interpretations for arrival at state 0 follow analogously.

Pairing $\delta + y_\alpha^{-1}\delta'$ with Individual Exponential Distributions $\mathcal{E}^{(\lambda_\gamma)}$

For want of a catchier title, this process describes evaluating convolutions between each object $\delta + y_\alpha^{-1}\delta'$ in Eq. (4.32) with a separate exponential distribution $\mathcal{E}^{(\lambda_\gamma)}$. We choose to couple the object $\delta + y_\alpha^{-1}\delta'$ with $\mathcal{E}^{(\lambda_{N-\alpha})}$. Using the result of Eq. (4.33), we can write Eq. (4.32) as

$$\begin{aligned} \frac{\dot{P}_{N|i_0}(t)}{\phi_{N|i_0}} &= \mathcal{E}^{(\lambda_1)} * \dots * \mathcal{E}^{(\lambda_{N-i_0})} * \\ &\quad \left[\frac{\lambda_{N-i_0+1}}{y_{i_0-1}} \delta + \left(1 - \frac{\lambda_{N-i_0+1}}{y_{i_0-1}} \right) \mathcal{E}^{(\lambda_{N-i_0+1})} \right] * \dots \\ &\quad \dots * \left[\frac{\lambda_{N-1}}{y_1} \delta + \left(1 - \frac{\lambda_{N-1}}{y_1} \right) \mathcal{E}^{(\lambda_{N-1})} \right]. \end{aligned} \quad (4.37)$$

We stress that the objects $\delta + y_\alpha^{-1}\delta'$ can be paired with any of the exponential distributions. We chose to match them in this way so that the reduced chains can be systematically compared. We now ensure that the eigenvalues are ordered such that $\lambda_{N-\alpha}/y_\alpha < 1$ for all α . Such an ordering is always possible due to the interlacing property of the eigenvalues [10].

We now can see that the arrival time distribution, $\dot{P}_{N|i_0}(t)/\phi_{N|i_0}$, is given by the distribution of the sum of $N - 1$ random variables drawn from exponential distributions or the mixed distribution in Eq. (4.33). This means that an arrival time [a sample of the distribution (4.37)] can be expressed as the sum of these $N - 1$ random variables. With this we can construct a forward-only process consisting of $N - 1$ jumps from states 1 to N in eigenspace. The first $N - i_0$ jumps are compulsory exponential steps, whereas the final $i_0 - 1$ steps are each exponentials that have a finite probability ($\lambda_{N-\alpha}/y_\alpha$) of being skipped.

We can make further simplifications to this process by multiplying out the brackets in Eq. (4.37). This creates a total of 2^{i_0-1} possible forward-only channels with up to $i_0 - 1$ exponential steps skipped, as shown in Fig. 4.3.

Arrival time samples of the original process are generated from Fig. 4.3 in the following way: One of the channels is chosen with probability determined by products of the terms $\lambda_{N-\alpha}/y_\alpha$ and $1 - \lambda_{N-\alpha}/y_\alpha$, which emerge from the expansion of the brackets in Eq. (4.37). After a channel has been selected, the clock is started and the forward-only process of the channel is traversed. The clock is stopped when the final state in the schematic is reached ('absorption').

Recursively Convoluting $\delta + y_\alpha^{-1}\delta'$ with the Exponential Chain E_ℓ

An alternative approach is to successively convolve the $i_0 - 1$ objects of the form $\delta + y_\alpha^{-1}\delta'$ with the full exponential chain E_{N-1} from the right. We note that

$$E_\kappa * (\delta + y_\alpha^{-1}\delta') = \left[\frac{\lambda_\kappa}{y_\alpha} E_{\kappa-1}(t) + \left(1 - \frac{\lambda_\kappa}{y_\alpha} \right) E_\kappa(t) \right], \quad (4.38)$$




Fig. 4.3 A set of forward-only processes in eigenspace. The λ_α are eigenvalues of $-\mathbb{A}$, and each arrow represents an exponential process with the rate indicated. In each run one channel is chosen with appropriate probability. Transitions indicated by *dashed arrows* are skipped (zero time). This process generates arrival times distributed identically to the process in Fig. 4.1. The case shown here is for arrival at N , starting from $i_0 = 3$ in the original space, such that a maximum of two possible steps can be skipped

which follows directly from Eq. (4.33). Using this we see that each of the recursive convolutions introduces a new exponential chain with one step less. Thus, by performing all the convolutions, the arrival time distribution can be expressed as a linear combination of the distributions $E_{N-\alpha}$ ($1 \leq \alpha \leq i_0$), which we write as

$$\frac{\dot{P}_{N|i_0}(t)}{\phi_{N|i_0}} = \sum_{\alpha=1}^{i_0} G_{N-\alpha}^{(y)} E_{N-\alpha}(t), \quad (4.39)$$

where the $G_{N-\alpha}^{(y)}$ are constants (independent of time). These coefficients are given by

$$G_{N-\alpha}^{(y)} = \frac{1}{\psi_{i_0}} \left(\prod_{\gamma=1}^{\alpha-1} \lambda_{N-\gamma} \right) \sum_{\sigma_1=1}^{\alpha} (y_{\sigma_1} - \lambda_{N-\sigma_1}) \sum_{\sigma_2=\sigma_1}^{\alpha} (y_{\sigma_2+1} - \lambda_{N-\sigma_2}) \sum_{\sigma_3=\sigma_2}^{\alpha} \dots \sum_{\substack{\sigma_{i_0-\alpha} = \\ \sigma_{i_0-\alpha-1}}}^{\alpha} (y_{\sigma_{i_0-\alpha}+i_0-\alpha-1} - \lambda_{N-\sigma_{i_0-\alpha}}). \quad (4.40)$$

We note here that the $G_{N-\alpha}^{(y)}$ must satisfy

$$\sum_{\alpha=1}^{i_0} G_{N-\alpha}^{(y)} = 1, \quad (4.41)$$

for Eq. (4.39) to be normalised.




Fig. 4.4 A single forward-only chain in eigenspace in which the final state can be reached directly from some of the intermediate states. The λ_α are eigenvalues of $-\mathbb{A}$, and each *arrow* represents an exponential process with the rate indicated. The quantity F_α denotes the probability that the next state of the dynamics in eigenspace is $\alpha + 1$, as opposed to N , if the current state is α . This process generates arrival times distributed identically to the process in Fig. 4.1. The case shown here is for arrival at N , starting from $i_0 = 3$ in the original space

From Eq. (4.39), we can generate samples from the arrival time distribution by choosing one of i_0 exponential channels with probability $G_{N-\alpha}^{(y)}$ in which the last $\alpha - 1$ exponential steps are skipped ($1 \leq \alpha \leq i_0$). We can again make further simplifications to this process by expressing the linear combination (4.39) as the single chain shown in Fig. 4.4. In this representation the system can transition to two possible states if currently in eigenstate α : either $\alpha \rightarrow \alpha + 1$ or $\alpha \rightarrow N$. These paths have transition rate $F_\alpha \lambda_\alpha$ and $(1 - F_\alpha) \lambda_\alpha$, respectively. The total transition rate out of eigenstate α is then λ_α , and the waiting time at α is an exponential distribution with parameter λ_α independent of whether the system transitions to $\alpha + 1$ or to N . The quantity F_α denotes the probability that the next state of dynamics in eigenspace is $\alpha + 1$, if the system is currently in eigenstate α . With probability $1 - F_\alpha$ the next state is eigenstate N . Evaluating the probability of a trajectory in terms of F_α , and then matching with Eq. (4.39) gives

$$F_\alpha = \frac{1 - \sum_{\kappa=1}^{\alpha} G_{\kappa}^{(y)}}{1 - \sum_{\kappa=1}^{\alpha-1} G_{\kappa}^{(y)}} \quad \text{for } \alpha < N - 1. \quad (4.42)$$

Arrival time samples are generated from Fig. 4.4 by traversing the forward-only chain, which can easily be simulated with the Gillespie algorithm [21], see Sect. 2.9. In practice, however, we find that evaluating the recursive sums in Eq. (4.40) is inefficient in comparison with the method described in Fig. 4.3.

Comparison of the Two Approaches

The representation shown in Fig. 4.3 corresponds to the picture obtained for a restricted set of processes by probabilistic methods in Ref. [10]. On the other hand, Fig. 4.4 reflects the findings of Refs. [13] and [14], derived from the construction of so-called ‘intertwining processes’. Our analysis shows that these different decompositions originate from one common structure, Eq. (4.32). The explicit schemes

in Figs. 4.3 and 4.4 provide a computational method to generate samples from the arrival time distribution, for example by carrying out simulations of these forward processes using the Gillespie algorithm [21]. It is important to keep in mind that the eigenstates shown in Figs. 4.3 and 4.4 cannot be mapped one-to-one to the real-space states in Fig. 4.1. The equivalence of the real and eigenspace representations only holds on the level of arrival time statistics.

‘Bottom-Line’ Arrival Time Distributions

The final expressions for the arrival time distributions follow directly from Eq. (4.39). First we note that the convolution of ℓ exponential distributions has the form

$$E_\ell(t) = \left(\prod_{\alpha=1}^{\ell} \lambda_\alpha \right) \sum_{\alpha=1}^{\ell} \prod_{\substack{\gamma=1 \\ \gamma \neq \alpha}}^{\ell} \frac{1}{\lambda_\gamma - \lambda_\alpha} e^{-\lambda_\alpha t}. \quad (4.43)$$

We note that this expression only holds if the eigenvalues λ_α are distinct. The convolution of two identical exponential distributions is given by a Gamma distribution, but there is no neat general expression for the convolution of multiple exponential and Gamma distributions with different parameters.

Substituting the result (4.43) into Eq. (4.39), we arrive at the final expression for the conditional arrival time distribution at state N ,

$$\begin{aligned} \frac{\dot{P}_{N|i_0}(t)}{\phi_{N|i_0}} &= \left(\prod_{\alpha=1}^{N-1} \lambda_\alpha \right) \frac{1}{\psi_{i_0}} \sum_{\alpha=1}^{N-1} \left[\frac{\prod_{\gamma=1}^{i_0-1} (y_\gamma - \lambda_\alpha)}{\prod_{\substack{\gamma=1 \\ \gamma \neq \alpha}}^{N-1} (\lambda_\gamma - \lambda_\alpha)} \right] e^{-\lambda_\alpha t} \\ &= \frac{\Lambda}{\psi_{i_0}} \sum_{\alpha=1}^{N-1} C_\alpha^{(y)} e^{-\lambda_\alpha t}. \end{aligned} \quad (4.44)$$

From Eq. (4.44) it is simple to compute the mean fixation time. This is given by

$$\langle t \rangle = \int_0^\infty t \frac{\dot{P}_{N|i_0}(t)}{\phi_{N|i_0}} dt = \frac{\Lambda}{\psi_{i_0}} \sum_{\alpha=1}^{N-1} \frac{C_\alpha^{(y)}}{\lambda_\alpha^2}. \quad (4.45)$$

In fact all higher moments can be computed just as easily, such that

$$\langle t^r \rangle = \int_0^\infty t^r \frac{\dot{P}_{N|i_0}(t)}{\phi_{N|i_0}} dt = \frac{\Lambda}{\psi_{i_0}} \sum_{\alpha=1}^{N-1} \frac{C_\alpha^{(y)} r!}{\lambda_\alpha^{r+1}}. \quad (4.46)$$

Finally, we state that the conditional arrival time distribution at state 0 is

$$\begin{aligned} \frac{\dot{P}_{0|i_0}(t)}{\phi_{0|i_0}} &= \left(\prod_{\alpha=1}^{N-1} \lambda_{\alpha} \right) \frac{1}{\chi_{i_0}} \sum_{\alpha=1}^{N-1} \left[\frac{\prod_{\substack{\gamma=1 \\ \gamma \neq \alpha}}^{N-i_0-1} (x_{\gamma} - \lambda_{\alpha})}{\prod_{\substack{\gamma=1 \\ \gamma \neq \alpha}}^{N-1} (\lambda_{\gamma} - \lambda_{\alpha})} e^{-\lambda_{\alpha} t} \right] \\ &= \frac{\Lambda}{\chi_{i_0}} \sum_{\alpha=1}^{N-1} C_{\alpha}^{(x)} e^{-\lambda_{\alpha} t}, \end{aligned} \quad (4.47)$$

and the unconditional distribution is given by $\dot{P}_{0|i_0}(t) + \dot{P}_{N|i_0}(t)$.

4.5 Application to Evolutionary Games

As an application of this theory we now consider examples of evolutionary dynamics with frequency-dependent selection, as described in Sect. 2.8. For this example we assume a pairwise-comparison process, leading to the birth and death rates

$$b_i = \frac{i(N-i)}{N} g[+\Delta\pi(i)], \quad (4.48a)$$

$$d_i = \frac{i(N-i)}{N} g[-\Delta\pi(i)], \quad (4.48b)$$

where $\Delta\pi(i) = \pi_A(i) - \pi_B(i)$ and $g(z) = (1 + \beta z)/2$ for selection intensity β . The payoffs are as described in Eq. (2.74), and the payoff matrix is parametrised by R , S , T , and P , as described in Eq. (2.73).

As shown in Fig. 4.5, arrival time distributions in the different games can be broad and skewed, such that the mean fixation time contains only limited information. Our theory, however, accurately captures the complete arrival time distributions which are measured from simulations, giving complete information about the arrival time statistics.

The generalisation of the Karlin–McGregor result to include arbitrary initial conditions was a crucial step. As seen in Fig. 4.6, the initial condition can have a great effect on the conditional arrival time distribution, illustrated here for the Prisoner's dilemma. When the initial condition i_0 is close to the final state N , the distribution is much more peaked than is the case when i_0 is close to state 0. The inset of Fig. 4.6 shows that the distribution tails all decay with the same exponent, given by the slowest (closest to zero) eigenvalue of the matrix $\mathbb{A}$. However, the amplitude of the distribution in the tails is vastly different across the initial conditions.





Fig. 4.5 The conditional fixation time distributions at $i = N$ for different games. *Lines* show results from the theory, Eq. (4.44). *Symbols* are from Gillespie simulations (10^6 runs per game) of the birth–death process. The mean fixation times (*arrows*) are not a good description of the distributions, especially for the coexistence game where the distribution is very broad. Parameters used in this figure are $N = 100$, $i_0 = 10$, $\beta = 0.1$. The payoff matrix parameters are: Coexistence game: $R = P = 1.0$, $S = T = 1.5$; Coordination game: $R = P = 1.5$, $S = T = 1.0$; Prisoner's dilemma: $R = -S = 0.5$, $T = 1.0$, $P = 0.0$)

Fig. 4.6 The conditional fixation time distributions at $i = N$ for different initial conditions in the Prisoner's dilemma. *Lines* show results from the theory, Eq. (4.44). Inset plot is the same data represented on a logarithmic vertical axis. Parameters used are the same as Fig. 4.5



4.6 Equilibration Processes in Systems with Mutation

Now that we have calculated the complete arrival time distribution, we will use it to establish a link between fixation processes and the approach to stationary distributions. For this we consider birth–death processes without absorbing states. As described in Sect. 3.5, this can be achieved by adding mutation occurring at a rate $u \ll 1$, such that $b_0 = \mathcal{O}(u)$ and $d_N = \mathcal{O}(u)$. All other transition rates depicted in Fig. 4.1 are $\mathcal{O}(u^0)$ and are only affected at sub-leading order by u . We will consider transition rates of the form

$$b_i = (1 - u) \frac{i(N - i)}{N} g[+\Delta\pi(i)] + \frac{u}{2} \frac{(N - i)^2}{N}, \quad (4.49a)$$

$$d_i = (1 - u) \frac{i(N - i)}{N} g[-\Delta\pi(i)] + \frac{u}{2} \frac{i^2}{N}, \quad (4.49b)$$

where we keep the same pairwise-comparison process as used in the previous section, i.e. $g(z) = (1 + \beta z)/2$, and the payoffs are not affected by u .

Throughout this section we will focus on two games: a symmetric coordination game and the Prisoner's dilemma. These games are described by the payoff matrices

$$\text{Coordination game:} \quad \begin{array}{c|cc} & A & B \\ \hline A & 1.5 & 1.0 \\ B & 1.0 & 1.5, \end{array} \quad (4.50a)$$

$$\text{Prisoner's dilemma:} \quad \begin{array}{c|cc} & A & B \\ \hline A & 0.5 & -0.5 \\ B & 1.0 & 0.0, \end{array} \quad (4.50b)$$

and the payoff functions are given in Eq. (2.74). In Fig. 4.7, we show the phase portraits of the deterministic dynamics of these games with and without mutation. As described in Sect. 2.6, we have $x = \lim_{N \rightarrow \infty} i/N$ and $\dot{x} = \lim_{N \rightarrow \infty} (b_i - d_i)/N$. The effect of mutation is to push the stable (previously absorbing) fixed points towards the centre of the domain. When mutation is present, there is always some ‘force’ pushing the population away from the boundaries.

As there are no absorbing states when $u > 0$, the dynamics reaches a stationary distribution, $\mathbf{P}^{\text{st}}$, with full support, i.e. $P_i^{\text{st}} > 0$ for all $0 \leq i \leq N$. From Eq. (2.54), this distribution can be expressed as

$$P_0^{\text{st}} = \left(\sum_{i=0}^N \prod_{j=1}^i \frac{b_{j-1}}{d_j} \right)^{-1}, \quad P_{i>0}^{\text{st}} = \left(\prod_{j=1}^i \frac{b_{j-1}}{d_j} \right) P_0^{\text{st}}. \quad (4.51)$$

Examples of these distributions are shown in Fig. 4.8a for the coordination game and Fig. 4.8b for the Prisoner's dilemma. For large values of u , the mutation dominates




Fig. 4.7 Phase portraits of processes described by the rates (4.49) with and without mutation. Arrows indicate the sign of $\dot{x}$, which represents the direction of selection. Panel (a) shows the deterministic dynamics of the (symmetric) coordination game, with payoffs from Eq. (4.50a). Panel (b) shows the deterministic dynamics of the Prisoner's dilemma, with payoffs from Eq. (4.50b). For both panels the selection intensity is $\beta = 0.1$.




Fig. 4.8 Stationary distributions of the birth–death process described by rates (4.49) for different values of u , as described in Eq. (4.51). Panel (a) shows stationary distributions of the (symmetric) coordination game, with payoffs from Eq. (4.50a). Panel (b) shows the stationary distributions of the Prisoner’s dilemma, with payoffs from Eq. (4.50b). For both panels the selection intensity is $\beta = 0.1$ and the system size is $N = 50$

the dynamics and this results in a peak in the centre of the domain. However, as u decreases so does the probability mass located in the interior of the domain in the stationary state. Hence as $u \rightarrow 0$, the distribution becomes peaked at the boundaries.

The timescale of the system with mutation is characterized by the so-called ‘mixing time’, t_{mix} , as described in Sect. 2.5. This is the time taken for the probability distribution, $\mathbf{P}(t)$, to come within a specified distance of the stationary distribution $\mathbf{P}^{\text{st}}$, i.e. t_{mix} is the first time at which $d[\mathbf{P}(t_{\text{mix}}), \mathbf{P}^{\text{st}}] = \varepsilon$. The distance between distributions $\mathbf{P}$ and $\mathbf{Q}$ commonly used in this context is $d[\mathbf{P}, \mathbf{Q}] = \sum_{i=0}^N |P_i - Q_i|/2$ with $\varepsilon = 1/2$ [22, 23]. Using our results we can determine if and when there is a correspondence between the mixing time in systems with mutation and the fixation time in systems without mutation. We now describe the dynamics in each of these cases (with and without mutation) separately.

Dynamics Without Mutation

In the system without mutation all realisations reach fixation eventually. If the dynamics is started from state i_0 [i.e. $P_i(t=0) = \delta_{i,i_0}$], the stationary state of the birth–death process is of the form

$$\Phi_{i|i_0} = (1 - \phi_{N|i_0})\delta_{i,0} + \phi_{N|i_0}\delta_{i,N}, \quad (0 \leq i \leq N). \quad (4.52)$$

The quantity $\phi_{N|i_0}$ is the probability that the process reaches the absorbing state N , as described by Eq. (4.34). The probability of being absorbed at state 0 is then $1 - \phi_{N|i_0}$.

Let us now consider the distance of the distribution $\mathbf{P}(t)$ from this distribution,

$$d[\mathbf{P}(t), \Phi] = \frac{1}{2} \left[|P_0(t) - (1 - \phi_{N|i_0})| + \sum_{i=1}^{N-1} P_i(t) + |P_N(t) - \phi_{N|i_0}| \right]. \quad (4.53)$$

Probability continuously flows into the absorbing states, hence $P_0(t) \leq 1 - \phi_{N|i_0}$ and $P_N(t) \leq \phi_{N|i_0}$ for all t . We can therefore simplify the above expression, and we are left with

$$\begin{aligned}
d[\mathbf{P}(t), \Phi] &= \frac{1}{2} \left[(1 - \phi_{N|i_0}) - P_0(t) + \sum_{i=1}^{N-1} P_i(t) + \phi_{N|i_0} - P_N(t) \right] \\
&= \frac{1}{2} \left[1 - P_0(t) - P_N(t) + \sum_{i=1}^{N-1} P_i(t) \right] \\
&= 1 - P_0(t) - P_N(t).
\end{aligned} \tag{4.54}$$

This means that the distance $d(t) = d[\mathbf{P}(t), \Phi]$ is given by the probability that the system has not yet reached fixation in either of the absorbing states by time t . This in turn means that $1 - d(t) = \Pr(t_{\text{fix}} \leq t)$ is the probability to have reached fixation by time t , i.e. $1 - d(t)$ is the cumulative distribution of the unconditional fixation time t_{fix} . Thus the first time when the distance is equal to the mixing time threshold, i.e. $d(t) = \varepsilon$, corresponds to the $(1 - \varepsilon)$ -th percentile of the fixation time distribution. In particular, the first time at which $d(t) = 1/2$ is the median fixation time.

Aside: As $1 - d(t) = \Pr(t_{\text{fix}} \leq t)$ is the cumulative distribution, it follows that $-\dot{d}(t)$ is the probability density function of the unconditional fixation time. With this we can express the mean unconditional fixation time as

$$\begin{aligned}
\langle t_{\text{fix}} \rangle &= \int_0^\infty t [-\dot{d}(t)] dt \\
&= [-td(t)]_0^\infty + \int_0^\infty d(t) dt \\
&= \int_0^\infty d(t) dt.
\end{aligned} \tag{4.55}$$

Thus the mean unconditional fixation time is the area under the curve $d(t)$.

Dynamics With Mutation

In the limit of small mutation rates ($0 < uN \ll 1$), it can be seen from Eq. (4.51) that

$$\frac{P_i^{\text{st}}}{P_0^{\text{st}}} = \frac{b_0}{b_i} \left(\prod_{j=1}^i \frac{b_j}{d_j} \right) = \mathcal{O}(u) \quad \text{for } 1 \leq i \leq N-1, \tag{4.56}$$

and

$$\frac{P_N^{\text{st}}}{P_0^{\text{st}}} = \frac{b_0}{d_N} \left(\prod_{j=1}^{N-1} \frac{b_j}{d_j} \right) = \mathcal{O}(1). \tag{4.57}$$

Together with the normalisation condition ($\sum_{i=0}^N P_i^{\text{st}} = 1$) we can determine that P_0^{st} and P_N^{st} must be $\mathcal{O}(u^0)$, and the remaining probability masses in the interior are $\mathcal{O}(u)$. This effect can be seen in Fig. 4.8. It is verified in Fig. 4.9, where we plot the probability masses at either boundary and the total interior mass in the stationary




Fig. 4.9 Probability masses at either boundary and the total probability mass in the interior of the domain in the stationary state. Panel (a) is the result for the (symmetric) coordination game (Eq. 4.50a). Panel (b) is the result for the Prisoner's dilemma (Eq. 4.50b). As selection is predominantly directed towards the boundary at $i = 0$ in the Prisoner's dilemma (Fig. 4.7b), we have $P_0^{\text{st}} > P_N^{\text{st}}$. For both panels the selection intensity is $\beta = 0.1$ and the system size is $N = 50$

state as a function of u . It is seen that both boundary values converge to a fixed value as $u \rightarrow 0$, whereas the interior mass decreases as a power of u (in fact it scales as u^1).

As the stationary distributions are peaked at the boundaries for small values of u , we can write

$$P_i^{\text{st}} = (1 - \sigma)\delta_{i,0} + \sigma\delta_{i,N} + \mathcal{O}(u), \quad (0 \leq i \leq N) \quad (4.58)$$

where $\sigma = P_N^{\text{st}}/(P_0^{\text{st}} + P_N^{\text{st}}) = \mathcal{O}(u^0)$. It is clear that this distribution is not the same as Φ given in Eq. (4.52): $\mathbf{P}^{\text{st}}$ in systems with $u > 0$ is independent of the initial condition, unlike Φ . Thus there is no obvious connection between fixation times (the approach to Φ) and mixing times (the approach to $\mathbf{P}^{\text{st}}$) in the limit $u \rightarrow 0$.

However, equilibration in many systems with rare mutations is a two-step process; the system first reaches a quasi-stationary distribution that is dependent on the initial condition, before ‘leaking’ on a longer timescale into the final stationary state [23]. As the interior birth and death rates of the systems with and without mutation differ only by corrections of $\mathcal{O}(u)$, we expect the dynamics on a short timescale ($t \ll u^{-1}$) to be essentially the same in both systems. The effects of mutation can only be seen on a longer timescale. We argue that the system initially approaches a distribution close to Φ before reaching its stationary distribution $\mathbf{P}^{\text{st}}$.

Let $\mathbf{P}^{(u=0)}(t) = [P_0^{(u=0)}(t), \dots, P_N^{(u=0)}(t)]^T$ be the probability distribution of the system without mutation. The time evolution is described by the master equation

$$\dot{\mathbf{P}}^{(u=0)} = \mathbb{W} \cdot \mathbf{P}^{(u=0)}, \quad (4.59)$$

where $\mathbb{W}$ is the $(N+1) \times (N+1)$ transition matrix from Eq. (4.1). In fact, this equation is exactly Eq. (4.1). Let $\mathbf{P}^{(u)}(t)$ be the distribution in the system with mutation whose evolution is described by

$$\dot{\mathbf{P}}^{(u)} = (\mathbb{W} + u\mathbb{Q}) \cdot \mathbf{P}^{(u)}, \quad (4.60)$$

where the matrix $\mathbb{Q}$ reflects the difference between the systems with and without mutation. The elements of both matrices $\mathbb{W}$ and $\mathbb{Q}$ are independent of u . Now, let $\mathbf{q}(t) = \mathbf{P}^{(u)}(t) - \mathbf{P}^{(u=0)}(t)$, such that

$$\dot{\mathbf{q}} = \mathbb{W} \cdot \mathbf{q} + u \mathbb{Q} \cdot \mathbf{P}^{(u)}. \quad (4.61)$$

We want to calculate how the separation, $\mathbf{q}$, grows in time given that both systems (with and without mutation) start from the same initial condition [i.e. $\mathbf{q}(0) = \mathbf{0}$]. For this purpose it is convenient to work in the eigenspace of $\mathbb{W}$. As described in Eq. (4.5), $\mathbb{W}$ has two zero eigenvalues, $\mu_0 = \mu_N = 0$, with eigenvectors $v_i^{(0)} = \delta_{i,0}$ and $v_i^{(N)} = \delta_{i,N}$. These are the absorbing states of the system without mutation. The remaining eigenvalues of $\mathbb{W}$, μ_α ($1 \leq \alpha \leq N-1$), are negative (c.f. Sect. 4.3).

Decomposing $\mathbf{q}(t) = \sum_\alpha \tilde{q}_\alpha(t) \mathbf{v}^{(\alpha)}$ into eigendirections $\mathbf{v}^{(\alpha)}$ of $\mathbb{W}$ we have

$$\dot{\tilde{q}}_\alpha = \mu_\alpha \tilde{q}_\alpha + u g_\alpha(t), \quad (4.62)$$

where we have written $\mathbb{Q} \cdot \mathbf{P}^{(u)}(t) = \sum_\alpha g_\alpha(t) \mathbf{v}^{(\alpha)}$ and we note that $g_\alpha(t) = \mathcal{O}(u^0)$. This can be integrated to give

$$\tilde{q}_\alpha(t) = u \int_0^t e^{\mu_\alpha(t-\tau)} g_\alpha(\tau) d\tau. \quad (4.63)$$

On short timescales ($t \ll u^{-1}$) we have $\tilde{\mathbf{q}}(t) = \mathcal{O}(u)$, and hence the separation $\mathbf{q}(t)$ is also $\mathcal{O}(u)$. That is to say in the limit $u \rightarrow 0$, both systems (with and without mutation) initially evolve along the same trajectory. On this timescale both systems approach the distribution Φ . This is shown in Fig. 4.10.

On a longer time scale [$t = \mathcal{O}(u^{-1})$], differences between the two systems become of $\mathcal{O}(u^0)$. However, these differences are concentrated on the states $i = 0$ and $i = N$ as this is where all the probability mass is located. Effectively, a redistribution of probability mass between the boundary states takes place. The distribution of the system with mutation evolves from $(1 - \phi_{N|i_0})\delta_{i,0} + \phi_{N|i_0}\delta_{i,N} + \mathcal{O}(u)$ to $P_i^{\text{st}} = (1 - \sigma)\delta_{i,0} + \sigma\delta_{i,N} + \mathcal{O}(u)$, as shown in Fig. 4.10.

Approximating the stationary distribution of the system with small mutation rate as $P_i^{\text{st}} = (1 - \sigma)\delta_{i,0} + \sigma\delta_{i,N}$ [i.e. neglecting terms $\mathcal{O}(u)$], the distance between the distribution $\mathbf{P}^{(u)}(t)$ at time t and the stationary distribution is

$$d[\mathbf{P}^{(u)}(t), \mathbf{P}^{\text{st}}] \approx \frac{1}{2} \left[\left| P_0^{(u)}(t) - (1 - \sigma) \right| + \sum_{i=1}^{N-1} P_i^{(u)}(t) + \left| P_N^{(u)}(t) - \sigma \right| \right]. \quad (4.64)$$

While $P_0^{(u=0)}(t)$ and $P_N^{(u=0)}(t)$ are monotonically increasing with time in the system without mutation, this is not necessarily the case if there is mutation. This can be seen for $P_N(t)$ in Fig. 4.10. Hence we cannot easily drop the absolute values in Eq. (4.64) as in the case without mutation.




Fig. 4.10 Probability masses as a function of time at either boundary in systems without (*solid lines*) and with (*symbols*) mutation. Without mutation, $P_N(t)$ is given by the integral of Eq. (4.44) and $P_0(t)$ by the integral of Eq. (4.47). These values converge to $\phi_{N|i_0}$ and $\phi_{0|i_0} = 1 - \phi_{N|i_0}$ (*dotted lines*) as $t \rightarrow \infty$. The probability masses in the system with mutation, found from numerical integration of the master equation, initially follow the trajectory of the system without mutation before approaching the stationary state (*dashed lines*). Panel (a) is the result for the (symmetric) coordination game (Eq. 4.50a), in which the probability masses at each boundary in the stationary state are equal (just below 0.5). Panel (b) is the result for the Prisoner's dilemma (Eq. 4.50b). For both panels the selection intensity is $\beta = 0.1$, the system size is $N = 50$, and the initial condition used is $i_0 = 40$. The mutation rate is $u = 10^{-4}$.

Relation Between Dynamics With and Without Mutation

We observe, though, that $P_0^{(u)}(t = 0) = 0$ and $P_N^{(u)}(t = 0) = 0$ for $0 < i_0 < N$. Hence there is an initial phase of the dynamics in which $P_0^{(u)}(t) < 1 - \sigma$ and $P_N^{(u)}(t) < \sigma$. Let us write t^* for the first time at which either $P_0(t^*) = 1 - \sigma$ or $P_N(t^*) = \sigma$ (whichever happens first). In Fig. 4.10a, we can identify this time as $t^* \sim 10^1 - 10^2$, i.e. when $P_N(t) = P_N^{\text{st}}$. In Fig. 4.10b, however, $t^* \sim 10^0 - 10^1$.

Prior to the time t^* we have

$$\begin{aligned} d[\mathbf{P}^{(u)}(t), \mathbf{P}^{\text{st}}] &\approx \frac{1}{2} \left[(1 - \sigma) - P_0^{(u)}(t) + \sum_{i=1}^{N-1} P_i^{(u)}(t) + \sigma - P_N^{(u)}(t) \right] \\ &= 1 - P_0^{(u)}(t) - P_N^{(u)}(t). \end{aligned} \quad (4.65)$$

This is the same [up to $\mathcal{O}(u)$] as the distance to the fixation distribution, Φ , in the system without mutation, given in Eq. (4.54). From this we can conclude that

$$d[\mathbf{P}^{(u)}(t), \mathbf{P}^{\text{st}}] \approx d[\mathbf{P}^{(u)}(t), \Phi] \quad \text{for } t < t^*. \quad (4.66)$$

This is illustrated in Fig. 4.11, where the distributions $\mathbf{P}^{\text{st}}$ and Φ are represented as single points in the (P_0, P_N) plane. The approximate equality (4.66) only holds for points on the trajectory $[P_0^{(u)}(t), P_N^{(u)}(t)]$ which lie inside the shaded region shown in Fig. 4.11. If the fixation distribution, Φ , and the stationary distribution are similar, then a lot of the trajectory $[P_0^{(u)}(t), P_N^{(u)}(t)]$ will be contained in this shaded region, and the equality (4.66) holds for a longer period of time.




Fig. 4.11 Approach to the stationary distribution for the coordination game described by payoff matrix (4.50a). Dots show the trajectory of $[P_0^{(u)}(t), P_N^{(u)}(t)]$ from Fig. 4.10a. The probability quickly approaches the fixation distribution, Φ , before slowly leaking to the stationary state, $\mathbf{P}^{\text{st}}$. For $0 < i_0 < N$, the trajectory starts at $(0, 0)$ and leaves the shaded area at time t^* . For any point inside the shaded area, the distance to the points $(1 - \phi_{N|i_0}, \phi_{N|i_0})$ and $(1 - \sigma, \sigma)$ in our metric (solid lines) are equal. Parameters are as in Fig. 4.10a




Fig. 4.12 Correspondence of mixing time and median fixation time for small mutation rates. Panel (a) shows the unconditional fixation time distribution for the coordination game, constructed from Eqs. (4.44) and (4.47). Panel (b) depicts the mixing time for $u > 0$ and $\epsilon = 1/2$ (from numerical integration of the master equation). Remaining parameters are as in Fig. 4.10a

As we have determined that the systems with and without mutation initially evolve along the same trajectory, we can extend the approximate equality (4.66) to

$$d[\mathbf{P}^{(u)}(t), \mathbf{P}^{\text{st}}] \approx d[\mathbf{P}^{(u)}(t), \Phi] \approx d[\mathbf{P}^{(u=0)}(t), \Phi] \quad \text{for } t < t^*. \quad (4.67)$$

Thus the first time at which $d[\mathbf{P}^{(u)}(t), \mathbf{P}^{\text{st}}] = \epsilon$, provided $t < t^*$, is approximately the $(1 - \epsilon)$ -th percentile of the unconditional fixation time distribution. Choosing $\epsilon = 1/2$, we have the equivalence between the mixing time and the *median* fixation

time. This is illustrated in Fig. 4.12 where we consider the coordination game. Again it is important to stress that this equivalence only holds if t^* satisfies $d[\mathbf{P}^{(u)}(t^*), \mathbf{P}^{\text{st}}] \leq \varepsilon$. It does not hold in the example of the Prisoner’s dilemma shown in Fig. 4.10b, as the distance to the stationary state (or the fixation distribution) when $P_N(t)$ crosses the dashed line representing the stationary distribution is close to one, i.e. larger than ε .

4.7 Efficiency of the Method

Finally we briefly comment on the efficiency of our approach. The forward-only process described in Fig. 4.3 provides a very simple method of sampling from the arrival time distributions. We compare the computational cost of this method with the cost of simulating the original birth–death process using the Gillespie algorithm. Specifically, we measure how long it takes, in absolute ‘wall-clock’ time,² to generate 1,000 samples from the unconditional and conditional arrival time distributions. These results are shown in Fig. 4.13. The time to generate samples using the forward-only process is always shorter than simulating the original process. This difference is particularly noticeable when sampling from the conditional distributions. When using the forward-only process we only need to generate the eigenvalues of the matrix $\mathbb{A}$ and one of the sub-matrices $\mathbb{A}^{(i_0-1)}$ or $\mathbb{A}_{(N-i_0-1)}$ (depending on the arrival state), and we have full control over the arrival state. On the other hand, when simulating the original birth–death process there is no control over the arrival state, and many realisations may arrive at the ‘wrong’ boundary. Thus the computation cost is related to both the fixation time, which is exponentially long in the coexistence game [19], and the fixation probability, which is exponentially small in the coordination game and Prisoner’s dilemma when the mutants have to overcome adverse selection [19].

An alternative approach to measure the computation cost is to compare how long it takes to obtain the arrival time distribution using our spectral method with the time taken to numerically integrate the master equation (4.1). In Fig. 4.14 the benefits of our approach over the conventional methods are clearly seen.

As a final remark, we state that our method is limited by the computational efficiency of evaluating the eigenvalues of the tridiagonal matrix $\mathbb{A}$. Numerical tests have shown that this computation can become unstable when $N \approx 1500$. For larger matrices computational precision limits the calculation of the smallest eigenvalues.

²All samples were generated on the same computer, a 2012 MacBook Air with 1.8 GHz i5 processor running OSX 10.10.4. Software used is Mathematica 9.0.1, and times were measured using the `AbsoluteTiming[]` function.




Fig. 4.13 Times to generate 1,000 samples of **a** the unconditional arrival time, and **b** the arrival time conditioned on reaching state N . Empty symbols/dashed lines show the time taken to generate samples by simulating the forward-only process in Fig. 4.3, including the time to compute all eigenvalues of $-\mathbb{A}$ and its sub-matrices. Filled symbols show the time taken to generate samples through Gillespie simulations of the original birth–death process. The payoff parameters used are the same as in Fig. 4.5. The remaining parameters are $\beta = 0.1$ and $i_0 = N/10$




Fig. 4.14 Times to generate the unconditional arrival time distribution. Empty symbols/dashed lines show the time taken using our spectral method. Filled symbols show the time taken to numerically integrate the master equation (4.1) until the median unconditional fixation time, i.e. up to the time at which $P_0(t) + P_N(t) = 1/2$. Here we use the simple explicit Euler numerical integration method with time-step $\Delta t = 1/N$ [24]. Parameters are as in Fig. 4.13

4.8 Summary

Birth–death processes have received significant attention partly because of their applicability [4, 25], but mainly because of their apparent simplicity. However, a clear understanding of fixation in these processes was lacking, or was lost deep within the literature of probability theory. Existing studies have been limited to investigating the distribution of fixation times in specific birth–death processes, or have simply been limited to finding the mean. As we have shown, however, sometimes the mean does not contain enough information to accurately represent the fixation statistics. In this chapter we have considered a birth–death process with two absorbing states that describes the evolution of a population featuring two types of individual, where

the types may be genotypes or strategies, for example. The probability that, a single type of individual or strategy takes over the population, and the time that this takes, is dependent on the initial condition. We formulated our theory to allow for a general starting point for the state of the population. This is in contrast to some of the probability theory literature, where reflecting boundaries or specific initial conditions are required to obtain results [6, 7].

The main result of this work is the closed-form solutions for the fixation time distributions, along with efficient methods of sampling from these distributions by reducing the birth–death process to a forward-only process in a different space of states. These results are expressed in terms of the spectrum of the original process. To illustrate these results we calculated the distributions of arrival times in a collection of evolutionary games. This showed that, especially when a coexistence of strategies is favoured by selection, the arrival time distributions can be broad and skewed, and vary greatly with the choice of initial condition.

Using these results, we have established a link between the time-to-fixation in the birth–death process with absorbing states, and the time-to-equilibration in systems where rare mutations can reintroduce lost strategies/genotypes. In the limit of small mutation rates, and under specific choices of game, we have demonstrated the equivalence of the so-called mixing time with the median fixation time.

The reduced forward-only processes provide methods to sample from the arrival time distributions, and these have been shown to be much more efficient than simulating the original birth–death process. This is emphasised when sampling from the conditional arrival time distribution as we have complete control of the terminal state, which is in contrast with simulations of the original process.

In this work we have placed existing representations for simpler cases into a wider and more coherent context [10, 13, 14]. We have established that different representations reported in the probability theory literature stem from a single common origin. Nevertheless, there are fundamental open questions. Claims of probabilistic interpretations of Karlin and McGregor’s theorem have been made [11, 12], but in our view this picture is still incomplete. We would argue that a full probabilistic interpretation of the representations in eigenspace is only reached when each time-step of the forward-only process can be constructed directly and uniquely from realisations of the original process alone. Whether or not this is possible is unclear.

References

1. D. Dingli, A. Traulsen, J.M. Pacheco, Stochastic dynamics of hematopoietic tumor stem cells. *Cell Cycle* **6**, 461 (2007)
2. P.M. Altrock, A. Traulsen, F.A. Reed, Stability properties of underdominance in finite subdivided populations. *PLoS Comput. Biol.* **7**, e1002260 (2011)
3. P. Ashcroft, A. Traulsen, T. Galla, When the mean is not enough: calculating fixation time distributions in birth-death processes. *Phys. Rev. E* **92**, 042154 (2015)
4. N. Goel, N. Richter-Dyn, *Stochastic Models in Biology* (Academic Press, New York, 1974)

5. W.J. Ewens, *Mathematical Population Genetics. I. Theoretical Introduction* (Springer, New York, 2004)
6. S. Karlin, J. McGregor, Coincidence properties of birth and death processes. *Pacific J. Math* **9**, 1109 (1959)
7. J. Keilson, *Markov Chain Models: Rarity and Exponentiality* (Springer, New York, 1979)
8. M. Brown, Y.-S. Shao, Identifying coefficients in the spectral representation for first passage time distributions. *Probab. Eng. Inform. Sc.* **1**, 69 (1987)
9. J.A. Fill, Time to stationarity for a continuous-time markov chain. *Probab. Eng. Inform. Sc.* **5**, 61 (1991)
10. M. Brown, Interlacing eigenvalues in time reversible markov chains. *Math. Oper. Res.* **24**, 847 (1999)
11. J.A. Fill, The passage time distribution for a birth-and-death chain: strong stationary duality gives a first stochastic proof. *J. Theoret. Probab.* **22**, 543 (2009)
12. P. Diaconis, L. Miclo, On times to quasi-stationarity for birth and death processes. *J. Theoret. Probab.* **22**, 558 (2009)
13. J.A. Fill, On hitting times and fastest strong stationary times for skip-free and more general chains. *J. Theoret. Probab.* **22**, 587 (2009)
14. L. Miclo, On absorption times and dirichlet eigenvalues. *ESAIM Probab. Stat.* **14**, 117 (2010)
15. Y. Gong, Y.-H. Mao, C. Zhang, Hitting time distributions for denumerable birth and death processes. *J. Theoret. Probab.* **25**, 950 (2012)
16. M. Barrio, A. Leier, T.T. Marquez-Lago, Reduction of chemical reaction networks through delay distributions. *J. Chem. Phys.* **138**, 104114 (2013)
17. A. Leier, M. Barrio, T.T. Marquez-Lago, Exact model reduction with delays: Closed-form distributions and extensions to fully bi-directional monomolecular reactions. *J. R. Soc. Interface* **11**, 20140108 (2014)
18. H. Risken, *The Fokker-Planck Equation* (Springer, Berlin, 1989)
19. T. Antal, I. Scheuring, Fixation of strategies for an evolutionary game in finite populations. *Bull. Math. Biol.* **68**, 1923 (2006)
20. C. Taylor, Y. Iwasa, M.A. Nowak, A symmetry of fixation times in evolutionary dynamics. *J. Theor. Biol.* **243**, 245 (2006)
21. D.T. Gillespie, Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* **81**, 2340 (1977)
22. D.A. Levin, Y. Peres, E.L. Wilmer, *Markov Chains and Mixing Times* (AMS Publishing, Providence RI, 2009)
23. A.J. Black, A. Traulsen, T. Galla, Mixing times in evolutionary games. *Phys. Rev. Lett.* **109**, 028101 (2012)
24. W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, Cambridge UK, 2007)
25. M.A. Nowak, *Evolutionary Dynamics* (Harvard University Press, Cambridge MA, 2006)

Chapter 5

Metastable States in a Model of Cancer Initiation

5.1 Introduction

We now turn our attention to a more applied subject, the accumulation of mutations and the initiation of cancer. With an ageing population, the prevalence of this genetic disease in the UK (and across developed countries) has sky-rocketed, with one in two people expected to be diagnosed with cancer in their lifetime [1]. This issue is close to the hearts of the British population. Cancer Research UK, the UK's leading healthcare charity, received donations in excess of £500m in the last financial year (2014–2015), 80 % of which was used to fund research [2]. Although the majority of research is clinical or experimental, theoretical approaches greatly contribute to our understanding of this malady.

The initiation and progression of cancer is a result of the accumulation of genetic alterations [3]. The dynamics of mutation acquisition is governed by evolutionary parameters such as the rate at which alterations arise, the selection effect that these alterations confer to cells, and the size of the population of cells that proliferate within a tissue. Therefore these processes are amenable to mathematical investigation, and much effort has been devoted to modelling these systems and analysing the rates at which mutations arise within pre-cancerous tissues [4–16].

Models of the initiation and progression of cancer vary dramatically in their complexity and tractability. At the most complex end of the scale, mathematical models consider explicit tissue structure and mechanics [17, 18], as well as resource competition and the creation of 'public goods' [19]. Simpler spatial models have been designed to replicate some tissue structures, such as the linear process for describing the accumulation of mutations in a colorectal crypt [20]. In terms of the dynamics of mutation acquisition, the effects of genetic instabilities [21] or the hierarchical organisation of cells within the population [22] have been considered.

At the more tractable end of the scale of cancer models are the well-mixed, birth and death representations of mutation acquisition. For tumour progression, modelling a growing population is crucial as, by definition, cancerous phenotypes

in the population grow in an uncontrolled manner [16]. For this reason branching processes have been used to describe the progression of cancer [9, 23]. During the initiation of cancer, however, the number of cells in pre-cancerous tissues fluctuates by only a small amount. One can then make the simplifying assumption that the population size is constant. This is the approach often used to describe the inactivation of tumour suppressor genes (TSG) [24, 25], which directly regulate the growth and differentiation pathways of cells [3]. This simple fixed-size model will be the subject of our analysis. The work in this chapter has been published in Ref. [26].

Although the birth–death processes that we have investigated in the previous Chapters capture a wide variety of applications, it is often the case that the population is described by more than one variable. Such scenarios occur when there are multiple (more than two) types of interacting individuals. This is evident in recent modelling works of mutation acquisition [21, 24, 25, 27–30], which have focused on the accumulation of two mutations in a fixed-size population of cells, such that there are three cell types present in the population. These studies have revealed a more detailed picture of the initiation of cancer; a homogeneous population harbouring no mutations can move to a homogeneous state in which all cells carry two mutations without ever visiting a homogeneous state in which all cells harbour just one mutation. This phenomenon is referred to as ‘stochastic tunnelling’, and is illustrated in Fig. 5.1.

The term ‘stochastic tunnelling’ is, however, an unfortunate use of vocabulary. The ‘stochastic’ part indicates that this route to fixation does not occur with certainty. As indicated in Fig. 5.1 there are two routes to the homogeneous state with two mutations. The sequential route is still available to the system, but it becomes less likely in certain parameter regimes. As will be shown below, the ‘stochastic’ route is actually predicted by the deterministic dynamics, adding further confusion to the terminology. The process we refer to as ‘tunnelling’ is not limited to the scenario




Fig. 5.1 The population can reach the all-type-2 state via two routes. The first is the sequential fixation route in which the first mutation takes over the population, and where this is then followed by the second mutation. The second route does not visit the all-type-1 state. This is the stochastic tunnelling route. The *arrows* are related to the change of state only, and imply nothing about the fitness of these states

of ‘going under’ a potential barrier as described by quantum-mechanical tunnelling. Here ‘tunnelling’ refers only to overlapping transitions between the homogeneous states.

The fixation of mutations is an important phenomenon if the cells harbouring mutations are deleterious, i.e. less fit than the wild-type un-mutated cells. If cells are less fit they will generally have low concentrations within the tissue. Then the chance of a cancerous phenotype emerging (further mutation) is very low. Demographic fluctuations can drive the disadvantageous cells to higher numbers, but these states are short-lived. If these mutations become fixated in the tissue, however, then the state is maintained until a further mutation occurs and the chance of a cancerous phenotype emerging is much greater.

The existing analytical approaches provide accurate approximations for the time to fixation of cells harbouring two mutations in a subset of the parameter space; there are extensive regions which, up to date, have been left unexplored. These are predominantly situations in which the cell harbouring two mutations is not the most advantageous in the sequence. Before the double mutant reaches fixation, the population has to travel across a fitness hill or move constantly downhill in fitness space. The dynamics can then become trapped in quasi-equilibria, or metastable states, which are a consequence of the adverse selection being balanced by the effect of mutation [31]. These long-lived equilibria support a heterogeneous population, and fixation is driven purely by demographic fluctuations. We address this regime based on ideas from mathematical physics. Specifically we employ the Wentzel–Kramers–Brillouin (WKB) method to derive quantitative predictions for fixation times. The WKB method has been used to describe the escape from metastable states in a variety of systems, including predator–prey and epidemic models [32–37]. It has also been used in models of mutation acquisition to describe the escape over a ‘recombination barrier’ in a sexually reproducing population [38], and in a model of Muller’s ratchet¹ [40].

In this chapter we will investigate a microscopic model that describes the accumulation and spread of two mutations in a population of cells. This model is described in Sect. 5.2, along with a brief summary of the previous analytical approaches used to investigate this model. We then investigate the underlying deterministic dynamics of this model in Sect. 5.3, where we compute the locations and stability properties of fixed points. In Sect. 5.4 we classify the stochastic behaviours that can occur. These are, to some extent, determined by the previous deterministic analysis. By employing the WKB method in Sect. 5.5, we obtain analytical and semi-analytical approximations for fixation times in parameter regimes which could not be captured by previous methods, i.e. regimes in which metastable states exist. We then discuss these results, and the implications they have on our understanding of stochastic tunnelling, the accumulation of mutations, and the initiation of cancer.

¹The process of repeatedly accumulating disadvantageous mutations [39], named after Hermann Joseph Muller (1890–1967).

5.2 Model

We consider the same microscopic model that has been used to describe the inactivation of a tumour suppressor gene in a pre-cancerous tissue, as presented in Ref. [24]. We first describe the microscopic model, and then comment on the approach of the previous analyses.

Microscopic Model

We consider a well-mixed, finite population of N cells. Each cell can be of one of three possible types:

- (i) Type 0—a wild-type cell harbouring no mutations;
- (ii) Type 1—a cell harbouring one mutation;
- (iii) Type 2—a cell harbouring two mutations.

Initially, all cells are of type 0, which we refer to as the all-wild-type state. The evolution of the population is determined by a Moran process in continuous time [41]. During each elementary time-step of this process, a cell is randomly chosen to reproduce proportional to its fitness. In the same time-step a cell is randomly removed, such that the total population size remains constant. The daughter cell can either inherit its type from the parent, or acquire a single mutation during reproduction. The relative fitness values of type-0, type-1 and type-2 cells are denoted by r_0 , r_1 and r_2 , respectively. These fitnesses can be thought of as basic reproductive rates. Without loss of generality, we use $r_0 = 1$ throughout, such that all fitness values are relative to the wild-type. This is equivalent to rescaling time such that it is measured in generations of the wild-type cells. The mutation rates u_1 and u_2 denote the probability that the daughter of a type-0 cell is of type 1, and the probability that the daughter of a type-1 cell is of type 2, respectively. We neglect all other combinations of mutations, i.e. back mutations and multiple mutations during a single reproduction event are not possible. The assumption of no back-mutation is commonly used in the population genetics literature [42]. It means that a cell with i mutations cannot produce an offspring with less than i mutations. The assumption is justifiable since the human genome is very large, $\sim 3 \times 10^9$ base pairs, and the probability of mutating a specific base per cell division is very small, $\sim 10^{-10}$ – 10^{-11} [43]. Therefore the chance of undoing a specific point mutation is vanishingly small. The probability that a second critical alteration occurs at a different locus is much higher.

In our model finite populations will eventually reach a state in which all cells have acquired two mutations. This state is ‘absorbing’, i.e. once this state has been reached, no further dynamics can occur. There are of course physical processes beyond the second mutation. In pre-cancerous tissues for example, there will be a finite probability that cells progress from this state to accumulate further changes. These processes are not the focus of our work though, and so are not included in the model.

Let us denote the number of type-0, type-1, and type-2 cells by n_0 , n_1 and n_2 , respectively. The total population, $N = n_1 + n_2 + n_3$, is constant. We label the

transition rates for the Moran process as $T^{i \rightarrow j}$, which is the rate per unit time that a cell of type i is replaced by a cell of type j . In a process labelled $T^{0 \rightarrow 1}$, for example, the state of the population changes from (n_0, n_1, n_2) to state $(n_0 - 1, n_1 + 1, n_2)$. There are six possible changes to the state, labelled by the stoichiometry coefficient $\mathbf{v}$. The set of transition rates are

$$T^{0 \rightarrow 1} = \frac{n_0}{N} \times \frac{u_1 r_0 n_0 + (1 - u_2) r_1 n_1}{\bar{r}}, \quad \mathbf{v} = (-1, +1, 0), \quad (5.1a)$$

$$T^{0 \rightarrow 2} = \frac{n_0}{N} \times \frac{u_2 r_1 n_1 + r_2 n_2}{\bar{r}}, \quad \mathbf{v} = (-1, 0, +1), \quad (5.1b)$$

$$T^{1 \rightarrow 0} = \frac{n_1}{N} \times \frac{(1 - u_1) r_0 n_0}{\bar{r}}, \quad \mathbf{v} = (+1, -1, 0), \quad (5.1c)$$

$$T^{1 \rightarrow 2} = \frac{n_1}{N} \times \frac{u_2 r_1 n_1 + r_2 n_2}{\bar{r}}, \quad \mathbf{v} = (0, -1, +1), \quad (5.1d)$$

$$T^{2 \rightarrow 0} = \frac{n_2}{N} \times \frac{(1 - u_1) r_0 n_0}{\bar{r}}, \quad \mathbf{v} = (+1, 0, -1), \quad (5.1e)$$

$$T^{2 \rightarrow 1} = \frac{n_2}{N} \times \frac{u_1 r_0 n_0 + (1 - u_2) r_1 n_1}{\bar{r}}, \quad \mathbf{v} = (0, +1, -1). \quad (5.1f)$$

The quantity $\bar{r} = (r_0 n_0 + r_1 n_1 + r_2 n_2)/N$ is the average fitness of the population. As an example, the first reaction rate, $T^{0 \rightarrow 1}$, in Eq. (5.1) can be broken down as follows: a type-0 cell is randomly chosen to be removed with probability n_0/N . Meanwhile, either a type-0 cell is chosen to reproduce at rate $r_0 n_0/\bar{r}$, and a mutation occurs (u_1) during the reproduction to produce a type-1 daughter cell, or a type-1 cell is chosen to reproduce with rate $r_1 n_1/\bar{r}$ without a mutation ($1 - u_2$). The rates for the other processes can be interpreted analogously. We choose a continuous-time setup, and correspondingly all rates in Eq. (5.1) scale linearly in the population size N . Simulations are carried out using a standard Gillespie algorithm as described in Sect. 2.9, and times are measured in generations of the wild-type cell.

The probability, $P_{\mathbf{n}}(t)$, that the population is in state $\mathbf{n} = (n_0, n_1, n_2)$ at time t is described by the master equation

$$\dot{P}_{\mathbf{n}}(t) = \sum_{\mathbf{v}} [T_{\mathbf{n}-\mathbf{v}}^{\mathbf{v}} P_{\mathbf{n}-\mathbf{v}}(t) - T_{\mathbf{n}}^{\mathbf{v}} P_{\mathbf{n}}(t)], \quad (5.2)$$

where $T_{\mathbf{n}}^{\mathbf{v}}$ is the transition rate from Eq. (5.1) that has the corresponding stoichiometric coefficient $\mathbf{v}$, evaluated when the population is in state $\mathbf{n}$.

Previous Analyses

The majority of analytical investigations of mutation acquisition and stochastic tunnelling [21, 24, 25, 27–30] have been limited to considering transitions between homogeneous states of the population, as illustrated in Fig. 5.1. If we let $\Pi_i(t)$ be the probability to be found in the state in which all cells in the population harbour i mutations ($i = 0, 1, 2$), and $R_{i \rightarrow j}$ ($j > i$) is the rate at which we move from the

homogeneous state with i mutations to the homogeneous state with j mutations, then we can write the three-state master equation as

$$\dot{\Pi}_0 = -R_{0 \rightarrow 1}\Pi_0 - R_{0 \rightarrow 2}\Pi_0, \quad (5.3a)$$

$$\dot{\Pi}_1 = R_{0 \rightarrow 1}\Pi_0 - R_{1 \rightarrow 2}\Pi_1, \quad (5.3b)$$

$$\dot{\Pi}_2 = R_{0 \rightarrow 2}\Pi_0 + R_{1 \rightarrow 2}\Pi_1. \quad (5.3c)$$

Here $R_{0 \rightarrow 2}$ is the so-called tunnelling rate, and the quantity $\dot{\Pi}_2(t)$ describes the probability density of arrival times at the homogeneous state with two mutations. Although Eq. (5.3) is easy to solve, the form of the rates $R_{i \rightarrow j}$ as a function of the underlying model parameters is non-trivial. It is finding these rates, in particular $R_{0 \rightarrow 2}$, that has been the focus of the above mentioned recent investigations.

In Ref. [24], the transition rates in Eq. (5.3) are computed as

$$R_{0 \rightarrow 1} = u_1 N \frac{1 - r_0/r_1}{1 - (r_0/r_1)^N}, \quad (5.4a)$$

$$R_{1 \rightarrow 2} = u_2 N \frac{1 - r_1/r_2}{1 - (r_1/r_2)^N}, \quad (5.4b)$$

$$R_{0 \rightarrow 2} = u_1 N \left[\frac{-(r_0 - r_1) + \sqrt{(r_0 - r_1)^2 + 2u_2 r_1 (r_0 + r_1) \frac{1 - r_0/r_2}{1 - (r_0/r_2)^N}}}{r_0 + r_1} \right]_+, \quad (5.4c)$$

where the notation $[\cdot]_+$ indicates this function cannot have a negative value, i.e. $[x]_+ = 0$ if $x \leq 0$ and $[x]_+ = x$ if $x > 0$. The first rate can be interpreted as follows: an entirely wild-type population produces a single mutant with rate $u_1 N$. Ignoring any mutational effects, this mutant takes over the population with probability $(1 - r_0/r_1)/[1 - (r_0/r_1)^N]$, which can be computed from Eq. (2.24). The rate $R_{1 \rightarrow 2}$ can be interpreted analogously. The tunnelling rate, $R_{0 \rightarrow 2}$, is constructed by considering the probability that an independent lineage of type-1 cells (which emerges with rate $u_1 N$) produces a further mutation which takes over the population. The predictions for the mean fixation time from this formulation are in very good agreement with simulation results when r_2 is large, or when the system size is small such that the tunnelling probability is low [24]. We will show later in this chapter that this approach does not accurately capture fixation times when the type-2 cells are deleterious and N is large.

5.3 Deterministic Analysis

As a starting point, we will first consider the deterministic limit of Eq. (5.2). This approach does not capture any of the stochastic effects. However, the types of stochastic trajectories that can be observed for different parameter sets are, to some

extent, set by the underlying deterministic dynamics. This approach will highlight the expected qualitative behaviours of our model.

We have derived the deterministic equations for a master equation of the form (5.2) in Sect. 2.6, i.e. Eq. (2.65). Writing $x_i = \lim_{N \rightarrow \infty} \langle n_i \rangle / N$, we have the relation $x_0 + x_1 + x_2 = 1$, and the average fitness is given by $\bar{r} = r_0x_0 + r_1x_1 + r_2x_2$. Inserting the transition rates (5.1) into Eq. 2.65, the deterministic equations governing the dynamics of the population can be written as

$$\bar{r}\dot{x}_0 = [(1 - u_1)r_0 - \bar{r}]x_0, \quad (5.5a)$$

$$\bar{r}\dot{x}_1 = u_1r_0x_0 + [(1 - u_2)r_1 - \bar{r}]x_1, \quad (5.5b)$$

$$\bar{r}\dot{x}_2 = u_2r_1x_1 + (r_2 - \bar{r})x_2. \quad (5.5c)$$

Given the relation $x_0 + x_1 + x_2 = 1$, there are only two independent degrees of freedom. Thus the state-space can be projected onto the two-dimensional plane, which we will refer to as the concentration simplex. This is illustrated in Fig. 5.2. Each point in the simplex represents one particular state $(1 - x_1 - x_2, x_1, x_2)$ of the population. At points in the interior of the simplex all three types of cells are present in the population ($x_i > 0$ for $i = 0, 1, 2$). Points on the edges of the simplex represent states in which one of the three types is not present, for example $x_0 = 0$ for points along the edge connecting the lower-right corner of the simplex with the upper corner. We will refer to this as the 1–2 edge in the following, and similarly for




Fig. 5.2 The concentration simplex, or barycentric coordinate system. The *bottom-left* corner corresponds to $x_0 = 1$. Labels along the *left* edge indicate the concentration of type-0 cells. The *bottom-right* corner corresponds to $x_1 = 1$ and the labels on the *bottom* edge indicate the concentration of type-1 cells. The *top* corner corresponds to $x_2 = 1$ and labels on the *right* edge indicate the concentration of type-2 cells. For the *point* shown, the population consists of 20 % type-0 cells, 50 % type-1 cells, and 30 % type-2 cells

the other edges. The three corners of the simplex represent the homogeneous states, i.e. $x_0 = 1$ (lower left corner), $x_1 = 1$ (lower right) and $x_2 = 1$ (upper corner).

Stability Analysis

The deterministic equations (5.5) can permit fixed points, that is a point $\mathbf{x}^*$ at which $\dot{x}_i = 0$ for all $i = 0, 1, 2$. The most obvious of these is at $(0, 0, 1)$, which corresponds to the absorbing state. The equations can have a further zero, one, or two non-trivial fixed points, depending on the values of the fitness parameters and the mutation rates. These additional fixed points correspond to points at which the ‘push’ of mutation balances the adverse effect of selection [28, 31]. The stability of the fixed points, as described in Sect. 2.6, is determined by the eigenvalues of the Jacobian of Eqs. (5.5). We define the two-dimensional Jacobian in terms of x_1 and x_2 , and impose $x_0 = 1 - x_1 - x_2$. The Jacobian is given by

$$\mathbb{J}(x_1, x_2) = \begin{pmatrix} \partial_1 \dot{x}_1 & \partial_2 \dot{x}_1 \\ \partial_1 \dot{x}_2 & \partial_2 \dot{x}_2 \end{pmatrix}, \quad (5.6)$$

where $\partial_i f = \partial f / \partial x_i$.

Along the 1–2 boundary of the concentration simplex ($x_0 = 0$) we have $x_2 = 1 - x_1$. A fixed point of Eqs. (5.5) can then be found at

$$x_1^* = \frac{(1 - u_2)r_1 - r_2}{r_1 - r_2}, \quad (5.7a)$$

$$x_2^* = \frac{u_2 r_1}{r_1 - r_2}. \quad (5.7b)$$

The parameter range in which this fixed point exists is determined by the condition $0 < x_1^* < 1$, which we can write as $(1 - u_2)r_1 > r_2$. The fixed point on the 1–2 edge therefore exists when type-1 cells have a fitness advantage over type-2 cells. The factor $1 - u_2$ accounts for the effect of the mutation ‘push’. Increasing this fitness advantage moves the fixed point towards $x_1 = 1$, or equivalently away from the absorbing state at $x_1 = 0$. For vanishing mutation rate u_2 , the fixed point approaches the $x_1 = 1$ state. Evaluating the Jacobian in Eq. (5.6) at the fixed point (5.7), we find that the determinant and trace satisfy

$$\det[\mathbb{J}(x_1^*, 1 - x_1^*)] = \frac{[(1 - u_2)r_1 - r_2][(1 - u_2)r_1 - (1 - u_1)r_0]}{[(1 - u_2)r_1]^2}, \quad (5.8a)$$

$$\text{Tr}[\mathbb{J}(x_1^*, 1 - x_1^*)] = \frac{(1 - u_1)r_0 + r_2 - 2(1 - u_2)r_1}{(1 - u_2)r_1}, \quad (5.8b)$$

and from this the discriminant is always positive such that the eigenvalues of $\mathbb{J}$ are real. Using the classification scheme in Fig. 2.4, we can characterise this boundary fixed point in the following way:

$$\left\{ \begin{array}{l} \text{Fixed point exists if } (1 - u_2)r_1 > r_2, \\ \text{and is stable if } (1 - u_1)r_0 < (1 - u_2)r_1, \\ \text{or is a saddle point if } (1 - u_1)r_0 > (1 - u_2)r_1. \end{array} \right\} \quad (5.9)$$

These inequalities contain regions I and II of the phase diagram shown in Fig. 5.3. The stable fixed point can be seen in Fig. 5.4a, and the saddle point can be seen in Fig. 5.4b.

A further fixed point of Eqs. (5.5), this time with $x_0^* > 0$, can be found at

$$x_1^* = \frac{[(1 - u_1)r_0 - r_2]u_1r_0}{u_2r_1(r_0 - r_2) + (r_0 - r_1)[(1 - u_1)r_0 - r_2]}, \quad (5.10a)$$

$$x_2^* = \frac{u_1u_2r_0r_1}{u_2r_1(r_0 - r_2) + (r_0 - r_1)[(1 - u_1)r_0 - r_2]}, \quad (5.10b)$$

which exists if the model parameters satisfy $(1 - u_1)r_0 > (1 - u_2)r_1$ and $(1 - u_1)r_0 > r_2$. If this is the case we find $\det[\mathbb{J}] > 0$ and $\text{Tr}[\mathbb{J}] < 0$, and the discriminant is always positive. Hence, according to Fig. 2.4 this fixed point is always stable. This is the case in regions II and III in Fig. 5.3, i.e. when cells harbouring one and two mutations are less fit than the wild-type. The fixed point can be seen in Fig. 5.4b, c. The fixed point moves closer to $x_0 = 1$ when the fitness advantage of the wild-type cells is increased (i.e. by lowering the fitness of type-1 and type-2 cells). Decreasing the mutation rates also moves the fixed point closer to $x_0 = 1$.




Fig. 5.3 Boundary fixed points occur when $(1 - u_2)r_1 > r_2$ (regions I and II; mutation–selection balance between types 1 and 2). Stable interior fixed points occur when $(1 - u_1)r_0 > (1 - u_2)r_1$ and $(1 - u_1)r_0 > r_2$ (regions II and III; mutation–selection balance between all three types). No additional fixed points are found in regions IV and V (beneficial type-2 mutation). In all regions, the point $x_2 = 1$ is an absorbing state and is therefore a fixed point as well




Fig. 5.4 The streamlines of Eqs. (5.5) in parameter regions which permit non-trivial fixed points, which are regions I, II and III of Fig. 5.3. The *thick line* is the deterministic flow from the all-wild-type initial condition. *Solid circles* indicate stable fixed points, and the *open circle* for region II corresponds to the saddle point that is stable along the 1–2 boundary. The fitness landscape that generates each type of behaviour is shown below each simplex. Region II corresponds to the case of Muller’s ratchet, where increasingly deleterious mutations are accumulated [39]. In these illustrations the mutation rates are $u_1 = u_2 = 10^{-2}$. **a** Region I. **b** Region II. **c** Region III




Fig. 5.5 The streamlines of Eqs. (5.5) in parameter regions where $x_2 = 1$ is the only fixed point, which are regions VI and V of Fig. 5.3. The *thick line* is the deterministic flow from the all-wild-type initial condition. The fitness landscape that generates each type of behaviour is shown below each simplex. Region IV is the classic valley crossing scenario [38, 44–47]. In these illustrations the mutation rates are $u_1 = u_2 = 10^{-2}$. **a** Region IV. **b** Region V

In the remaining regions of parameter space, namely regions IV and V in Fig. 5.3, there are no non-trivial fixed points. As seen in Fig. 5.5a, b, the flow from the all-wild-type state is directly towards the all-type-2 state. Returning to the discussion of

stochastic tunneling, we can see that the deterministic trajectories in Fig. 5.5a, b do not pass through the all-type-1 state. Hence the deterministic dynamics predicts the stochastic tunnelling route.

5.4 Types of Stochastic Behaviour

The deterministic equations (5.5) can provide some intuition into the qualitative stochastic behaviours that can occur. We now analyse this behaviour in each region of parameter space outlined in Fig. 5.3.

Region I

In region I, the deterministic dynamics flows towards the stable fixed point on the 1–2 edge of the concentration simplex ($x_0 = 0$). The type-0 cells have the lowest fitness, and are deterministically lost by selection. The fixed point is a consequence of the mutation–selection balance between type-1 and type-2 cells [31]; selection acts to reduce the number of type-1 cells, but mutation from cells of type 1 acts to increase their number. Writing $r_2 = (1 - s)r_1$, the existence condition for the equilibrium, $(1 - u_2)r_1 > r_2$, reduces to $u_2 < s$. It is a well-known result from population genetics that this condition prevents the deterministic loss of the type-1 cells [31, 48]. The deterministic system gets stuck at this fixed point, but a finite population will eventually reach the all-type-2 state.

At large but finite population sizes, the stochastic dynamics are expected to approximately follow the deterministic path such that type-0 cells quickly become extinct. This is shown in Fig. 5.6. The lack of backwards mutations means the population cannot depart from the 1–2 edge and the problem reduces to one degree




Fig. 5.6 **a** Thin line shows a single stochastic trajectory in the concentration simplex for parameters in region I. The trajectory closely follows the deterministic trajectory and type-0 cells become extinct. **b** The individual components of the same trajectory as a function of time. Type-0 cells reach extinction very quickly, and the population composition fluctuates about the fixed point, indicated by dashed lines. Eventually type-1 cells are lost and type-2 cells reach fixation. A moving average over a time window of 200 generations has been taken to improve clarity. Parameters are as in Fig. 5.4a, i.e. $r_0 = 1.00$, $r_1 = 1.05$, $r_2 = 1.00$, $u_1 = u_2 = 10^{-2}$, and $N = 300$

of freedom. The mutation–selection balance maintains the heterogeneous population state of type-1 and type-2 cells, which fluctuates about the fixed point location as shown in Fig. 5.6b. The intrinsic noise then has to drive the system from this metastable state into the absorbing all-type-2 state against the direction of selection. Fixation times are expected to grow exponentially with the population size in-line with existing results [38, 49, 50].

Region II

In region II, the deterministic flow from the all-wild-type state is towards a stable fixed point in the interior of the simplex. This point corresponds to the mutation–selection balance point of all three cell types. There is a second fixed point located on the 1–2 edge, which corresponds to mutation–selection balance between types 1 and 2 in the absence of type-0 cells (analogous to region I). As type-0 cells have the highest fitness in this regime, selection is directed away from the 1–2 edge. Thus the fixed point on this edge is a saddle.

As before the stochastic dynamics in finite populations will reach the all-type-2 state eventually. The population will closely follow the deterministic trajectory before reaching the metastable state about the stable interior fixed point. This can be seen in Fig. 5.7a. Here the mutation–selection balance maintains the heterogeneous state with all three species present. The population fluctuates about this fixed point as shown in Fig. 5.7b until it eventually overcomes the adverse selection and escapes. There are two possibilities for the subsequent behaviour:

- (i) Type-0 cells become extinct and the population settles into the metastable state on the 1–2 edge. Intrinsic fluctuations enable the population to overcome the




Fig. 5.7 **a** Thin line shows a single stochastic trajectory in the concentration simplex for parameters in region II. The population fluctuates about the interior stable fixed point before type-0 cells become extinct. **b** The individual components of the same trajectory as a function of time. The population quickly reaches the metastable state and fluctuates about the fixed point location, indicated by dashed lines. Eventually type-0 cells are lost at $t \approx 5500$. After this the population of type-1 and type-2 cells fluctuates about the boundary fixed point location, before type-1 cells are lost and type-2 cells reach fixation. A moving average over a time window of 200 generations has been taken to improve clarity. Parameters are as in Fig. 5.4b, i.e. $r_0 = 1.00$, $r_1 = 0.98$, $r_2 = 0.95$, $u_1 = u_2 = 10^{-2}$, and $N = 300$

adverse selection along the edge and eventually reach the absorbing all-type-2 state. This corresponds to sequential extinction, first of type-0 cells, then of type-1 cells. This process is equivalent to a minimal model of Muller's ratchet [39], in which the most advantageous phenotypes are sequentially lost. A trajectory of this type is illustrated in Fig. 5.7b.

- (ii) Cells of type 0 and type 1 can, in principle, go extinct (almost) simultaneously. The trajectory of the system then hits the 1–2 edge close to the all-type-2 corner of the simplex. It does not become trapped in the metastable state located on the 1–2 edge.

We observe that this second path to extinction is realised only very rarely, which is in agreement with similar studies of predator–prey dynamics [32]. Hence our further analysis will only consider the sequential fixation path.

Region III

In region III the deterministic dynamics has a single stable fixed point in the interior of the concentration simplex. This point again corresponds to the mutation–selection balance point of all three cell types. Large, but finite populations will behave as discussed in case (ii) for region II: They will initially become trapped in the metastable state about the mutation–selection balance point, before intrinsic fluctuations eventually drive the system to the absorbing all-2 state. In region III, type-0 and type-1 cells go extinct at essentially the same time. The type-0 cells can reach extinction first, and then type-1 cells quickly follow as selection along the 1–2 edge is directed towards the absorbing state $[(1 - u_2)r_1 < r_2]$. This is illustrated in Fig. 5.8.




Fig. 5.8 **a** Thin line shows a single stochastic trajectory in the concentration simplex for parameters in region III. The population fluctuates about the interior stable fixed point before type-2 cells eventually fixate. **b** The individual components of the same trajectory as a function of time. The population quickly reaches the metastable state and fluctuates about the fixed point location, indicated by dashed lines. Eventually type-0 and type-1 cells are lost almost simultaneously. A moving average over a time window of 200 generations has been taken to improve clarity. Parameters are as in Fig. 5.4c, i.e. $r_0 = 1.00$, $r_1 = 0.95$, $r_2 = 0.98$, $u_1 = u_2 = 10^{-2}$, and $N = 300$

Regions IV and V

In a subset of the parameter space, shown as regions IV and V in Fig. 5.3, the deterministic flow from the all-wild-type state is directly to the absorbing all-2 state, as shown in Fig. 5.5. For such model parameters we expect that fixation in finite populations will be quick (relative to fixation times in regions with metastable states) as type-2 cells are favoured by selection (and mutation). These scenarios agree with the theory of natural selection, in which the populations fitness increases over time [51]. In region IV this is achieved by crossing a fitness valley, and in region V it is achieved by sequentially selecting the most advantageous phenotype.

Summary of Previous Literature

Figure 5.9 illustrates in which parameter regimes fixation has previously been studied in the stochastic tunnelling literature. These existing studies almost exclusively focus on regions IV and V, i.e. cases in which fixation is driven not primarily by demographic noise, but by the underlying deterministic flow. As mentioned above fixation is typically fast in regions IV and V. Based on similar studies in evolutionary game theory one would expect the fixation time to grow logarithmically with the population size, $\tau \sim \ln N$ [49]. The regions containing non-trivial fixed points are largely unexplored by previous investigations. Fixation is controlled by stochastic effects so that fixation times are large and broadly distributed. As we will discuss below, fixation times grow exponentially with the population size in such cases. This




Fig. 5.9 The coloured area roughly corresponds to the regions in which the probability of fixation at a given time has been successfully predicted in Ref. [30]. The southwest–northeast striped region, with r_1 deleterious or slightly advantageous, and r_2 very advantageous, is the approximate region of interest of Refs. [21, 25, 27]. These studies focused on the time to emergence of a single type-2 cell. The northwest–southeast striped region, with r_1 neutral or deleterious, and r_2 advantageous, is approximately the region of interest of Refs. [24, 29]. These studies were concerned with computing fixation times of the advantageous type-2 cells and rely on the assumption that the number of type-1 cells is small. Finally, the horizontal striped region approximately corresponds to the literature of fitness valley crossings, notably Refs. [38, 44–47]. These studies are concerned with $1/N < r_2 - r_0 \ll 1$

is perfectly in-line with the findings of Ref. [30], who point out that fixation in these regions takes a very long time. Efficient measurements of fixation time in simulations are hence difficult. Methods which require the numerical solutions of, for example, the backward Fokker-Planck equation or a backward master equation reach their limits here as well [30]. This is because forward integration of these equations has to be carried out for increasingly long periods of time. The contribution of this work is to analyse precisely these previously inaccessible cases.

5.5 WKB Analysis

In this section we compute the fixation properties of systems in which the underlying deterministic dynamics has one or more attracting fixed points away from the absorbing all-type-2 state. For this analysis we use the celebrated WKB method. We here present this method as an ‘off-the-shelf approach’, using existing WKB studies to guide our analysis. In Chap. 6 we further discuss the WKB method, and in particular its mathematical origins and relation to other fields, such as large deviations theory.

To proceed with the analysis of our model we need to make the following simplifying assumptions, which are justified by the previous deterministic analysis:

- (i) We assume that the population first settles into a distribution about the fixed point, as shown in Figs. 5.6, 5.7 and 5.8.
- (ii) We assume that the population will ‘leak’ into the absorbing state on a very long timescale from this distribution. With this assumption we can also say that the time taken for the population to reach the metastable state is negligible when compared to the escape time. This is most clearly seen in Fig. 5.6b.

With these assumptions we can compute the distribution about the fixed points, and the escape rate from these states, from the master equation (5.2). These assumptions (and hence the subsequent analysis) are only valid when the selective pressure is greater than the effect of noise, such that the metastable states are long-lived. For this reason, the approach described here is only valid for large values of N which satisfy this condition.²

Mathematically we formulate the problem as follows: let $\pi_{\mathbf{n}}$ be the quasi-stationary distribution (QSD), which is independent of time. This is the distribution that, prior to the system reaching the absorbing state $\mathbf{n}_{\text{abs}}$, we would observe for the state of the population. This distribution satisfies $\sum_{\mathbf{n} \neq \mathbf{n}_{\text{abs}}} \pi_{\mathbf{n}} = 1$. We expect that it is peaked about the stable fixed points of the underlying deterministic dynamics. The mean time taken to escape from this metastable state, τ , is much greater than the time taken to initially reach the metastable state t_r , i.e. $\tau \gg t_r$. Provided this condition holds, we can assume that after a short time the probability to find the population in state $\mathbf{n}$ is given by

²The minimum value of N for which our analysis is valid is dependent on the remaining model parameters, but comparisons with simulation results in the next section show it is accurate for $N \gtrsim 100$.

$$P_{\mathbf{n} \neq \mathbf{n}_{\text{abs}}}(t \gg t_r) \simeq \pi_{\mathbf{n}} e^{-t/\tau}, \quad P_{\mathbf{n}_{\text{abs}}}(t \gg t_r) \simeq 1 - e^{-t/\tau}. \quad (5.11)$$

The exponential decay factor, $e^{-t/\tau}$, describes the ‘leaking’ process from the metastable state into the absorbing state. The second equation follows from the normalisation of $P_{\mathbf{n}}(t)$.

To find the mean fixation time of the type-2 cells, we substitute Eq. (5.11) into the master equation (5.2) to obtain the quasi-stationary master equation (QSME)

$$-\frac{1}{\tau} \pi_{\mathbf{n} \neq \mathbf{n}_{\text{abs}}} = \sum_{\mathbf{v}} [T_{\mathbf{n}-\mathbf{v}}^{\mathbf{v}} \pi_{\mathbf{n}-\mathbf{v}} - T_{\mathbf{n}}^{\mathbf{v}} \pi_{\mathbf{n}}]. \quad (5.12)$$

For $\mathbf{n} = \mathbf{n}_{\text{abs}}$ (the absorbing state) we have

$$\begin{aligned} \frac{1}{\tau} &= \sum_{\mathbf{v}} T_{\mathbf{n}_{\text{abs}}-\mathbf{v}}^{\mathbf{v}} \pi_{\mathbf{n}_{\text{abs}}-\mathbf{v}} \\ &= T_{(0,1,N-1)}^{1 \rightarrow 2} \pi_{(0,1,N-1)} + T_{(1,0,N-1)}^{0 \rightarrow 2} \pi_{(1,0,N-1)}, \end{aligned} \quad (5.13)$$

where we have used $T_{\mathbf{n}_{\text{abs}}}^{\mathbf{v}} = 0$ for all $\mathbf{v}$. Hence if we find the QSD, $\pi_{\mathbf{n}}$, by solving the QSME (5.12), we can determine the mean fixation time, τ , and the probability to have reached fixation by time t , $P_{\mathbf{n}_{\text{abs}}}(t) = 1 - e^{-t/\tau}$. By separating variables in Eq. (5.11), we have reduced the complexity of the master equation (5.2); time does not feature in the QSME (5.12).

We now change variables from $\mathbf{n}$ to $\mathbf{x} = \mathbf{n}/N$, and we will interpret $\mathbf{x}$ as a continuous variable. This approximation is valid as we have already stated that we require N to be large. The continuous version of the QSD is the probability density $\psi(\mathbf{x}) = N\pi_{N\mathbf{x}}$. We employ the WKB ansatz [52] to represent the QSD as

$$N\pi_{N\mathbf{x}} = \psi(\mathbf{x}) = C \exp[-NS_0(\mathbf{x}) - S_1(\mathbf{x}) + \mathcal{O}(N^{-1})], \quad (5.14)$$

where $S_0(\mathbf{x}) = \mathcal{O}(1)$ for all $\sigma \geq 0$ [53], and we have introduced C as a normalisation constant. The discussion of the origin of this ansatz follows in Chap. 6. To find the QSD, we substitute the ansatz into the QSME (5.12), and then follow the existing approaches for solving similar problems (see Ref. [53], for example) by expanding the resulting equation in powers of N^{-1} .

Further analytical progress can be made if the QSME has only one variable. This is relevant in regions I and II of our model, where the population must escape from a metastable state on the 1–2 edge of the concentration simplex. Escape from an interior metastable state can also be studied using the WKB approach. However, the QSME then retains two degrees of freedom, and explicit expressions for the QSD and escape time cannot be obtained. In this scenario we must resort to numerical methods. In the following we describe the WKB approach for the three regions of parameter space that contain additional fixed point separately.

Region I

We first consider the case in which there exists a single additional fixed point of the deterministic dynamics, which is located on the 1–2 boundary of the concentration simplex. In this scenario type-0 cells become extinct very quickly and the population settles into the quasi-stationary distribution along this edge. As the population cannot depart from the 1–2 boundary, the system reduces to one degree of freedom. We parametrise the system in terms of the concentration of type-1 cells, x_1 . We then have $x_2 = 1 - x_1$, and $x_1 = 0$ is the absorbing state (in which all cells are of type 2). Along the 1–2 boundary there are only two reactions from Eq. (5.1) which have non-zero rate; $T^{1 \rightarrow 2}$ and $T^{2 \rightarrow 1}$. We express these as the intensive quantities

$$f_+(x_1) = \frac{T_{Nx_1}^{2 \rightarrow 1}}{N} = (1 - x_1) \times \frac{(1 - u_2)r_1x_1}{r_1x_1 + r_2(1 - x_1)}, \quad (5.15a)$$

$$f_-(x_1) = \frac{T_{Nx_1}^{1 \rightarrow 2}}{N} = x_1 \times \frac{u_2r_1x_1 + r_2(1 - x_1)}{r_1x_1 + r_2(1 - x_1)}. \quad (5.15b)$$

The phase portrait in Fig. 5.10, as described in Sect. 2.6, highlights the location of the fixed point and the direction of selection away from the absorbing boundary at $x_1 = 0$. In the figure we have used $\dot{x}_1 = f_+(x_1) - f_-(x_1)$. This is exactly Eq. (5.5b) with $x_0 = 0$ and $x_2 = 1 - x_1$.

The aim of this subsection is to compute an explicit expression for the mean escape time from the metastable state on the 1–2 boundary. The main steps of this calculation are:

- (i) Find an expression for the QSD on the boundary. This is achieved by substituting the WKB ansatz (5.14) into the QSME (5.12) and solving the equations in descending powers of N .
- (ii) Close to the $x_1 = 0$ boundary we must consider the flux to the absorbing state, and so we construct a boundary-layer solution in this region. We consider a




Fig. 5.10 The phase portrait of the deterministic dynamics along the 1–2 boundary in region I. The stable fixed point (filled circle), as described in Eq. (5.7), is found at $x_1^* = 1 - u_2r_1/(r_1 - r_2)$. The absorbing state at $x_1 = 0$ is an unstable fixed point (empty circle). Parameters are $r_1 = 1.05$, $r_2 = 1.00$ and $u_2 = 10^{-2}$. Other parameters do not feature along this boundary

Taylor expansion of Eq. (5.12) about $x_1 = 0$ without imposing the WKB ansatz as the solution.

- (iii) The two solutions described above are matched within the boundary-layer to provide an accurate approximation for the mean escape time, τ .

Once this analysis is complete, we compare the predictions with the results of stochastic simulations of the original model.

(i) Calculating the Quasi-Stationary Distribution

We will now compute the approximate distribution about the metastable state. On the boundary the QSME (5.12) is

$$-\frac{1}{N\tau}\psi(x_1) = \sum_{v=\pm 1} \left[f_v \left(x_1 - \frac{v}{N} \right) \psi \left(x_1 - \frac{v}{N} \right) - f_v(x_1) \psi(x_1) \right], \quad (5.16)$$

which we can Taylor expand about x_1 . We first replace $\psi(x_1)$ by the ansatz (5.14) and perform the expansion

$$\begin{aligned} \psi \left(x_1 - \frac{v}{N} \right) &= C \exp \left[-N S_0 \left(x_1 - \frac{v}{N} \right) - S_1 \left(x_1 - \frac{v}{N} \right) \right] \\ &\approx C \exp \left[-N \left(S_0(x_1) - \frac{v S'_0(x_1)}{N} + \frac{v^2 S''_0(x_1)}{2N^2} \right) - \left(S_1(x_1) - \frac{v S'_1(x_1)}{N} \right) \right] \\ &= \psi(x_1) \exp \left[v S'_0(x_1) \right] \exp \left[-\frac{1}{N} \left(\frac{v^2 S''_0(x_1)}{2} - v S'_1(x_1) \right) \right] \\ &\approx \psi(x_1) \exp \left[v S'_0(x_1) \right] \left[1 + v \frac{S'_1(x_1)}{N} - \frac{v^2 S''_0(x_1)}{2N} \right]. \end{aligned} \quad (5.17)$$

We have here used $S'_\sigma(x_1)$ to represent differentiation with respect to x_1 , and we have assumed that the $S_\sigma(x_1)$ are differentiable. The expansion of the transition rates is simply $f_v(x_1 - v/N) \approx f_v(x_1) - v f'_v(x_1)/N$. Equation (5.16) can hence be approximated by

$$\begin{aligned} 0 &= \sum_{v=\pm 1} f_v(x_1) \left[e^{v S'_0(x_1)} - 1 \right] \\ &\quad + \frac{1}{N} \sum_{v=\pm 1} e^{v S'_0(x_1)} \left[v f_v(x_1) S'_1(x_1) - \frac{v^2 f_v(x_1) S''_0(x_1)}{2} - v f'_v(x_1) \right] \\ &\quad + \mathcal{O}(N^{-2}), \end{aligned} \quad (5.18)$$

where we have ignored the term $\mathcal{O}((N\tau)^{-1})$. We will soon confirm that this term is much less than $\mathcal{O}(N^{-2})$ as τ scales as e^N .

The leading-order terms in the system size of Eq. (5.18) satisfy

$$f_+(x_1) (e^p - 1) + f_-(x_1) (e^{-p} - 1) = 0, \quad (5.19)$$

where we have introduced $p = S'_0(x_1)$. Equation (5.19) can be expressed as a quadratic equation for e^p , from which we find two values of p . One of these is $p = 0$, and the other is

$$p(x_1) = \ln \left[\frac{f_-(x_1)}{f_+(x_1)} \right]. \quad (5.20)$$

We can then integrate this second expression to recover $S_0(x_1)$, which we define as

$$S_0(x_1) = \int_{x_1^*}^{x_1} \ln \left[\frac{f_-(q)}{f_+(q)} \right] dq. \quad (5.21)$$

The choice of integration constant is arbitrary as it will be absorbed into the normalisation coefficient, C . Here we chose the constant such that $S_0(x_1^*) = 0$, where x_1^* is the location of the fixed point as described in Eq. (5.7). Substituting in the transition rates from Eq. (5.15), we obtain

$$S_0(x_1) = \left[(1-q) \ln[(1-u_2)r_1(1-q)] - \frac{u_2r_1q + r_2(1-q)}{r_2 - u_2r_1} \ln[u_2r_1q + r_2(1-q)] \right]_{q=x_1^*}^{q=x_1}. \quad (5.22)$$

This function is well-behaved at all points $0 \leq x_1 \leq 1$, as shown in Fig. 5.11a. We find that $S_0(x_1^*)$ is a minimum of $S_0(x_1)$, such that the leading-order contribution (in the system size) to the QSD,

$$\psi(x_1) = C \exp[-N S_0(x_1)], \quad (5.23)$$

is peaked about the fixed point x_1^* . To determine the normalisation coefficient, C , we can expand the QSD about the fixed point




Fig. 5.11 **a** The leading-order term in the system-size of the WKB expansion, $S_0(x_1)$, from Eq. (5.22). **b** The next-leading-order term, $S_1(x_1)$ from Eq. (5.29), diverges as $x_1 \rightarrow 0$. This divergence is cancelled by replacing $S_1(x_1)$ by $\phi(x_1)$ from Eq. (5.30). Arrows indicate the location of the stable fixed point at $x_1^* = 1 - u_2r_1/(r_1 - r_2)$. Parameters are as in Fig. 5.10

$$\psi(x_1) \approx C \exp \left[-\frac{N}{2} (x_1 - x_1^*)^2 S_0''(x_1^*) \right], \quad (5.24)$$

where we have used $S_0(x_1^*) = S_0'(x_1^*) = 0$. By applying the normalisation condition $\int_0^1 \psi(q) dq = 1$, we find that

$$C \approx \sqrt{\frac{N S_0''(x_1^*)}{2\pi}}, \quad (5.25)$$

where we have assumed that $\int_0^1 \psi(q) dq \approx \int_{-\infty}^{\infty} \psi(q) dq$, or equivalently that $\psi(x_1)$ is sharply peaked about x_1^* .

The next-leading-order terms in the system size of Eq. (5.18) satisfy

$$\sum_{\nu=\pm 1} e^{\nu p(x_1)} \left[\nu f_\nu(x_1) S_1'(x_1) - \frac{\nu^2 f_\nu(x_1) p'(x_1)}{2} - \nu f_\nu'(x_1) \right] = 0. \quad (5.26)$$

Substituting Eq. (5.20) into this expression, it can be shown that $S_1'(x_1)$ satisfies

$$S_1'(x_1) = \frac{1}{2} \left[\frac{f_+'(x_1)}{f_+(x_1)} + \frac{f_-'(x_1)}{f_-(x_1)} \right]. \quad (5.27)$$

Integrating this equation gives

$$S_1(x_1) = \left[\frac{1}{2} \ln [f_+(q) f_-(q)] \right]_{q=x_1^*}^{q=x_1}, \quad (5.28)$$

where again we have chosen the arbitrary constant such that $S_1(x_1^*) = 0$. Substituting in the transition rates from Eq. (5.15), we find

$$S_1(x_1) = \left[\frac{1}{2} \ln \left(\frac{(1 - u_2) r_1 q (1 - q) \times [u_2 r_1 q^2 + r_2 q (1 - q)]}{[r_1 q + r_2 (1 - q)]^2} \right) \right]_{q=x_1^*}^{q=x_1}. \quad (5.29)$$

This function diverges at $x_1 = 0$ and $x_1 = 1$, as shown in Fig. 5.11b. We can ignore the divergence at $x_1 = 1$ as this state is not of interest. The divergence at $x_1 = 0$ cannot be ignored if we want to compute the time to reach this absorbing state. To overcome this we introduce the function

$$\phi(x_1) = S_1(x_1) - \ln(x_1), \quad (5.30)$$

which is well-behaved at $x_1 = 0$, as shown in Fig. 5.11b. With this we can express the QSD as

$$\pi_{N x_1} = \frac{\psi(x_1)}{N} = \frac{C}{N x_1} \exp[-N S_0(x_1) - \phi(x_1)], \quad (5.31)$$




Fig. 5.12 The measured distribution of concentrations of type-1 cells, x_1 , at given times (symbols) in region I of parameter space. These are obtained from an ensemble of 10^5 Gillespie simulations of the microscopic model. Inset is the same data with a logarithmic vertical axis. We show the WKB approximation (5.45) (solid lines; filled bars for $x_1 = 0$), and the boundary-layer solution (5.46) (dashed lines), which are in good agreement with simulation results in their respective regimes. The distributions away from $x_1 = 0$ have been re-scaled by a factor 50 for optical convenience. The arrow indicates the location of the deterministic fixed point, given by Eq. (5.7). Parameters are $r_0 = 1.0$, $r_1 = 1.05$, $r_2 = 1.0$, $u_1 = u_2 = 10^{-2}$, and $N = 200$

which is illustrated in Fig. 5.12. The x_1^{-1} divergence has now been isolated outside of the exponent, which will prove to be necessary when we construct the boundary-layer solution below.

We can attempt to determine the mean fixation time, τ . As an initial guess we can simply rearrange Eq. (5.13) to write

$$\begin{aligned} \tau &= \frac{1}{f_-(1/N)\psi(1/N)} \\ &= \left(\frac{1}{N} \frac{u_2 r_1 / N + r_2(1 - 1/N)}{r_1 / N + r_2(1 - 1/N)} \times \frac{C}{1/N} \exp[-N S_0(1/N) - \phi(1/N)] \right)^{-1} \\ &= \frac{r_1 + r_2(N - 1)}{u_2 r_1 + r_2(N - 1)} \times \sqrt{\frac{2\pi}{N S_0''(x_1^*)}} \exp[N S_0(1/N) + \phi(1/N)]. \end{aligned} \quad (5.32)$$

However we expect the QSD (5.31) to be a less accurate description of the true distribution as the distance from the fixed point increases. This is because we have neglected any flux into the absorbing state. Instead we can construct a boundary-layer solution close to $x_1 = 0$.

(ii) Boundary-Layer Solution

The boundary-layer solution is calculated by expanding the QSME (5.12) [and (5.13)] about $x_1 = 0$ without imposing a specific form for the QSD (i.e. we do not use the WKB ansatz). We return to the discrete coordinates $n_1 = N x_1$ and $\pi_{n_1} = \psi(n_1/N)/N$. For $n_1 = 0$ we have

$$\begin{aligned} \frac{1}{\tau} &= N f_- \left(\frac{1}{N} \right) \pi_1 \approx N \left(f_-(0) + \frac{1}{N} f'_-(0) \right) \pi_1 = f'_-(0) \pi_1 \\ \Rightarrow \pi_1 &\approx \frac{1}{\tau f'_-(0)}. \end{aligned} \quad (5.33)$$

For $1 \leq n_1 \ll N$ we have

$$0 = \sum_{v=\pm 1} f'_v(0) [(n_1 - v) \pi_{n_1-v} - n_1 \pi_{n_1}], \quad (5.34)$$

or equivalently

$$\theta_{n_1+1} = \left[1 + \frac{f'_+(0)}{f'_-(0)} \right] \theta_{n_1} - \frac{f'_+(0)}{f'_-(0)} \theta_{n_1-1}, \quad (5.35)$$

where $\theta_{n_1} = n_1 \pi_{n_1}$. This recursive system can be solved to arrive at

$$\begin{aligned} \theta_{n_1} &= \sum_{i=0}^{n_1-1} \left(\frac{f'_+(0)}{f'_-(0)} \right)^i \theta_1 = \frac{1 - \left(\frac{f'_+(0)}{f'_-(0)} \right)^{n_1}}{1 - \frac{f'_+(0)}{f'_-(0)}} \theta_1 \\ \Rightarrow \pi_{n_1} &= \frac{1 - \left(\frac{f'_+(0)}{f'_-(0)} \right)^{n_1}}{1 - \frac{f'_+(0)}{f'_-(0)}} \frac{\pi_1}{n_1}. \end{aligned} \quad (5.36)$$

This is the boundary-layer solution of the QSME (5.12) for dynamics constrained to the 1–2 edge of the concentration simplex. This is shown in Fig. 5.12. The derivatives of the transition rates (5.15) at $x_1 = 0$ are

$$f'_+(0) = \frac{(1 - u_2)r_1}{r_2} > 1, \quad (5.37a)$$

$$f'_-(0) = 1, \quad (5.37b)$$

where the inequality arises from the existence condition of the fixed point, Eq. (5.9). Now using $[f'_+(0)]^{n_1} \gg 1$ for suitably large n_1 , i.e. at the edge of the boundary layer, as well as Eq. (5.33), we can approximate the solution (5.36) as

$$\pi_{n_1} \approx \frac{[f'_+(0)]^{n_1}}{[f'_+(0) - 1]} \frac{1}{\tau n_1}. \quad (5.38)$$

We now want to match Eq. (5.38) with the WKB solution (5.31) to determine an accurate value of τ .

(iii) Matching WKB and Boundary-Layer Solutions

By expanding the WKB solution (5.31) about $x_1 = 0$, we recover

$$\begin{aligned}\pi(Nx_1) &\approx \frac{C}{Nx_1} \exp[-NS_0(0) - Nx_1 S'_0(0) - \phi(0)] \\ &= \frac{C}{Nx_1} \exp[-Nx_1 p(0)] \exp[-NS_0(0) - \phi(0)].\end{aligned}\quad (5.39)$$

Using Eq. (5.20), we can evaluate $\exp[-p(0)]$ as

$$\exp[-p(0)] = \lim_{q \rightarrow 0} \frac{f_+(q)}{f_-(q)} = \lim_{q \rightarrow 0} \frac{f'_+(q)}{f'_-(q)} = f'_+(0). \quad (5.40)$$

Likewise, the function $\exp[-\phi(0)]$ can be evaluated as

$$\begin{aligned}\exp[-\phi(0)] &= \lim_{q \rightarrow 0} \left[\frac{f_+(q)f_-(q)}{q^2} \right]^{-1/2} \times [f_+(x_1^*)f_-(x_1^*)]^{1/2} \\ &= \lim_{q \rightarrow 0} \left[\frac{f''_+(q)f_-(q) + 2f'_+(q)f'_-(q) + f_+(q)f''_-(q)}{2} \right]^{-1/2} \times f_+(x_1^*) \\ &= \frac{1}{\sqrt{f'_+(0)}} \times f_+(x_1^*),\end{aligned}\quad (5.41)$$

where we have used $f_+(x_1^*) = f_-(x_1^*)$ and $f_+(0) = f_-(0) = 0$. Hence we can write Eq. (5.39) as

$$\pi(Nx_1) \approx \frac{C}{Nx_1} [f'_+(0)]^{Nx_1} \frac{f_+(x_1^*)}{\sqrt{f'_+(0)}} \exp[-NS_0(0)]. \quad (5.42)$$

We can now equate this with Eq. (5.38) for $x_1 > 0$, and from this we obtain the mean fixation time

$$\tau = \frac{1}{C} \frac{\sqrt{f'_+(0)}}{f_+(x_1^*)[f'_+(0) - 1]} \exp[NS_0(0)]. \quad (5.43)$$

Inserting the explicit expressions for $S_0(0)$, the transition rates, and the normalisation coefficient, we can write the escape time from the metastable on the 1–2 boundary as

$$\begin{aligned}\tau &= \sqrt{\frac{2\pi r_1 r_2}{Nu_2}} \frac{(1 - u_2)(r_1 - r_2)}{[(1 - u_2)r_1 - r_2]^2} \\ &\times \exp \left\{ N \left[\frac{u_2 r_1}{r_2 - u_2 r_1} \ln \left[\frac{u_2(1 - u_2)r_1^2}{r_1 - r_2} \right] + \ln[(1 - u_2)r_1] - \frac{r_2}{r_2 - u_2 r_1} \ln[r_2] \right] \right\}.\end{aligned}\quad (5.44)$$

As expected, the mean escape time increases exponentially with the system size, but the dependence on the remaining model parameters is very messy.

Results

We can now compare these results with probability distributions and fixation times obtained from Gillespie simulations of the microscopic model described in Sect. 5.2. In Fig. 5.12, we show our theoretical approximation for the distribution of the concentration of type-1 cells, which is given by

$$P_{[0, N x_1, N(1-x_1)]}(t) = \frac{C}{N x_1} \exp[-N S_0(x_1) - \phi(x_1)] \times \exp(-t/\tau) \text{ for } x_1 > 0, \quad (5.45a)$$

$$P_{(0,0,N)}(t) = P_{\text{abs}}(t) = 1 - \exp(-t/\tau) \text{ for } x_1 = 0. \quad (5.45b)$$

Initially this distribution will be an inaccurate representation of the true distribution as the type-0 cells will not yet have become extinct, but it will become increasingly accurate as time progresses. This approximation is compared with the distribution obtained from simulations in Fig. 5.12. Specifically, we measure the probability to observe a fraction x_1 of type-1 cells at a given time t (independent of whether type-0 cells have become extinct or not). The data in the figure reveals good agreement between theory and simulation data away from the region close to $x_1 = 0$. There is also good agreement between Eq. (5.45b) and the probability to have reached the absorbing state by time t measured from the simulations.

Close to the absorbing state $x_1 = 0$, the boundary-layer approximation for the distribution is given by

$$P_{[0, N x_1, N(1-x_1)]}(t) = \frac{1 - [f'_+(0)]^{N x_1}}{1 - f'_+(0)} \frac{1}{\tau N x_1} \exp(-t/\tau) \text{ for } 1/N < x_1 \ll 1. \quad (5.46)$$

This shows improved agreement with simulation results close to $x_1 = 0$ when compared to WKB solution (5.45). This is emphasised in the inset logarithmic plot of Fig. 5.12.

Results for the mean fixation time in region I are shown in Fig. 5.13a. In Fig. 5.13b we plot the probability that type-2 cells have reached fixation by time $t^* = 3 \times 10^3$ (including fixation earlier than that), and compare this with the prediction of Eq. (5.45b). If a cellular generation lasts for one day, then this time is $\mathcal{O}(10)$ years, which is the same length of time that appears in multiple studies of mutation acquisition [14, 16]. The mean fixation times increase exponentially with the fitness of type-1 cells, r_1 . This is a consequence of the increasing height of the selection ‘barrier’ which must be overcome for type-2 cells to reach fixation. Also, increasing r_1 pushes the boundary fixed point towards the all-type-1 state, which results in a further increase in fixation time (or decrease in probability of fixation by time t^*). If the fixed point approaches the all-type-1 state, the probability of the population reaching this corner of the simplex due to demographic fluctuations increases. Thus




Fig. 5.13 **a** Mean fixation time of type-2 cells in region I. *Symbols* correspond to mean fixation times from simulations of the model (averaged over an ensemble of $2\text{--}5 \times 10^3$ samples) initiated in the all-wild-type state. Shape of *symbol* indicates fitness of type-2 cells (see legend); *filled symbols* are for $u_1 = u_2 = 10^{-2}$, *empty symbols* are for $u_1 = u_2 = 10^{-3}$. *Solid lines* (high mutation) and *dashed lines* (low mutation) are the WKB prediction for fixation time, Eq. (5.44). The approximation breaks down when $(1 - u_2)r_1 \simeq r_2$, which is when the fixed point approaches the absorbing state. **b** The probability that type-2 cells have reached fixation by time $t^* = 3 \times 10^3$. *Lines* correspond to the WKB prediction, Eq. (5.45b). Colours and symbols follow the same convention as in panel (a). Remaining parameters are $r_0 = 1.0$ and $N = 100$

increasing r_1 decreases the probability of stochastic tunnelling occurring. Increasing the fitness of type-2 cells, on the other hand, pushes the metastable state closer to the absorbing state. This leads to a significant reduction in the fixation time (increase in fixation probability by time t^*). Increasing the mutation rate u_2 has a similar effect to increasing r_2 ; the fixed point approaches the absorbing state, and the net effect of selection away from the absorbing state is reduced, leading to a decrease in the fixation time. In line with the previous literature [24, 30], increasing the mutation rate increases the probability of tunnelling.

In both panels of Fig. 5.13 the theoretical predictions from the WKB method are in excellent agreement with simulation results. This is the case even at the moderate population size of $N = 100$. In Fig. 5.14a we show that this accuracy is retained for increasing N . Small deviations between the theory and simulation results occur when mutation rates are low (dashed lines and open symbols in Fig. 5.13). This is seen more clearly in Fig. 5.14b. The theory underestimates the fixation time (overestimates the probability of fixation by time t^*) at small values of u . This is a consequence of assuming that the population approaches the metastable state in a negligible amount of time. For very small mutation rates, it takes an increasing period of time for successful (i.e. non-vanishing) mutant lineages to appear. Deviations between theory and simulation results occur when $(1 - u_2)r_1 \simeq r_2$. At this point the theory breaks down as the fixed point on the 1–2 edge approaches the absorbing state. The barrier associated with adverse selection is then negligible and the assumptions underlying the WKB approximation are no longer justified.

Finally we can compare the predictions of our theory with those based on the homogeneous-state approach, Eq. (5.3). In particular we compare against the results




Fig. 5.14 **a** Mean fixation time of type-2 cells in region I as a function of system size. *Symbols* correspond to mean fixation times from simulations of the model (averaged over an ensemble of 10^3 samples) initiated in the all-wild-type state. *Filled symbols* are for $u_1 = u_2 = 10^{-2}$, *empty symbols* are for $u_1 = u_2 = 10^{-3}$. *Solid lines* (high mutation) and *dashed lines* (low mutation) are the WKB prediction for fixation time, Eq. (5.44). **b** Mean fixation time of type-2 cells in region I as a function of mutation rates. Here we have used $r_2 = 1$ and $N = 100$. Symbols again correspond to mean fixation times from simulations of the model (averaged over an ensemble of 10^3 samples) initiated in the all-wild-type state




Fig. 5.15 Mean fixation time of type-2 cells in region I as a function of system size. *Symbols* correspond to mean fixation times from simulations of the model (averaged over an ensemble of 10^3 samples) initiated in the all-wild-type state. *Solid lines* are the prediction of Eq. (5.44). *Dashed lines* are the prediction of Ref. [24], which is computed from Eq. (5.3) using the transition rates shown in Eq. (5.4)

of Ref. [24]. In Fig. 5.15 we show mean fixation times as a function of the system size. This plot shows the homogeneous-state approach (dashed) of Ref. [24] breaks down in region I for large values of N . The figure reveals that the WKB predictions are of good accuracy for population sizes greater than 50–200 or so, depending on model parameters. This confirms that the approximations we make (large population size) only limit the range of validity of the WKB approach to a relatively minor degree. If the population of cells is too small then the strength of the noise is greater than the effect of selection, and the idea of a metastable state breaks down. If this is the case the process is mutation-limited, and similar fixation times are observed when

the mutation rate is fixed, but fitness parameters are varied (circles and diamonds converge at small N , as do squares and triangles, in Fig. 5.15). The WKB solution becomes more accurate with increasing population size. For $N \gtrsim 100$, Eq. (5.44) is in much better agreement with simulation results than the existing theoretical estimates.

Region II

We now analyse the case where the system first approaches a fixed point in the interior of the concentration simplex, before reaching the 1–2 boundary and settling into the QSD as described above. When a fixed point exists away from the state-space boundaries, the solution procedure described above is no longer viable. This is because the QSME (5.12) retains two degrees of freedom. The ansatz (5.14), however, is still a valid approximation to the QSD about this interior fixed point. Substituting this into the QSME, and taking the leading-order terms in the system size gives the expression

$$\sum_{\mathbf{v}} f_{\mathbf{v}}(\mathbf{x}) (\exp[\mathbf{v} \cdot \mathbf{p}] - 1) = 0, \quad (5.47)$$

where $f_{\mathbf{v}}(\mathbf{x}) = T_{N\mathbf{x}}^{\mathbf{v}}/N$ and $\mathbf{p} = \nabla S_0(\mathbf{x})$. Unlike the boundary case described above, we cannot discern a non-zero $\mathbf{p}$ which satisfies Eq. (5.47). Instead, we identify Eq. (5.47) as a Hamilton–Jacobi³ equation of the form $H(\mathbf{x}, \mathbf{p}) = 0$ and consider the characteristic equations, also known as Hamilton’s equations [54]. These are of the form

$$\dot{x}_i = \frac{\partial H}{\partial p_i}, \quad \dot{p} = -\frac{\partial H}{\partial x_i}. \quad (5.48)$$

This formulation will be discussed in much more detail in the next Chapter, but for now we will state that these equations define characteristic curves which satisfy the principle of least action, i.e. they define the most-likely path between two points in space [54]. If these trajectories can be found, subject to initial and final conditions, then the quasi-stationary distribution can be reconstructed by calculating

$$S_0(\mathbf{x}) = \int \mathbf{p}(t) \cdot \dot{\mathbf{x}}(t) dt = \int \mathbf{p} \cdot d\mathbf{x}, \quad (5.49)$$

which is the integral of $\mathbf{p}$ along the trajectory described by $\mathbf{x}$.

For the initial conditions of these trajectories we take $\mathbf{x}$ to be the deterministic fixed point in the interior of the domain, i.e. $\mathbf{x} = \mathbf{x}^*$ given by Eq. (5.10), and $\mathbf{p} = \mathbf{0}$. For the final conditions, we must find a value of $\mathbf{p}$. If the final position $\mathbf{x}$ is not a fixed point, $\mathbf{p}$ is not uniquely determined, and it becomes a variational problem to find a suitable $\mathbf{p}$. We leave this discussion for Chap. 6. If $\mathbf{x}$ is a fixed point, the trajectory must end at this point and we have $\dot{\mathbf{x}} = \mathbf{0}$ and $\dot{\mathbf{p}} = \mathbf{0}$. Hence the fixed

³After William Rowan Hamilton (1805–1865) and the previously introduced Carl Gustav Jacob Jacobi.

points of Eqs. (5.48) determine the final values of $\mathbf{p}$, and we are then left to solve a two-boundary-value problem.

We can find three fixed points of Eqs. (5.48) with $\mathbf{p} = \mathbf{0}$. These are the fixed points of the deterministic equations (5.5), which we label as $M_1 = (0, 1, 0, 0)$ for the absorbing state, $M_2 = (x_1^*, 1 - x_1^*, 0, 0)$ for the boundary fixed point [with x_1^* given by Eq. (5.7)], and $M_3 = (x_1^*, x_2^*, 0, 0)$ for the interior fixed point [given by Eq. (5.10)]. The labelling of these states follows the convention of Ref. [32], who studied a similar scenario in a predator–prey system. As stated, all trajectories start at the interior point M_3 . The end points require a non-zero value of $\mathbf{p}$, otherwise we will end up with $S_0(\mathbf{x}) = 0$. These so-called ‘fluctuational fixed points’ of Eqs. (5.48) are labelled as F_1 for the absorbing state and F_2 for the boundary fixed point.

Problems of this type can be tackled using different techniques: firstly, the equations of motion (5.48) could be integrated using a shooting method to find the least-action trajectory with a given final point [32, 34, 36]; secondly, the equations can be integrated using an iterative scheme which converges to the optimal trajectory [55] connecting given start and end points. Alternatively, minimisation techniques could be used to find the least-action trajectory [56]. A detailed analysis of these different methods follows in Chap. 6. In the present application we find that the iterative method quickly converges for our problem. Results presented in the following use this method, which is outlined below.

As stated earlier, the most probable path to the absorbing state is the sequential extinction path, where first type-0 cells are lost, and then type-1 cells. Thus we only focus on the trajectory from the interior fixed point M_3 to the boundary fluctuational fixed point F_2 , and do not consider the simultaneous extinction path from M_3 to F_1 . To determine the least-action trajectory, we initially fix the values of $\mathbf{p}$ for all times to the values at F_2 , which are found by numerically solving $\dot{\mathbf{x}} = 0$ and $\dot{\mathbf{p}} = 0$ in Eq. (5.48). We then numerically integrate the equations of motion (5.48) for the position vector $\mathbf{x}$ forward in time, starting at M_3 and keeping $\mathbf{p}$ constant. This integration is carried out for a sufficient range of time to reach the vicinity of the fixed point F_2 , but not too long to avoid numerical errors building up. In the next step the relations for $\mathbf{p}$ in Eq. (5.48) are integrated backwards in time using the trajectory $\mathbf{x}(t)$ found in the previous iteration. The $\mathbf{p}$ values at the start of this backwards integration are chosen as those corresponding to F_2 . This procedure is then iterated, with alternating forward and backward integration of Hamilton’s equation. At each step of the procedure the ‘action’ of the path is calculated using Eq. (5.49). Following the convention of Ref. [32], this value is labelled S_{32} as it begins at point M_3 and ends at point F_2 . The iteration of alternating forward and backward integration is then repeated until S_{32} converges to a stable value. The resulting trajectory through the concentration simplex [i.e. the projection of the four-dimensional trajectory $(\mathbf{x}, \mathbf{p})$ onto the two dimensional plane of $\mathbf{x}$] is shown in Fig. 5.16a.

The final value of S_{32} characterises the mean time taken to escape from the interior fixed point to the boundary through the relation,

$$\tau_{32} \sim \frac{C_{32}}{\sqrt{N}} e^{NS_{32}}, \quad (5.50)$$




Fig. 5.16 **a** The dominant escape path from the interior metastable state to the 1–2 boundary in region II. Parameters are $r_0 = 1.00$, $r_1 = 0.98$, $r_2 = 0.95$, and $u_1 = u_2 = 10^{-2}$. The *thin line* is the trajectory shown in Fig. 5.7 for $N = 300$. **b** The dominant escape path from the interior metastable state to the absorbing state in region III. Parameters are $r_0 = 1.00$, $r_1 = 0.95$, $r_2 = 0.98$, and $u_1 = u_2 = 10^{-2}$. The *thin line* is the trajectory shown in Fig. 5.8 for $N = 300$

where C_{32} is a constant that is found by fitting to simulation data for the time taken to reach the boundary [33]. This expression has the same functional dependence on N as the one given in Eq. (5.44). The mean fixation time is given by $\tau_{32} + \tau_{21}$, which is the time to escape from the interior metastable state to the boundary plus the time to escape from the boundary to the absorbing state; it is a two-hit process. The latter time, τ_{21} , is given by Eq. (5.44). Thus in region II the mean fixation time is given by

$$\tau = \frac{C_{32}}{\sqrt{N}} e^{Ns_{32}} + \tau_{21}. \quad (5.51)$$

Small changes to the parameters now have significant effects on the fixation time, as shown in Fig. 5.17 (filled symbols/solid lines). Increasing the fitness of the type-2 cells moves both the interior fixed point and the boundary fixed point towards the absorbing all-type-2 state. It also reduces the strength of selection away from the absorbing state. These combined effects dramatically reduce the mean fixation time, and its rate of increase with the population size.

As in region I, the probability of tunnelling decreases as the fitness advantage of type-1 cells over type-2 cells increases. This is because the fixed point on the 1–2 edge approaches the all-type-1 state. For the same reason, the tunnelling rate decreases as the mutation rates decrease.

Region III

In this region fixation is controlled solely by the escape from the interior metastable state; it is a one-hit process. Type-0 cells are the most advantageous in the sequence, and the stable interior fixed point is located close to the all-wild-type state. The mean fixation time is calculated again by considering the least-action trajectories described by Hamilton's equations (5.48), which are computed using the iterative method described above. We now focus on trajectories from the stable interior fixed




Fig. 5.17 Fixation time in region II (*filled symbols*) and region III (*empty symbols*) of Fig. 5.3 as a function of system size from simulations, averaged over 100 realisations. *Lines* are from the theory, see Eq. (5.51) for region II (*solid lines*) and Eq. (5.52) for region III (*dashed lines*). Remaining parameters are $r_0 = 1.0$ and $u_1 = u_2 = 10^{-2}$

point M_3 to the absorbing fluctuational fixed point F_1 . This direct path from the metastable state to the all-type-2 state is shown in Fig. 5.16b. The probability of tunnelling is higher than in the previous cases. It increases as the fitness of type-2 cells and the mutation rates increase as the stable interior fixed point moves to lower numbers of type-1 cells (i.e. away from the all-type-1 state).

The ‘action’, S_{31} , is computed from Eq. (5.49) by integrating along the Hamiltonian trajectory. The mean fixation time is then given by

$$\tau_{31} \sim \frac{C_{31}}{\sqrt{N}} e^{N S_{31}}, \quad (5.52)$$

where again C_{31} is a constant which is found by fitting to simulation results. We see in Fig. 5.17 (empty symbols/dashed lines) that varying the model parameters has a lesser effect on fixation times than in region II. In region III, fixation is a one-hit process—the population only has to escape the interior stable fixed point—and not a two-hit process as in region II where the effects of the two steps are compounded. Contrary to the results for region I, the mean fixation time is a decreasing function of r_1 in region III. This can be explained as follows: by increasing r_1 , the selection strength away from the 1–2 boundary decreases and the stable state moves to higher type-1 numbers, such that the population has an improved chance of reaching the 1–2 boundary. From there selection is directed towards the absorbing state, and the time spent on the 1–2 boundary is negligible compared to the time to reach this edge. Hence, the fixation time reduces as type-1 cells become more fit. The rate of increase of the fixation time with the population size reduces as well.

There are systematic deviations between theory and simulation results in the data set shown as open triangles in Fig. 5.17, and to a lesser extent also for the data shown

as open diamonds. This is attributed to the fact that the fitness parameters r_1 and r_2 are very similar to each other or equal for these instances, and they are also close to the fitness of the wild-type. Selection is then close neutral and the metastable state is only weakly attracting. The WKB approach then reaches its limits as the assumption of a long-lived metastable state begins to break down.

Comparison Across Parameter Space

Finally, we bring together the different theoretical approximations to the mean fixation time in Fig. 5.18. We fix all model parameters except for r_1 , and we sweep across regions I, II and III. The theoretical predictions of Eqs. (5.44), (5.51), and (5.52) are in excellent agreement with the simulation results within their respective regimes. The expression for the time to escape from the metastable state on the 1–2 edge of the concentration simplex, Eq. (5.44), is only valid for $(1 - u_2)r_1 > r_2$. As Eq. (5.51) is dependent on this expression, it too reaches its limits close to the boundary of regions II and III. However, the direct simultaneous extinction path used to characterise absorption in region III accurately predicts the mean fixation time in part of region II. This suggests that for $(1 - u_2)r_1 \simeq r_2$, when the boundary fixed point approaches the absorbing state, it is the simultaneous extinction path, not the sequential extinction path, that is dominant. As previously mentioned, the prediction based on the homogeneous-state assumption (from Ref. [24]) does not accurately represent the data.

We have mentioned through this discussion that the probability of stochastic tunnelling occurring is dependent on the locations of the fixed points. In Fig. 5.19 we plot the probability of tunnelling (measured from simulations) across the r_1 – r_2





Fig. 5.18 Mean fixation times in regions I, II and III (circles) averaged over 10^3 realisations. The thick line in good agreement in region I is from Eq. (5.44). The thick line in region III is prediction of Eq. (5.52), where the coefficient C_{31} is found by fitting to numerically obtained fixation times as a function of N . The dashed line in region II is the prediction of Eq. (5.51), where the coefficient C_{32} is found by fitting Eq. (5.50) to the time taken to reach the 1–2 boundary as a function of N . The faint dotted line is the prediction of Ref. [24], which is based on the homogeneous-state assumption. Model parameters are fixed to $r_0 = 1.0$, $r_2 = 0.97$, $u_1 = u_2 = 10^{-2}$ and $N = 300$

Fig. 5.19 This heat map shows the probability that fixation occurs via the stochastic tunnelling route, as opposed to sequential fixation route. The probability is calculated as the fraction of realisations (out of an ensemble of 10^4 simulations) that do not pass through the all-type-1 state before reaching fixation. The remaining model parameters are $r_0 = 1.0$, $u_1 = u_2 = 10^{-2}$ and $N = 200$



parameter space. In regions I and II, for parameters sufficiently far from the region boundaries, the probability of tunnelling is zero. This is because the fixed point on the 1–2 boundary is close the all-type-1 state, and demographic fluctuations are likely to push the population to this homogeneous state. Close to the region boundaries, the fixed point is located sufficiently far away from the all-type-1 state and this helps to ensure the homogeneous type-1 state is not visited. Hence the probability of tunnelling approaches one. In region III, when there is only a single fixed point which is located in the interior of the domain, the probability of tunnelling is one. This agrees with the predictions of the WKB approach, where the dominant escape path from this state proceeds directly to the absorbing state, as shown in Fig. 5.16b. The probability of tunneling is also one in regions IV and V, where the tunnelling route is predicted by the deterministic dynamics as shown in Fig. 5.5.

5.6 Summary

Previous analysis of models describing the accumulation of multiple mutations in a tissue have predominantly focussed on the homogeneous-state assumption. That is the population spends the majority of the time in states where all cells harbour the same number of mutations. These assumptions are only valid if the final mutant in the sequence is the most advantageous cell-type, or if population numbers are small. If this is not the case, then metastable states are found away from the homogeneous states. Our analysis identified the escape from these metastable states as the key bottleneck to fixation of cells with two mutations. For parameter values for which there are no metastable states (i.e. when type-2 cells have the highest fitness), the fixation dynamics is largely governed by the deterministic flow. The rate-limiting

steps are then the appearance of successful mutant lineages [25], and the subsequent fixation of cells with two mutations is a zero-hit process for large population sizes. As such the progression from healthy tissue (all wild-type) to susceptible tissue (all type-2; inactivated TSG) will be fast relative to the cases in which a metastable state exists. If there is one stable fixed point in the deterministic dynamics, the process becomes a one-hit phenomenon limited by the escape from the corresponding metastable state. In regions with two fixed points one observes a two-hit process. The population becomes trapped in a first metastable state, escapes to a second metastable state, and then reaches full fixation.

In addition to this qualitative classification, we used the WKB method to calculate fixation times in parameter regimes previously inaccessible to existing analytical approaches. Our theoretical predictions in principle rely on a limit of large populations, however comparison against simulation results demonstrates the accuracy of our theory even at moderate population sizes of $N = 100$ cells. For populations much smaller than this the assumptions of the WKB method break down. The rate-limiting step is then the occurrence of a successful lineage of mutants and not the escape from metastable states. The expressions obtained from the WKB approach become more accurate as the population size increases.

This analysis allowed us to classify how changes to the fitness landscape, mutation rates, and population size affect the probability of tunnelling and the time-to-fixation of cells harbouring two mutations. In terms of the development of tumours, our analysis shows that the path to accumulating mutations is not simply limited by the mutation rates, but also by the escape from metastable states. Populations can exist in a heterogeneous state for very long periods of time before fluctuations eventually drive the second mutation to fixation. The probability with which stochastic tunnelling occurs is, in part, determined by the location of these metastable states. If they are located close to the all-type-1 state, then the probability of tunnelling is low.

Although our theory is aimed at large population sizes and exponentially growing fixation times, we have shown that it can also make accurate predictions on biologically relevant timescales. Assuming a cell generation lasts for one day, our theory can capture fixation times of around 3 years or more ($> 10^3$ generations). Related studies on the progression of cancer suggest a typical timescale on the order of 10 years to accumulate a sufficient number of mutations [14, 16], which is well within the scope of our theory. However, the times predicted by our theory are extremely sensitive to parameter variation. This limits the parameter ranges for which biologically relevant timescales can be generated. Specifically, the selective (dis)advantages need to be small ($\lesssim 10\%$). This is in agreement with selection coefficients in related studies [16]. Of course the length of a cellular generation can vary by an order of magnitude or so, depending on the specific cell type [3].

Our results do, however, allow an extrapolation to situations when fixation times become very long, for instance for very large populations and/or when selection is strongly against the invading mutants. In these scenarios, stochastic simulations can become too expensive computationally to provide meaningful measurements. Analytical methods based on backward master equations or backward Fokker–Planck equations suffer from computational limitations as well in such cases. Our mathe-

mathematical work complements existing analytical approaches to the Moran model of cells acquiring two successive mutations. This work fills the gap left by the existing literature and leads to a more comprehensive understanding of mutation acquisition and stochastic tunnelling in evolving populations.

References

1. Cancer Research UK, All cancers combined: key facts. Cancer Research UK (2014)
2. Cancer Research UK, Annual report and accounts 2014/15. Cancer Research UK (2015)
3. R.A. Weinberg, *The Biology of Cancer* (Garland Science, New York, 2013)
4. C. Nordling, A new theory on the cancer-inducing mechanism. *Br. J. Cancer* **7**, 68 (1953)
5. P. Armitage, R. Doll, The age distribution of cancer and a multi-stage theory of carcinogenesis. *Br. J. Cancer* **8**, 1 (1954)
6. J. Fisher, Multiple-mutation theory of carcinogenesis. *Nature* **181**, 651 (1958)
7. A.G. Knudson, Mutation and cancer: statistical study of retinoblastoma. *Proc. Natl. Acad. Sci. U.S.A.* **68**, 820 (1971)
8. S.H. Moolgavkar, The multistage theory of carcinogenesis and the age distribution of cancer in man. *J. Natl. Cancer Inst.* **61**, 49 (1978)
9. S.H. Moolgavkar, A.G. Knudson, Mutation and cancer: a model for human carcinogenesis. *J. Natl. Cancer Inst.* **66**, 1037 (1981)
10. S.H. Moolgavkar, E.G. Luebeck, Multistage carcinogenesis: population-based model for colon cancer. *J. Natl. Cancer Inst.* **84**, 610 (1992)
11. L. Nunney, Lineage selection and the evolution of multistage carcinogenesis. *Proc. R. Soc. Lond. B* **266**, 493 (1999)
12. R.A. Gatenby, T.L. Vincent, An evolutionary model of carcinogenesis. *Cancer Res.* **63**, 6212 (2003)
13. F. Michor, Y. Iwasa, M.A. Nowak, Dynamics of cancer progression. *Nat. Rev. Cancer* **4**, 197 (2004)
14. N. Beerewinkel, T. Antal, D. Dingli, A. Traulsen, K.W. Kinzler, V.E. Velculescu, B. Vogelstein, M.A. Nowak, Genetic progression and the waiting time to cancer. *PLoS Comput. Biol.* **3**, e225 (2007)
15. H. Haeno, R.L. Levine, D.G. Gilliland, F. Michor, A progenitor cell origin of myeloid malignancies. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 16616 (2009)
16. I. Bozic, T. Antal, H. Ohtsuki, H. Carter, D. Kim, S. Chen, R. Karchin, K.W. Kinzler, B. Vogelstein, M.A. Nowak, Accumulation of driver and passenger mutations during tumor progression. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 18545 (2010)
17. I. Van Leeuwen, H. Byrne, O. Jensen, J. King, Crypt dynamics and colorectal cancer: advances in mathematical modelling. *Cell Prolif.* **39**, 157 (2006)
18. T. Antal, P.L. Krapivsky, M.A. Nowak, Spatial evolution of tumors with successive driver mutations. *Phys. Rev. E* **92**, 022705 (2015)
19. M. Archetti, D.A. Ferraro, G. Christofori, Heterogeneity for IGF-II production maintained by public goods dynamics in neuroendocrine pancreatic cancer. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 1833 (2015)
20. M.A. Nowak, F. Michor, Y. Iwasa, The linear process of somatic evolution. *Proc. Natl. Acad. Sci. U.S.A.* **100**, 14966 (2003)
21. N.L. Komarova, A. Sengupta, M.A. Nowak, Mutation-selection networks of cancer initiation: tumor suppressor genes and chromosomal instability. *J. Theor. Biol.* **223**, 433 (2003)
22. B. Werner, D. Dingli, A. Traulsen, A deterministic model for the occurrence and dynamics of multiple mutations in hierarchically organized tissues. *J. R. Soc. Interface* **10**, 20130349 (2013)

23. T. Antal, P. Krapivsky, Exact solution of a two-type branching process: models of tumor progression. *J. Stat. Mech.* **2011**, P08018 (2011)
24. Y. Iwasa, F. Michor, M.A. Nowak, Stochastic tunnels in evolutionary dynamics. *Genetics* **166**, 1571 (2004)
25. M.A. Nowak, F. Michor, N.L. Komarova, Y. Iwasa, Evolutionary dynamics of tumor suppressor gene inactivation. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 10635 (2004)
26. P. Ashcroft, F. Michor, T. Galla, Stochastic tunneling and metastable states during the somatic evolution of cancer. *Genetics* **199**, 1213 (2015)
27. Y. Iwasa, F. Michor, N.L. Komarova, M.A. Nowak, Population genetics of tumor suppressor genes. *J. Theor. Biol.* **233**, 15 (2005)
28. F. Michor, Y. Iwasa, Dynamics of metastasis suppressor gene inactivation. *J. Theor. Biol.* **241**, 676 (2006)
29. S.R. Proulx, The rate of multi-step evolution in Moran and Wright-Fisher populations. *Theor. Popul. Biol.* **80**, 197 (2011)
30. H. Haeno, Y.E. Maruvka, Y. Iwasa, F. Michor, Stochastic tunneling of two mutations in a population of cancer cells. *PLoS ONE* **8**, e65724 (2013)
31. J.F. Crow, M. Kimura, *An Introduction to Population Genetics Theory* (Harper and Row, New York, 1970)
32. O. Gottesman, B. Meerson, Multiple extinction routes in stochastic population models. *Phys. Rev. E* **85**, 021140 (2012)
33. O.A. van Herwaarden, J. Grasman, Stochastic epidemics: major outbreaks and the duration of the endemic period. *J. Math. Biol.* **33**, 581 (1995)
34. A. Kamelev, B. Meerson, Extinction of an infectious disease: a large fluctuation in a nonequilibrium system. *Phys. Rev. E* **77**, 061107 (2008)
35. M.I. Dykman, I.B. Schwartz, A.S. Landsman, Disease extinction in the presence of random vaccination. *Phys. Rev. Lett.* **101**, 078101 (2008)
36. A.J. Black, A.J. McKane, WKB calculation of an epidemic outbreak distribution. *J. Stat. Mech.* **2011**, P12006 (2011)
37. L. Billings, L. Mier-Y-Teran-Romero, B. Lindley, I.B. Schwartz, Intervention-based stochastic disease eradication. *PLoS ONE* **8**, e70211 (2013)
38. A. Altland, A. Fischer, J. Krug, I.G. Szendro, Rare events in population genetics: stochastic tunneling in a two-locus model with recombination. *Phys. Rev. Lett.* **106**, 088101 (2011)
39. H.J. Muller, The relation of recombination to mutational advance. *Mut. Res.* **1**, 2 (1964)
40. J.J. Metzger, S. Eule, Distribution of the fittest individuals and the rate of Muller's ratchet in a model with overlapping generations. *PLoS Comput. Biol.* **9**, e1003303 (2013)
41. P.A.P. Moran, *The Statistical Processes of Evolutionary Theory* (Clarendon Press, Oxford UK, 1962)
42. J. Ma, A. Ratan, B.J. Raney, B.B. Suh, W. Miller, D. Haussler, The infinite sites model of genome evolution. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 14254 (2008)
43. T.A. Kunkel, K. Bebenek, DNA replication fidelity. *Annu. Rev. Biochem.* **69**, 497 (2000)
44. D.M. Weinreich, L. Chao, Rapid evolutionary escape by large populations from local fitness peaks is likely in nature. *Evolution* **59**, 1175 (2005)
45. D.B. Weissman, M.M. Desai, D.S. Fisher, M.W. Feldman, The rate at which asexual populations cross fitness valleys. *Theor. Popul. Biol.* **75**, 286 (2009)
46. D.B. Weissman, M.W. Feldman, D.S. Fisher, The rate of fitness-valley crossing in sexual populations. *Genetics* **186**, 1389 (2010)
47. M. Lynch, Scaling expectations for the time to establishment of complex adaptations. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 16577 (2010)
48. W.J. Ewens, *Mathematical Population Genetics I. Theoretical Introduction* (Springer, New York, 2004)
49. T. Antal, I. Scheuring, Fixation of strategies for an evolutionary game in finite populations. *Bull. Math. Biol.* **68**, 1923 (2006)
50. M. Mobilia, Oscillatory dynamics in rock-paper-scissors games with mutations. *J. Theor. Biol.* **264**, 1 (2010)

51. R.A. Fisher, *The Genetical Theory of Natural Selection* (Clarendon Press, Oxford UK, 1930)
52. C.M. Bender, S.A. Orszag, *Advanced Mathematical Methods for Scientists and Engineers* (Springer, New York, 1999)
53. M. Assaf, B. Meerson, Extinction of metastable stochastic populations. *Phys. Rev. E* **81**, 021116 (2010)
54. L.D. Landau, E.M. Lifshitz, *Mechanics*, vol. 1 (Pergamon Press, Oxford UK, 1976)
55. I. Lohmar, B. Meerson, Switching between phenotypes and population extinction. *Phys. Rev. E* **84**, 051901 (2011)
56. M. Heymann, E. Vanden-Eijnden, The geometric minimum action method: a least action principle on the space of curves. *Comm. Pure Appl. Math.* **61**, 1052 (2008)

Chapter 6

The WKB Method: A User-Guide

6.1 Introduction

The Wentzel–Kramers–Brillouin (WKB) method has been used to address a variety of problems in physics and at the interface of biosciences, from problems in optics, quantum mechanics and General Relativity to estimating the lifetime of a disease outbreak. In this chapter we explore the mathematical basis of the method in its application to stochastic processes. The aim of this work is to create a self-contained tutorial that will introduce the reader to the concepts that may be familiar to those who have worked in this field for a long time, but may seem bewildering when looking from the outside. This is currently work in progress, and although an article is in preparation, this is not a complete representation of the document we hope to produce.

The WKB method has its origins in mathematics in the early 19th century. George Green (1793–1841) and Joseph Liouville (1809–1882) first applied the method to find approximate solutions to second order differential equations [1, 2]. It was extended by Harold Jeffreys (1891–1989) to handle turning points [3], before it was popularised by Gregor Wentzel (1898–1978), Hendrik Kramers (1894–1952) and Léon Brillouin (1889–1969), who used the method to find approximate solutions to the Schrödinger equation of quantum mechanics [4–6].

In systems subject to random dynamics, such as those described by stochastic differential equations (SDEs) or Markov jump processes, the WKB method was popularised by Refs. [7–10]. It has been used to compute the statistics of rare events, as reviewed in Ref. [11]. These events could be switching between attractive states [10, 12–15], or reaching an absorbing boundary, such as the extinction of a population [16–21]. As such it is natural to make comparisons with large deviations theory (LDT) [22, 23]. On the other hand, the WKB approach can also be used to make predictions about the stationary states of these stochastic systems [9, 24–26].

The WKB method is built around the expansion of the probability distribution of the stochastic process in terms of a small parameter, which is the amplitude of the ‘noise’ in these systems. The ‘traditional approach’, in the sense that this procedure is followed by the majority of the WKB literature, is to replace the probability distribution with the WKB ansatz, also known as the eikonal approximation¹ [10]. The Fokker–Planck or master equation describing the stochastic process can then be expanded in powers of the small parameter, and the equations obtained at each order can be solved. The resulting leading-order equation in terms of the small parameter is of the form of a Hamilton–Jacobi equation, and when analytic solutions are not possible, it is solved by either the method of characteristics or minimisation schemes. This approach recovers the leading-order contribution to the stationary distribution in terms of the small parameter. However, many questions are left unanswered by this procedure: Why are the characteristic trajectories obtained from the Hamilton–Jacobi equation called ‘most-likely paths’? In systems with absorbing boundaries, how can we define the quasi-stationary distribution? How does this method compare with the theory of large deviations and the concept of quasi-potentials?

In this chapter we will present the method in such a way that shows it is a rigorous mathematical procedure, but it can be understood and applied across a wide range of problems. Through the analysis of examples of increasing complexity, we will define and illustrate the multitude of terms and expressions that appear throughout the literature. We will show that the WKB approach provides an efficient method for computing (quasi-)stationary states in high-dimensional systems. Perhaps most usefully, it allows the construction of quasi-potential landscapes, from which one can extract transition statistics between basins of attraction, such as exit times and most-likely transition paths.

In Sect. 6.2 we introduce a very simple toy model with one degree of freedom which allows us to discuss and compute most-likely paths and quasi-stationary distributions. This model is analysed using a perturbative approach, but this method is limited to cases in which there exists an explicit small parameter in the forward operator of the stochastic process. In Sect. 6.3 we show that the WKB method can be used when the small parameter is the intrinsic noise intensity to generate expressions for the quasi-stationary distribution. In Sect. 6.4 we use a four-state, ‘two-dimensional’ toy model to introduce the concept of a landscape and multiple transition paths. Sect. 6.5 then contains the main body of this chapter, in which we discuss how to generate landscapes and (quasi-)stationary distributions, as well as how the WKB method relates to other approaches, such as path-integral formulations or the theory of large deviations. Finally, we illustrate some of the possible results that can be obtained in Sect. 6.6.

¹This was the approximation scheme used by physicists to investigate wave-scattering phenomena in problems such as optics and quantum mechanics.




Fig. 6.1 One-dimensional asymmetric random walk on a lattice. State 0 is a reflecting boundary; the only transition from state 0 is to state 1. At the opposite end, state N is an absorbing boundary, and when N is reached the process is terminated

6.2 Toy Model with One Degree of Freedom

To introduce some of the terminology and concepts that feature throughout the WKB literature, we will first consider the asymmetric random walk on a one-dimensional lattice, as shown in Fig. 6.1. The state of the system is given by a single stochastic variable, i , which describes the position of a particle on the lattice. This can also be interpreted as a birth–death process in which the birth and death rates are independent of the state of the population. In this process a single particle located at site i ($0 < i < N$) can hop to state $i - 1$ with a rate 1, or to state $i + 1$ with a rate $\varepsilon \ll 1$. Additionally, state 0 is a reflecting boundary such that the only transition is to state 1 with rate ε . Once state N has been reached the process is terminated. State N is effectively an absorbing boundary, i.e. there are no transitions out of this state.

Denoting the probability to be found in state i at time t as $P_i(t)$, the process in Fig. 6.1 is described by the continuous-time master equation

$$\dot{\mathbf{P}}(t) = \mathbb{W} \cdot \mathbf{P}(t), \quad (6.1)$$

where $\mathbf{P} = (P_0, P_1, \dots, P_N)^T$ and $\mathbb{W}$ is the $(N + 1) \times (N + 1)$ transition matrix which has elements $w_{i,i} = -(1 - \delta_{i,0} + \varepsilon)$, $w_{i+1,i} = \varepsilon$ and $w_{i-1,i} = 1$ for $0 \leq i < N$. All other matrix entries are zero. In particular the entire final column of $\mathbb{W}$ is zero, indicating that there are no transitions out of state N .

Most-Likely Path

One feature that is mentioned throughout the WKB literature is the ‘most-likely path’ between two states. In this model we can consider the trajectories connecting the reflecting state 0 to the absorbing state N . As shown in Fig. 6.2, a typical trajectory of our toy model with $N = 5$ spends the majority of the time hopping between states 0 and 1, with frequent visits to state 2, occasional visits to state 3 and very rare visits to state 4. After $\sim 15,000$ state transitions, or a time of $\sim 60,000$ units, this realisation reaches state $N = 5$. But how can we measure the likelihood that we observe this trajectory, and in what space is the probability of a path defined?

The space of trajectories connecting states 0 and N is uncountably infinite. It is uncountable because the transitions can occur after any positive length of time.




Fig. 6.2 Example trajectory from state 0 to state N of the model shown in Fig. 6.1 and described by the master equation (6.1). **a** The full trajectory. **b** The last 100 steps. **c** The last 10 steps. Parameters used are $N = 5$ and $\varepsilon = 10^{-1}$

This space then lends itself to the path-integral description, which we will discuss in Sect. 6.5. We can, however, define the most-likely path in a countable space of trajectories by considering only the state of the system ($i = 0, 1, \dots, N$) and ignoring the time at which the transitions occur. This space is a projection of the former onto the state space, but still path probabilities are not clear due to the large degeneracy of trajectories hopping to and from state 0 multiple times. If a path visits 0 m -times before reaching state N , then there are m paths from 0 to N contained within this single path. To avoid this degeneracy we only focus on the *final* path from 0 to N . That is we compute a realisation of the birth–death process until state N is reached. We then look *from* the final time to the point at which we first reach state 0. In Fig. 6.3a, we show ten realisations of this final path from state 0 to state $N = 5$. Just focusing on the state-space (ignoring time), we see that nine out of ten of the trajectories shown follow the path $0 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5$. We refer to this as the forward-only path. Only one of the paths shown has a backwards step where $i \rightarrow i - 1$.

We can now define our ensemble of paths as those which leave state 0 (and do not return) and reach state N with $k \in \{0, 1, 2, \dots\}$ backwards steps. The most-likely of these paths is then the forward-only path ($k = 0$), as seen in Fig. 6.3b. We can determine how likely it is to observe paths with k backwards steps by considering the probability of stepping forwards or backwards, as we do when we execute the Gillespie algorithm. For a particle at site $0 < i < N$, the probability that the next transition is to $i + 1$ is $q_{i+1,i} = \varepsilon/(1 + \varepsilon)$. The probability that the next transition is to $i - 1$ is $q_{i-1,i} = 1/(1 + \varepsilon)$. As state 0 is a reflecting boundary, we have $q_{1,0} = 1$.

The forward-only path ($k = 0$) from state 0 to state N is observed with probability

$$\Pr(k = 0) = \frac{1}{Z} \prod_{i=0}^{N-1} q_{i+1,i} = \frac{1}{Z} \left(\frac{\varepsilon}{1 + \varepsilon} \right)^{N-1}, \quad (6.2)$$




Fig. 6.3 **a** Examples of final trajectories from state 0 to state N in our toy model, Fig. 6.1. Here the ‘time to absorption’ is defined as $s = t - t_{\text{fix}}$, such that all runs reach state N at $s = 0$. The trajectory marked with a *thick line* is the only one out of the ten samples shown in which a step $i \rightarrow i - 1$ is observed. **b** Histogram of final trajectories with k backwards steps, computed from an ensemble of 10^4 realisations. The crosses are the predictions of Eq. (6.4), where the number of possible paths are $M_0 = 1$, $M_1 = 3$, $M_2 = 8$, $M_3 = 21$, $M_4 = 55$, $M_5 = 144$. We only consider this set of paths when we normalise the probability distribution. Parameters used are $N = 5$ and $\varepsilon = 10^{-1}$

where Z is a normalisation factor for this probability space (to be determined below). A trajectory with a single backward step, as shown in Fig. 6.3a, is observed with probability

$$\begin{aligned}
 \Pr(k = 1) &= \frac{1}{Z} \sum_{j=2}^{N-1} \left\{ \left(\prod_{i=0}^{j-1} q_{i+1,i} \right) \times q_{j-1,j} \times \left(\prod_{i=j-1}^{N-1} q_{i+1,1} \right) \right\} \\
 &= \frac{1}{Z} \sum_{j=2}^{N-1} \left\{ \left(\frac{\varepsilon}{1+\varepsilon} \right)^{j-1} \times \frac{1}{1+\varepsilon} \times \left(\frac{\varepsilon}{1+\varepsilon} \right)^{N-j+1} \right\} \\
 &= \frac{1}{Z} \sum_{j=2}^{N-1} \frac{\varepsilon}{(1+\varepsilon)^2} \left(\frac{\varepsilon}{1+\varepsilon} \right)^{N-1} \\
 &= \frac{1}{Z} (N-2) \frac{\varepsilon}{(1+\varepsilon)^2} \left(\frac{\varepsilon}{1+\varepsilon} \right)^{N-1}. \tag{6.3}
 \end{aligned}$$

We note that the backwards step $1 \rightarrow 0$ is not included in the above sum as we only consider trajectories which leave state 0 and do not return. In general, a path with k backward steps is observed with probability

$$\Pr(k) = \frac{1}{Z} M_k \left(\frac{\varepsilon}{(1+\varepsilon)^2} \right)^k \left(\frac{\varepsilon}{1+\varepsilon} \right)^{N-1}, \tag{6.4}$$

where M_k is the number of possible paths that have k backward steps.² We can calculate the normalisation factor by imposing $\sum_{k=0}^{\infty} \text{Pr}(k) = 1$. The predictions of Eq. (6.4) are in excellent agreement with numerical results, as shown in Fig. 6.3b.

If we were to consider the timing of the transitions, then the probabilities of observing specific paths become zero. Instead we must consider the probability density of paths, which is closely linked with the theory of large deviations and the path-integral framework. These links will be discussed in Sect. 6.5.

Quasi-Stationary Distribution

As $q_{i-1,i} \gg q_{i+1,i}$, we expect the system to be found close to state 0 if it has not already reached the absorbing state. Indeed this is what we see in the time-series in Fig. 6.2. The distribution in which the system is found prior to fixation, which in this case is peaked about state 0, is referred to as the quasi-stationary distribution (QSD). A mathematician would describe this as “*the distribution that is invariant under time-evolution when the process is conditioned on survival (non-absorption)*.” [27]. The system leaks from this distribution to the absorbing state on a very long but finite timescale.

To identify the QSD, we decompose the solution of the master equation (6.1) onto the eigen-basis of the matrix $\mathbb{W}$. This gives

$$\mathbf{P}(t) = \sum_{\alpha=0}^N c_{\alpha} \mathbf{v}^{(\alpha)} e^{\lambda_{\alpha} t}, \quad (6.5)$$

where λ_{α} and $\mathbf{v}^{(\alpha)}$ are eigenvalues and eigenvectors of $\mathbb{W}$, and the c_{α} are coefficients determined by the initial condition. As state N is absorbing, one eigenvalue (λ_0) of $\mathbb{W}$ is zero and the corresponding eigenvector is $\mathbf{v}_i^{(0)} = \delta_{i,N}$. We must have $c_0 = 1$ as $P_N(t \rightarrow \infty) = 1$. All other eigenvalues are negative (and real), as argued in Sect. 4.3. We order these remaining eigenvalues by their magnitudes, such that $|\lambda_1| \leq |\lambda_2| \leq \dots \leq |\lambda_N|$.

In our toy model there is a separation of timescales. This is characterised by a separation of the eigenvalues, $|\lambda_1| \ll |\lambda_2|$, which can be seen in Fig. 6.4a. One eigenvalue (λ_1) behaves as a power of the small parameter ε , whereas the other eigenvalues ($\lambda_{\alpha \geq 2}$) are much larger in magnitude and are almost independent of ε . For times $t \gg |\lambda_2|^{-1}$, all contributions to Eq. (6.5) from the eigenvectors $\mathbf{v}^{(\alpha \geq 2)}$ will be exponentially small. This fast timescale corresponds to the relaxation to the QSD. On this timescale we can approximate the distribution as

$$\mathbf{P}(t) \approx \mathbf{v}^{(0)} + c_1 \mathbf{v}^{(1)} e^{\lambda_1 t}. \quad (6.6)$$

The eigenvalue λ_1 characterises the slow timescale, which corresponds to the leaking to the absorbing state. After a short period of time which satisfies $|\lambda_2|^{-1} \ll t \ll$

²The quantity M_k is found by enumerating the paths with k backwards steps. We have not found a general expression for this number.

$|\lambda_1|^{-1}$, we have $P_N(t) \approx 0$ and $e^{\lambda_1 t} \approx 1$. Thus from Eq. (6.6) we require $c_1 v_N^{(1)} = -v_N^{(0)} = -1$. Without loss of generality, the coefficient c_1 can be set to unity, and we change the normalisation of the eigenvector $\mathbf{v}^{(1)}$ such that $v_N^{(1)} = -1$. Using $v_i^{(0)} = \delta_{i,N}$, we can now express $\mathbf{P}(t)$ as

$$P_{0 \leq i \leq N-1}(t) \approx v_i^{(1)} e^{\lambda_1 t}, \quad P_N(t) \approx 1 - e^{\lambda_1 t}. \quad (6.7)$$

This is exactly Eq. (5.11) from the previous chapter. As the QSD, $\mathbf{P}^*$, is conditioned on non-absorption, we can write $P_i^* = (1 - \delta_{i,N}) v_i^{(1)}$ for $0 \leq i \leq N$. Thus, the QSD is given by the eigenvector of the master equation which corresponds to the slowest eigenvalue.

We can now determine the scaling of the slowest eigenvalue with the small parameter ε by considering the mean fixation time. First, using Eq. (6.7), we observe that the mean fixation time is given by

$$\langle t_{\text{fix}} \rangle = \int_0^\infty t \dot{P}_N(t) dt \approx -\lambda_1 \int_0^\infty t e^{\lambda_1 t} dt = -\lambda_1^{-1}. \quad (6.8)$$

Second, we can calculate the mean fixation time from the backward master equation, as described in Sect. 2.4. Following the derivation of Eq. (2.25), the mean fixation time conditioned on starting in state $i < N$, t_i , is given by

$$t_i = \sum_{k=i}^{N-1} \varepsilon^{-(k+1)} \left(1 + \sum_{\ell=1}^k \varepsilon^\ell \right). \quad (6.9)$$

For all initial conditions $i < N$, there is a term $\mathcal{O}(\varepsilon^{-N})$. Hence, for small ε , we arrive at $\langle t_{\text{fix}} \rangle \sim \varepsilon^{-N}$. Equating this with Eq. (6.8), we can conclude that $\lambda_1 \sim -\varepsilon^N$. This scaling is shown in Fig. 6.4a.




Fig. 6.4 **a** Absolute value of the non-zero eigenvalues of the matrix $\mathbb{W}$ in Eq. (6.1), along with the predicted scaling of the slowest eigenvalue, $|\lambda_1| \sim \varepsilon^N$. **b** Elements of the eigenvector $\mathbf{v}^{(1)}$ (or equivalently the QSD) are approximately ε^i , in agreement with our perturbative treatment. Eigenvectors are scaled such that $v_N^{(1)} = -1$. In this illustration we have used $N = 5$

The eigenvector $\mathbf{v}^{(1)}$, and hence the QSD, can be approximated by considering a perturbative expansion of the forward operator $\mathbb{W} = \mathbb{W}^{(0)} + \varepsilon \mathbb{W}^{(1)}$ and of the probability $\mathbf{P}(t) = \sum_{k=0}^{\infty} \varepsilon^k \mathbf{P}^{(k)}(t)$. At leading order, $\mathcal{O}(\varepsilon^0)$, one recovers

$$\dot{\mathbf{P}}^{(0)}(t) = \mathbb{W}^{(0)} \cdot \mathbf{P}^{(0)}(t) \Rightarrow P_i^{(0)}(t) = \exp[\mathbb{W}^{(0)}t] \cdot \mathbf{P}^{(0)}(0). \quad (6.10)$$

This is the solution to the master equation in which particles can only hop from state i to state $i - 1$ with a rate 1 (this is the process described by the matrix $\mathbb{W}^{(0)}$). Starting with the initial condition $i_0 = 0$, no dynamics can happen and hence $P_i^{(0)}(t) = \delta_{i,0}$. For initial conditions $i_0 > 1$, we see that $P_{0 < i \leq i_0}^{(0)}(t) = t^{i_0-i} e^{-t} / (i_0 - i)! \sim e^{-t}$, which quickly collapses to $\delta_{i,0}$. For this reason we only consider the initial condition $i_0 = 0$, and hence $P_i^{(0)}(t) = \delta_{i,0}$ for $t \geq 0$. By considering the higher-order equations in terms of ε , which are of the form

$$\dot{\mathbf{P}}^{(k)}(t) = \mathbb{W}^{(0)} \cdot \mathbf{P}^{(k)}(t) + \mathbb{W}^{(1)} \cdot \mathbf{P}^{(k-1)}(t), \quad (1 \leq k < N) \quad (6.11)$$

we can recover $P_i^{(i)}(t) \sim \varepsilon^i$ for $1 \leq i < N$. Hence, from our definition of the QSD we obtain $v_i^{(1)} \sim \varepsilon^i$ for $0 \leq i < N$. This scaling is shown in Fig. 6.4b.

Although this perturbative approach is very intuitive, its applications are limited. This is because the method relies on the presence of an explicit small parameter in the forward operator $\mathbb{W}$. In many processes, as we will see in the next section, there are no explicit small parameters in the reaction scheme. Instead it is the amount of noise that is the small parameter of interest. If this is the case, how do we extract the quasi-stationary distribution? For this we introduce the WKB method.

6.3 The WKB Method in One Dimension

The WKB method is most easily illustrated by considering problems described by a single stochastic variable. In these one-dimensional scenarios we are able to make analytical progress and obtain closed-form expressions which approximate the quasi-stationary distribution (QSD). We apply the WKB method to two distinct classes of problem: those described by a Fokker–Planck equation, and the familiar individual based-models described by a master equation. As discussed in Sect. 2.7, the Fokker–Planck equation can be used to approximately describe an individual-based model. However, we will first consider it here as the outright description of a continuous stochastic process. Comparisons with results obtained from the master equation will be made at the end of the section. We will start with the Fokker–Planck equation as the mathematical motivation for the WKB method is more clear in the differential equation framework. The treatment of the master equation is then effectively a generalisation of this method to jump processes.

WKB and the Fokker–Planck Equation

The evolution of a continuous variable that is affected by noise can be described by a Fokker–Planck equation [28]. If the strength of this noise is given by a parameter $\varepsilon \ll 1$, then the Fokker–Planck equation is given by

$$\dot{\rho}(x, t) = -\frac{\partial}{\partial x} [A(x)\rho(x, t)] + \frac{\varepsilon}{2} \frac{\partial^2}{\partial x^2} [B(x)\rho(x, t)], \quad (6.12)$$

where $\rho(x, t)$ is the probability density function for the continuous variable x at time t . The drift term $A(x)$ describes the deterministic evolution of x , and the diffusion term $B(x) > 0$ describes the influence of the noise.

For simplicity we assume that the drift term permits a unique stable fixed point at $x^* > 0$, i.e. $A(x^*) = 0$. The scenario of multiple stable states will be discussed in Sect. 6.5. We also assume that there is an absorbing boundary located $x = 0$.³ Considering the deterministic dynamics, or using the arguments of the previous section, the system will relax to the QSD about the stable fixed point x^* . The escape to the absorbing state is a rare event, and can be pictured as the leaking of the probability density function from the QSD to $x = 0$. Assuming this leaking timescale is very long, we can determine the QSD by setting $\dot{\rho}(x, t) \approx 0$ in Eq. (6.12). We can also arrive at the condition $\dot{\rho}(x, t) \approx 0$ by replacing the probability density function with an ansatz of the form of Eq. (6.7). This approach was described in the previous chapter for jump processes. We label the QSD as $\rho^*(x)$, which satisfies the second-order ODE

$$0 \approx -\frac{d}{dx} [A(x)\rho^*(x)] + \frac{\varepsilon}{2} \frac{d^2}{dx^2} [B(x)\rho^*(x)]. \quad (6.13)$$

This equation is singularly perturbative in the small parameter ε ; the behaviour of the solution in the $\varepsilon \rightarrow 0$ limit cannot be approximated by setting $\varepsilon = 0$ and then solving Eq. (6.13) [30]. To analyze the behaviour of the solution in the small- ε limit, one considers an ansatz of the form $\rho^*(x) \sim \exp[-\sum_{\sigma=0}^{\infty} \delta^{\sigma-1} S_{\sigma}(x)]$. Here $\delta > 0$ is a small parameter that contains the ε -dependence of the solution, such that the $S_{\sigma}(x)$ are independent of both δ and ε . The parameter δ acts as a rescaling of the amplitude of $\rho^*(x)$, allowing us to zoom in to regions where there is exponential or dissipative behaviour [30]. In our process, there is a rapid *dissipation*, or fast decay, of the probability density away from the stable fixed point. To determine the relation between δ and ε , we substitute the ansatz into Eq. (6.13). Applying the product rule for differentiation gives

³We assume that this boundary is ‘regular’, i.e. the probability of reaching the boundary is non-zero and the expected time to reach the boundary is finite [29].

$$\begin{aligned}
0 = & -A'(x) + A(x) \left[\sum_{\sigma=0}^{\infty} \delta^{\sigma-1} S'_{\sigma}(x) \right] + \frac{\varepsilon}{2} B''(x) - \varepsilon B'(x) \left[\sum_{\sigma=0}^{\infty} \delta^{\sigma-1} S'_{\sigma}(x) \right] \\
& - \frac{\varepsilon}{2} B(x) \left[\sum_{\sigma=0}^{\infty} \delta^{\sigma-1} S''_{\sigma}(x) \right] + \frac{\varepsilon}{2} B(x) \left[\sum_{\sigma=0}^{\infty} \delta^{\sigma-1} S'_{\sigma}(x) \right]^2, \quad (6.14)
\end{aligned}$$

where we have used the notation $\mathcal{F}'(x) = d\mathcal{F}/dx$, and we assume the fields $S_{\sigma}(x)$, as well as the drift and diffusion terms, are smooth. As $A(x)$, $B(x)$, and $S_{\sigma}(x)$ are independent of the parameters δ and ε , we can be sure that the largest terms in this expression are $\delta^{-1}A(x)S'_0(x)$ and $\delta^{-2}\varepsilon B(x)[S'_0(x)]^2/2$. Through dominant balance these terms must be of the same order and cancel each other out, hence $\delta = \mathcal{O}(\varepsilon)$ is required. For simplicity we choose $\delta = \varepsilon$, such that the QSD is of the form

$$\rho^*(x) \sim \exp \left[- \sum_{\sigma=0}^{\infty} \varepsilon^{\sigma-1} S_{\sigma}(x) \right]. \quad (6.15)$$

The values of $S_{\sigma}(x)$ are then found by substituting the ansatz (6.15) into Eq. (6.13) [i.e. replace δ with ε in Eq. (6.14)], and considering equations at different powers of ε . Taking only the leading-order terms in ε gives

$$A(x)p(x) + \frac{1}{2}B(x)p^2(x) = 0, \quad (6.16)$$

where we have introduced $p(x) = S'_0(x)$. Equation (6.16) is quadratic in $p(x)$. One solution is simply $p(x) = 0$, and the other is $p(x) = -2A(x)/B(x)$. Integrating this second solution gives

$$S_0(x) = - \int_{x^*}^x \frac{2A(y)}{B(y)} dy, \quad (6.17)$$

where the arbitrary constant of integration is chosen such that $S_0(x^*) = 0$. A similar procedure can be used to find the next-leading-order correction, $S_1(x) = \ln[B(x)/B(x^*)]$. The approximate solution to Eq. (6.13) is then

$$\rho^*(x) \sim \frac{\mathcal{N}}{B(x)} \exp \left[\frac{1}{\varepsilon} \int_{x^*}^x \frac{2A(y)}{B(y)} dy \right], \quad (6.18)$$

where $\mathcal{N}$ is an overall normalisation constant.

If the system does not have an absorbing state at $x = 0$, then Eq. (6.18) is the exact stationary distribution of the Fokker–Planck equation found by setting the probability current to zero [31], i.e. $\rho^*(x)$ is given by the solution of

$$J(x) = A(x)\rho^*(x) - \frac{\varepsilon}{2} \frac{d}{dx} [B(x)\rho^*(x)] = 0. \quad (6.19)$$

WKB and the Master Equation

We can now describe the application of this method to individual-based stochastic processes that are described by a discrete variable, n , which can take the values $0 \leq n < \infty$. If there are M possible reactions, and the r -th reaction occurs with rate $T_n^{v_r}$ and has stoichiometric effect v_r on the population, then the master equation can be written as

$$\dot{P}_n(t) = \sum_{r=1}^M \left[T_{n-v_r}^{v_r} P_{n-v_r}(t) - T_n^{v_r} P_n(t) \right] = \sum_{r=1}^M (\mathcal{E}^{-v_r} - 1) T_n^{v_r} P_n(t), \quad (6.20)$$

where $\mathcal{E}^{v_r}$ is the shift operator: $\mathcal{E}^{v_r} \mathcal{F}(n) = \mathcal{F}(n + v_r)$ [29]. We assume the system has a typical (large) number of individuals, Ω , which is referred to as the system size. This parameter determines the influence of the intrinsic noise in these discrete processes [29]. As in the Fokker–Planck scenario, we will assume there exists a unique stable fixed point in the deterministic dynamics at x^* , where $x = \lim_{\Omega \rightarrow \infty} n/\Omega$, and that there exists an absorbing boundary at $n = 0$.

The WKB treatment of Eq. (6.20) was first achieved in Ref. [9] by replacing the shift operators in Eq. (6.20) by their continuous counterparts

$$\mathcal{E}^{v_r} \rightarrow \exp \left[\frac{v_r}{\Omega} \frac{\partial}{\partial x} \right] = \sum_{\ell=0}^{\infty} \frac{1}{\ell!} \left(\frac{v_r}{\Omega} \right)^{\ell} \frac{\partial^{\ell}}{\partial x^{\ell}}, \quad (6.21)$$

which is valid provided that the transition rates vary smoothly between states. The master equation in the continuum limit can then be written as

$$\frac{1}{\Omega} \dot{\rho}(x, t) = \sum_{r=1}^M \sum_{\ell=1}^{\infty} \frac{1}{\ell!} \left(\frac{-v_r}{\Omega} \right)^{\ell} \frac{\partial^{\ell}}{\partial x^{\ell}} [f_r(x) \rho(x, t)], \quad (6.22)$$

where $\rho(x, t) = \Omega P_{\Omega x}(t)$, and $f_r(x) = T_{\Omega x}^{v_r}/\Omega$ as before.⁴

As can be seen, derivatives of order ℓ are multiplied by the small parameter $\Omega^{-\ell}$ and Eq. (6.22) is singularly perturbative. To characterise the behaviour, the ansatz (6.15) has previously been employed [9], where we now replace the noise strength ε with the inverse system size, Ω^{-1} . Setting $\dot{\rho}(x, t) \approx 0$ in Eq. (6.22), we seek to solve

$$0 \approx \sum_{r=1}^M \sum_{\ell=1}^{\infty} \frac{1}{\ell!} \left(\frac{-v_r}{\Omega} \right)^{\ell} \frac{d^{\ell}}{dx^{\ell}} \left\{ f_r(x) \exp \left[- \sum_{\sigma=0}^{\infty} \Omega^{1-\sigma} S_{\sigma}(x) \right] \right\}. \quad (6.23)$$

⁴We could expand the reactions rates in further powers of Ω , such that $T_{\Omega x}^r = \Omega f_r(x) + g_r(x) + h_r(x)/\Omega + \dots$, as described in Ref. [19]. However in this section we only consider the leading-order contributions.

At leading order in the system size we have the expression

$$\sum_{r=1}^M w_r(x) (e^{v_r p(x)} - 1) = 0, \quad (6.24)$$

where again $p(x) = S'_0(x)$. Equation (6.16) can be recovered by replacing the exponential with $1 + v_r p(x) + v_r^2 p^2(x)/2$. However this approximation is not necessarily justified as we cannot assume terms $p^3(x)$ or greater are negligible.

For the case in which the process can only step from state i to $i \pm 1$, i.e. $v = \pm 1$, Eq. (6.24) can be solved to find $p(x)$. Similar to Eq. (6.16), there are two possible solutions, $p(x) = 0$ and $p(x) = \ln[f_-(x)/f_+(x)]$. Integrating this second solution leads to

$$S_0(x) = \int_{x^*}^x \ln \left(\frac{f_-(y)}{f_+(y)} \right) dy, \quad (6.25)$$

where again the constant of integration is defined by $S_0(x^*) = 0$. This gives the leading-order contribution in the system size to the QSD as

$$\rho^*(x) \sim \mathcal{N} \exp \left[\Omega \int_{x^*}^x \ln \left(\frac{f_+(y)}{f_-(y)} \right) dy \right], \quad (6.26)$$

where $\mathcal{N}$ is an overall normalisation constant. This is the same expression we obtained in the previous chapter when we discussed the escape from the metastable state located on the state-space boundary. In that case we estimated $\mathcal{N}$ by considering a Gaussian approximation to the QSD about the fixed point x^* .

Equation (6.26) can be compared with the stationary solution of Eq. (6.20) in the absence of the absorbing state [29, 32].⁵ From Eq. (2.54), this stationary solution is given by

$$P_n(\infty) = \mathcal{N} \prod_{i=2}^n \frac{T_{i-1}^+}{T_i^-} = \mathcal{N} \exp \left[\sum_{i=2}^n \ln \left(\frac{T_{i-1}^+}{T_i^-} \right) \right]. \quad (6.27)$$

Comparison

We can now compare the results of the WKB treatment of the master equation with those obtained from the Fokker–Planck equation that approximately describes the same individual-based process. The construction of the Fokker–Planck equation is described in Sect. 2.7, and for the birth–death process we have $A(x) = f_+(x) - f_-(x)$, $B(x) = f_+(x) + f_-(x)$, and $\varepsilon = \Omega^{-1}$. We summarise the WKB results in Table 6.1. As pointed out in Ref. [33], the values of $S_0(x)$ in the Fokker–Planck and master equation formalisms are not in exact agreement with each other. These differences will be inflated in the QSD, $\rho^*(x)$, where $S_0(x)$ features in the exponent. However, at the deterministic fixed point x^* , $S_0(x^*)$ and its first and second derivatives

⁵The absorbing state can be removed by setting $T_1^- = 0$.

Table 6.1 Comparison of results of the WKB approach from the Fokker–Planck and master equations for the birth–death individual-based process ($\nu = \pm 1$)


	Fokker–Planck	Master equation
$p(x)$	$-\frac{2A(x)}{B(x)}$	$\ln\left(\frac{f_-(x)}{f_+(x)}\right)$
$S_0(x)$	$-\int_{x^*}^x \frac{2A(y)}{B(y)} dy$	$\int_{x^*}^x \ln\left(\frac{f_-(y)}{f_+(y)}\right) dy$
$\rho^*(x)$	$\sim \exp\left[\Omega \int_{x^*}^x \frac{2A(y)}{B(y)} dy\right]$	$\sim \exp\left[-\Omega \int_{x^*}^x \ln\left(\frac{f_-(y)}{f_+(y)}\right) dy\right]$

We have replaced the small parameter ε with Ω^{-1} in the Fokker–Planck results, and the drift and diffusion terms are given by $A(x) = f_+(x) - f_-(x)$ and $B(x) = f_+(x) + f_-(x)$, respectively

are the same between the two formalisms [33]. Thus, close to the fixed point, we expect the functions to be very similar. This agreement has also been pointed out in Refs. [13, 31].

Example: Logistic Population Growth

An illustrative example of a stochastic process which features a quasi-stationary distribution is the well-studied birth/death/competition process [19, 34–37], which is based on the Verhulst model of population growth [38].⁶ This process is described by the reactions



where A represents an individual, α , β , and γ are rates and Ω is the (large) system-size parameter. The state of the system is described by the stochastic variable n , which is the number of individuals present in the population. The time evolution of the probability distribution of n is described by the master equation (6.20) with transition rates

$$T_n^+ = \alpha n, \quad (6.29a)$$

$$T_n^- = \beta n + \frac{\gamma n(n-1)}{\Omega}, \quad (6.29b)$$

where we have combined the death and competition reactions in Eq. (6.28) as they have the same stoichiometric coefficient, $\nu = -1$.

The deterministic dynamics of this model for the concentration $x = \lim_{\Omega \rightarrow \infty} n/\Omega$, satisfy

$$\dot{x} = (\alpha - \beta)x \left[1 - \frac{\gamma}{\alpha - \beta} x \right]. \quad (6.30)$$

This is simply the logistic equation with a stable fixed point at $x^* = (\alpha - \beta)/\gamma$ (provided $\alpha > \beta$) and an unstable (but absorbing) fixed point at $x = 0$.

In Fig. 6.5 we show the leading-order contributions to the QSDs in terms of the system size, as described in Table 6.1. The distribution obtained from simulations is in

⁶Pierre François Verhulst (1804–1849).




Fig. 6.5 b The quasi-stationary distributions (QSDs) calculated from the master equation (*solid line*) and from the Fokker–Planck equation (*dashed line*), as described in Table 6.1. These are compared with the QSD calculated from simulations of the individual-based model (*dots*). Simulations are initialised at the fixed point and run until a fixed time ($t = 10^4$). The final state is then recorded, and runs that have reached fixation are discarded. The *dots* are a histogram of the remaining data points from 10^7 realisations. *Inset* plot shows the same data on a linear vertical axis. Parameters used are $\alpha = 1.0$, $\beta = 0.05$, $\gamma = 1.0$, $\Omega = 20$

very good agreement with these results. As the distance from the fixed point increases, so does the deviation between results from the master equation and Fokker–Planck approximation.

It is worth noting here that naively we expect the results obtained from the Fokker–Planck approximation to become increasingly accurate as the system size is increased. This is true in the bulk of the distribution close to the fixed point x^* . However, in the tails of the distribution the separation between the Fokker–Planck results and those obtained from the master equation grow exponentially with the system size. Although the absolute difference between the two quasi-stationary distributions is decreasing due to narrowing of the tails as Ω increases, the relative difference can grow to several orders of magnitude.

6.4 Four-State Toy Model

To further illustrate the concepts of quasi-stationary distributions and most-likely paths, we introduce the four-state toy model shown in Fig. 6.6. The system starts in state 0, and will jump to either state 1 or 2 with equally small transition rates $\varepsilon \ll 1$. From these states there is a large transition rate back to state 0, or a small transition rate ($f_1\varepsilon$ or $f_2\varepsilon$, respectively) to state 3 which is absorbing. There are now two separate routes to the absorbing state, either proceeding through state 1 or state 2, and hence it can be thought of as a process in two dimensions (even though it can be described by a single stochastic variable).

As the reaction rates have a similar structure to the one-dimensional toy model introduced in Sect. 6.2, it is clear that prior to absorption the system spends the majority of the time in state 0. This is seen in the sample trajectory shown in Fig. 6.7. The master equation describing the process in Fig. 6.6 can be written as





Fig. 6.6 Illustration of the four-state, ‘two-dimensional’ toy model. The *arrows* indicate the possible transitions between the states. *Bold arrows* indicate that those transitions occur with a larger rate. *No arrows* out of state 3 indicate that it is an absorbing state

Fig. 6.7 Time series of the toy model shown in Fig. 6.6. The *spikes* indicate that time spent in states 1 and 2 is very small. The absorbing state is reached through state 2 in this simulation. The parameters used are $\varepsilon = 10^{-2}$, $f_1 = 1$, and $f_2 = 3$



$$\begin{pmatrix} \dot{P}_0 \\ \dot{P}_1 \\ \dot{P}_2 \\ \dot{P}_3 \end{pmatrix} = \begin{pmatrix} -2\varepsilon & 1 & 1 & 0 \\ \varepsilon & -(1 + f_1\varepsilon) & 0 & 0 \\ \varepsilon & 0 & -(1 + f_2\varepsilon) & 0 \\ 0 & f_1\varepsilon & f_2\varepsilon & 0 \end{pmatrix} \cdot \begin{pmatrix} P_0 \\ P_1 \\ P_2 \\ P_3 \end{pmatrix}. \quad (6.31)$$

Through a simple perturbative expansion in the small parameter ε (similar to the procedure described in Sect. 6.2) it can be shown that the probability to be found in state 0 (prior to fixation) is $1 - 2\varepsilon + \mathcal{O}(\varepsilon^2)$ and the probability to be found in either of states 1 or 2 (prior to fixation) is $\varepsilon + \mathcal{O}(\varepsilon^2)$. This is the quasi-stationary distribution.

Once the system reaches the absorbing state, its final trajectory from state 0 to state 3 will be $0 \rightarrow 1 \rightarrow 3$ with probability $f_1/(f_1 + f_2)$ and $0 \rightarrow 2 \rightarrow 3$ with probability $f_2/(f_1 + f_2)$. In Fig. 6.7 we have used $f_1 = 1$ and $f_2 = 3$. Hence we expect the first path through state 1 to be realised 25 % of the time, and the second path through state 2 to be realised 75 % of the time. This is confirmed in Fig. 6.8a.

The arrival times at the absorbing state are determined by the slowest eigenvalue of the matrix in Eq. (6.31), λ_1 . This eigenvalue controls the leaking from the QSD, as described in Sect. 6.2, and we expect λ_1 to scale as ε^2 (as we require two consecutive steps of rate $\sim \varepsilon$). To adjust for the increased rate of reaching the absorbing state, we expect the eigenvalue to be of the form $\lambda_1 \approx -c(f_1 + f_2)\varepsilon^2$, for constant c . The approximation for the arrival time density is then $\dot{P}_3(t) \sim e^{\lambda_1 t}$. The accuracy of this approximation is confirmed in Fig. 6.8b, where we find $c \approx 1$.




Fig. 6.8 **a** The probability of observing each of the path to absorption, obtained from 10^5 simulations of the process shown in Fig. 6.6. **b** The arrival time distribution at the absorbing state from the same data used in (a). The leading eigenvalue of the matrix in Eq. (6.31) scales as $(f_1 + f_2)\varepsilon^2 + \mathcal{O}(\varepsilon^3)$. The parameters are the same as Fig. 6.7, i.e. $\varepsilon = 10^{-2}$, $f_1 = 1$, and $f_2 = 3$

In more general systems which do not exhibit the small parameter in the forward operator, we can again use the WKB method to predict the quasi-stationary distribution and also most-likely paths.

6.5 The WKB Method in Higher Dimensions

For stochastic processes described by two or more variables, analytic progress is almost impossible.⁷ If the number of stochastic variables in the system is given by $d \geq 2$, then we call this a d -dimensional system. As described in Sect. 6.3, the small parameter in these systems corresponds to the strength of the noise, rather than explicitly appearing in the dynamics. Following Sect. 6.3, we start by analysing a continuous stochastic process described by a Fokker–Planck equation, before generalising the approach to discrete jump processes.

An important consideration, which will be illustrated through an example at the end of this chapter, is whether there exist multiple stable states in our system. In the one-dimensional system analysed in Sect. 6.3, we restricted ourselves to the case of a single stable state, and hence we had a unique QSD. If there are multiple stable states, labelled $\{\mathbf{x}_1^*, \mathbf{x}_2^*, \dots\}$, then there are a corresponding number of QSDs. The distribution of probability mass between these states on the fast relaxation timescale is determined by the initial condition, and on longer timescales by the leaking from one state to another. If we start in the basin of attraction of $\mathbf{x}_a^*$, then we will initially relax to the QSD about this stable state. As time progresses, the probability slowly (i.e. more slowly than the relaxation timescale) leaks to neighbouring stable states. The quantities of interest are now the QSDs about each of the stable states, the time to

⁷Analytical solutions are available for the two-type branching process [39], but, in general, such a description is usually lacking.

escape from these states, the most-likely transition paths, and the overall stationary state. Note that the scenario of systems with absorbing states is included in this framework.

The Fokker–Planck equation in d -dimensions, with a small noise strength $\varepsilon \ll 1$, is given by

$$\dot{\rho}(\mathbf{x}, t) = - \sum_i \frac{\partial}{\partial x_i} [A_i(\mathbf{x})\rho(\mathbf{x}, t)] + \frac{\varepsilon}{2} \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} [B_{ij}(\mathbf{x})\rho(\mathbf{x}, t)], \quad (6.32)$$

where $\rho(\mathbf{x}, t)$ is the probability density function at time t for the continuous variables $\mathbf{x} = (x_1, \dots, x_d)^T$. The drift term $\mathbf{A}(\mathbf{x})$ describes the deterministic evolution $\mathbf{x}$, and the diffusion matrix $B_{ij}(\mathbf{x})$ describes the influence of the noise, including the correlations of the noise between the stochastic variables which are described by the off-diagonal terms of this matrix [28].

To find the QSD about the stable state $\mathbf{x}_a^*$, we restrict the process to the basin of attraction of this state and assume that it is approximately stationary such that $\dot{\rho}(\mathbf{x}, t) \approx 0$. We can then write Eq. (6.32) as

$$0 \approx - \sum_i \frac{\partial}{\partial x_i} [A_i(\mathbf{x})\rho_a^*(\mathbf{x})] + \frac{\varepsilon}{2} \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} [B_{ij}(\mathbf{x})\rho_a^*(\mathbf{x})], \quad (6.33)$$

where $\rho_a^*(\mathbf{x})$ is the QSD about the stable state $\mathbf{x}_a^*$. Equation (6.33) is singularly perturbative as the small parameter ε multiplies the highest-order derivatives. As a result, the quasi-stationary solution ansatz will be analogous to the one-dimensional scenario, Eq. (6.15). Thus within the basin of attraction of $\mathbf{x}_a^*$, the QSD takes the form

$$\rho_a^*(\mathbf{x}) \sim \exp \left[- \sum_{\sigma=0}^{\infty} \varepsilon^{\sigma-1} S_{\sigma}^{(a)}(\mathbf{x}) \right]. \quad (6.34)$$

Substituting this ansatz into Eq. (6.32), and taking leading-order terms in the small parameter ε , we arrive at

$$H^{(\text{FPE})}(\mathbf{x}, \mathbf{p}) = \mathbf{A}(\mathbf{x}) \cdot \mathbf{p} + \frac{1}{2} \mathbf{p} \cdot \mathbb{B}(\mathbf{x}) \cdot \mathbf{p} = 0, \quad (6.35)$$

where $\mathbf{p} = \nabla S_0^{(a)}(\mathbf{x})$. The one-dimensional example is included in this framework. However, in $d \geq 2$ dimensions we cannot directly identify a non-zero $\mathbf{p}$ which satisfies Eq. (6.35). In fact the problem is non-integrable and underdetermined; we have $2d$ degrees of freedom ($x_1, \dots, x_d; p_1, \dots, p_d$), but only one constraint ($H = 0$). The equation is labelled by $H^{(\text{FPE})}(\mathbf{x}, \mathbf{p})$ as it is a Hamilton–Jacobi equation; it is a first-order differential equation of the form $H[\mathbf{x}, \nabla S_0^{(a)}(\mathbf{x})] = 0$ [40].

For individual-based models in $d \geq 2$ dimensions, the master equation can again be expressed in terms of the shift operators as

$$\dot{P}_{\mathbf{n}}(t) = \sum_r [T_{\mathbf{n}-\mathbf{v}_r}^{\mathbf{v}_r} P_{\mathbf{n}-\mathbf{v}_r}(t) - T_{\mathbf{n}}^{\mathbf{v}_r} P_{\mathbf{n}}(t)] = \sum_r (\mathcal{E}^{-\mathbf{v}_r} - 1) T_{\mathbf{n}}^{\mathbf{v}_r} P_{\mathbf{n}}(t), \quad (6.36)$$

where $\mathbf{n} = (n_1, \dots, n_d)^T$ and $\mathbf{v}_r$ is a d -dimensional vector that describes the change of state due to reaction r . We assume there exists a large parameter Ω that characterises the typical system size, and hence the strength of the noise through $\varepsilon \rightarrow \Omega^{-1}$. With this we can write the master equation in the continuum limit by replacing the shift operators with differentials as described in Sect. 6.2. The analysis proceeds as before, and we arrive at the leading-order equation in the system size

$$H^{(\text{ME})}(\mathbf{x}, \mathbf{p}) = \sum_r f_r(\mathbf{x}) (e^{\mathbf{v}_r \cdot \mathbf{p}} - 1) = 0, \quad (6.37)$$

where $\mathbf{p} = \nabla S_0^{(a)}(\mathbf{x})$ is defined in the basin of attraction of the stable state $\mathbf{x}_a^*$.

Solving the Hamilton–Jacobi Eq. (6.37) [or Eq. (6.35)] gives the function $S_0^{(a)}(\mathbf{x})$, and this allows us to compute the dominant contribution to the QSDs in terms of the system size. The solution to these equations can be found using the method of characteristics, as will be described below. For the transitions between the stable states, we must construct trajectories from one basin to another. These paths can be computed by considering the path-integral approach to the stochastic process, as will be discussed below. Once the functions $S_0^{(a)}(\mathbf{x})$ have been evaluated, we want to combine these objects in such a way that we create a landscape over all of the basins of attraction, which spans the full domain. To do this we match at the $S_0^{(a)}(\mathbf{x})$ at the separatrix [10]. This means we must shift the $S_0^{(a)}(\mathbf{x})$ up or down by an additive constant. This is analogous to matching potential wells, and hence the matched function $S_0(\mathbf{x})$ is like a potential landscape. In large deviations theory, the quantity $S_0(\mathbf{x})$ is called the quasi-potential. We explore this link further in this section. The stationary distribution of the process is then characterised by $S_0(\mathbf{x})$, i.e. $\rho^{\text{st}}(\mathbf{x}) \sim \exp[-\Omega S_0(\mathbf{x})]$.

Characteristic Solutions

To obtain the solutions of the Hamilton–Jacobi equations, we construct a $2d$ -dimensional space $(x_1, \dots, x_d; \pi_1, \dots, \pi_d)$. The function $H(\mathbf{x}, \boldsymbol{\pi}) = 0$ [either Eq. (6.35) or Eq. (6.37)] specifies a $2d - 1$ -dimensional surface in this space [8]. We want to find the values of $\boldsymbol{\pi}$ which satisfy $H = 0$. We define these values as $\mathbf{p}(\mathbf{x}) = \nabla S_0^{(a)}(\mathbf{x})$. We can use the method of characteristics to identify each $p_i(\mathbf{x})$. As $dH/dx_i = 0$ on the surface $H = 0$, we can write

$$\begin{aligned} \frac{dH}{dx_i} &= \frac{\partial H}{\partial x_i} + \sum_j \frac{\partial H}{\partial \pi_j} \frac{\partial p_j}{\partial x_i} \\ &= \frac{\partial H}{\partial x_i} + \sum_j \frac{\partial H}{\partial \pi_j} \frac{\partial p_i}{\partial x_j} = 0. \end{aligned} \quad (6.38)$$

In the last step we have used $\partial p_j / \partial x_i = \partial^2 S_0^{(a)} / \partial x_i \partial x_j = \partial p_i / \partial x_j$.

We can construct a surface $\pi_i = p_i(\mathbf{x})$ in the $d + 1$ -dimensional space $(x_1, \dots, x_d, \pi_i) = \Lambda_i$. A normal to this surface in the space Λ_i is

$$\begin{aligned}\boldsymbol{\psi}_{\perp}^{(i)} &= \left(\frac{\partial}{\partial x_1}, \dots, \frac{\partial}{\partial x_d}, \frac{\partial}{\partial \pi_i} \right) [p_i(\mathbf{x}) - \pi_i] \\ &= \left(\frac{\partial p_i}{\partial x_1}, \dots, \frac{\partial p_i}{\partial x_d}, -1 \right).\end{aligned}\quad (6.39)$$

A tangent vector along this surface, $\boldsymbol{\psi}_{\parallel}^{(i)}$, satisfies

$$0 = \boldsymbol{\psi}_{\parallel}^{(i)} \cdot \boldsymbol{\psi}_{\perp}^{(i)} = \sum_{j=1}^d \left[\boldsymbol{\psi}_{\parallel}^{(i)} \right]_j \frac{\partial p_i}{\partial x_j} - \left[\boldsymbol{\psi}_{\parallel}^{(i)} \right]_{d+1}. \quad (6.40)$$

We can then use Eq. (6.38) to see that the components

$$\left[\boldsymbol{\psi}_{\parallel}^{(i)} \right]_{1 \leq j \leq d} = \frac{\partial H}{\partial \pi_j}, \quad \left[\boldsymbol{\psi}_{\parallel}^{(i)} \right]_{d+1} = -\frac{\partial H}{\partial x_i}, \quad (6.41)$$

fulfil the orthogonality condition between the normal and tangent vectors. The vector $\boldsymbol{\psi}_{\parallel}^{(i)}$ describes the characteristic curves along the surface $\pi_i = p_i(\mathbf{x})$, and this holds for all $1 \leq i \leq d$. If these characteristic curves are parametrised by s , they satisfy the differential equations

$$\frac{dx_i}{ds} = \frac{\partial H}{\partial \pi_i}, \quad \frac{dp_i}{ds} = -\frac{\partial H}{\partial x_i}, \quad (1 \leq i \leq d). \quad (6.42)$$

These are the familiar Hamilton's equations which describe characteristic curves along the $H = 0$ surface [40]. From this $S_0^{(a)}(\mathbf{x})$ can be found,

$$\begin{aligned}\frac{dS_0^{(a)}}{ds} &= \sum_{i=1}^d \frac{\partial S_0^{(a)}}{\partial x_i} \frac{dx_i}{ds} = \sum_{i=1}^d p_i \frac{dx_i}{ds} \\ \Rightarrow S_0^{(a)}(\mathbf{x}) &= \int \sum_{i=1}^d p_i \frac{dx_i}{ds} ds = \int \mathbf{p} \cdot d\mathbf{x}.\end{aligned}\quad (6.43)$$

Therefore the value of $S_0^{(a)}(\mathbf{x})$ is given by the integral of the field $\mathbf{p}$ along the characteristic trajectories (6.42), which are bound to the surface $H = 0$. For now we will define $S_0^{(a)}(\mathbf{x}_a^*) = 0$, which is equivalent to having all trajectories start at the stable state $\mathbf{x}_a^*$.

Aside: We note that $\nabla_{\mathbf{p}}H(\mathbf{x}, \mathbf{0})$ recovers the deterministic equations of motion,

$$\nabla_{\mathbf{p}}H^{(\text{FPE})}(\mathbf{x}, \mathbf{0}) = \mathbf{A}(\mathbf{x}) = \frac{d\langle \mathbf{x} \rangle}{dt}, \quad (6.44a)$$

$$\nabla_{\mathbf{p}}H^{(\text{ME})}(\mathbf{x}, \mathbf{0}) = \sum_r \mathbf{v}_r f_r(\mathbf{x}) = \frac{d\langle \mathbf{x} \rangle}{dt}. \quad (6.44b)$$

Hence it is convenient to replace the parameter s with the time t , and dx_i/ds with $\dot{x}_i$. This apparent equivalence of s and t is coincidental; there is no time in Eqs. (6.35) and (6.37). As a counter-example to the equivalence, one could consider scaling all reaction rates in the master equation (6.36) by a constant factor k . This factor would appear in the mean-field dynamics, but could be removed from H , such that s and t no-longer coincide.

Path-Integral Formulation

The Hamilton–Jacobi equations (6.35) and (6.37) are not unique to the WKB formalism. The same functions can be recovered by considering the path-integral approach to the processes described by the master equation (6.36) and the Fokker–Planck equation (6.32). As described below, the functions (6.35) and (6.37) also characterise the probability density of trajectories, and can in turn lead us again to the discussion of the ‘most-likely path’.

Continuous Process:

The Fokker–Planck equation (6.32) describes the continuous stochastic process that is given by the (Itô) stochastic differential equation

$$\dot{\mathbf{x}} = \mathbf{A}(\mathbf{x}) + \mathbb{G}(\mathbf{x}) \cdot \boldsymbol{\eta}(t), \quad (6.45)$$

where the noise correlation matrix satisfies $\mathbb{G} \cdot \mathbb{G}^T = \varepsilon \mathbb{B}$, and $\boldsymbol{\eta}(t)$ are Gaussian white-noise variables which satisfy

$$\langle \eta_i(t) \rangle = 0, \quad \langle \eta_i(t) \eta_j(t') \rangle = \delta_{i,j} \delta(t - t'). \quad (6.46)$$

Again we are using the small parameter ε to characterise the noise strength.

To define a ‘path’ it is natural to consider a discrete-time approach, where time is divided into steps of length Δ . At time t , the state of the system is given by $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \dots, x_{d,t})^T$, where d is the dimension of our system, i.e. the number of stochastic variables. We use subscripts, $\mathbf{x}_t$, to distinguish the discrete-time variables from the continuous-time variables, $\mathbf{x}(t)$. We can represent the discrete-time SDE by using the Euler–Maruyama⁸ numerical integration scheme [28],

$$\mathbf{x}_{t+\Delta} = \mathbf{x}_t + \Delta \mathbf{A}(\mathbf{x}_t) + \sqrt{\Delta} \mathbb{G}(\mathbf{x}_t) \cdot \boldsymbol{\eta}_t. \quad (6.47)$$

⁸Leonhard Euler (1707–1783) and Gisiro Maruyama (1916–1986).

Here η_t is the discrete-time analogue of $\eta(t)$, i.e. it is a d -dimensional vector of independent Gaussian random numbers of unit variance. The correlation of the noise between species is contained in the matrix $\mathbb{G}$.

A path in discrete-time can be defined as the set $\{\mathbf{x}\} = \{\mathbf{x}_0, \mathbf{x}_\Delta, \dots, \mathbf{x}_T\}$, where T is the final time. As each x_i ($1 \leq i \leq d$) is a continuous variable, the probability to observe a specific path $\{\mathbf{x}\}$ is zero. We can, however, use Eq. (6.47) to define a path density,

$$\mathcal{P}[\{\mathbf{x}\}] = \int \prod_t d\eta_t \delta[\mathbf{x}_{t+\Delta} - \mathbf{x}_t - \Delta \mathbf{A}(\mathbf{x}_t) - \sqrt{\Delta} \mathbb{G}(\mathbf{x}_t) \cdot \eta_t] \mathcal{P}[\eta_t], \quad (6.48)$$

where $\mathcal{P}[\eta_t]$ is the joint probability distribution function of the Gaussian variables η_t , which can be factorised as the $\eta_{i,t}$ are independent. To evaluate the path density we introduce the auxiliary field $\tilde{\mathbf{x}}$ such that we can express the delta-functions in their exponential representation,

$$\mathcal{P}[\{\mathbf{x}\}] = \int \prod_t \frac{d\tilde{\mathbf{x}}_t d\eta_t}{(2\pi)^d} \exp \left\{ i \tilde{\mathbf{x}}_t \cdot [\mathbf{x}_{t+\Delta} - \mathbf{x}_t - \Delta \mathbf{A}(\mathbf{x}_t) - \sqrt{\Delta} \mathbb{G}(\mathbf{x}_t) \cdot \eta_t] \right\} \mathcal{P}[\eta_t]. \quad (6.49)$$

The integral over the noise variables at each time-point can be evaluated to give

$$\begin{aligned} \int d\eta_t \exp \left[i \sqrt{\Delta} \tilde{\mathbf{x}}_t \cdot \mathbb{G} \cdot \eta_t \right] \mathcal{P}[\eta_t] &= \int \frac{d\eta_t}{(2\pi)^{d/2}} \exp \left[i \sqrt{\Delta} \tilde{\mathbf{x}}_t \cdot \mathbb{G} \cdot \eta_t - \frac{1}{2} \eta_t \cdot \eta_t \right] \\ &= \exp \left[-\frac{1}{2} \Delta \tilde{\mathbf{x}}_t \cdot \mathbb{G} \cdot \mathbb{G}^T \cdot \tilde{\mathbf{x}}_t \right] \\ &= \exp \left[-\frac{1}{2} \varepsilon \Delta \tilde{\mathbf{x}}_t \cdot \mathbb{B} \cdot \tilde{\mathbf{x}}_t \right], \end{aligned} \quad (6.50)$$

and hence we can write Eq. (6.49) as

$$\mathcal{P}[\{\mathbf{x}\}] = \int \prod_t \frac{d\tilde{\mathbf{x}}_t}{(2\pi)^d} \exp \left\{ i \tilde{\mathbf{x}}_t \cdot \left[\mathbf{x}_{t+\Delta} - \mathbf{x}_t - \Delta \mathbf{A}(\mathbf{x}_t) + \frac{i}{2} \varepsilon \Delta \mathbb{B}(\mathbf{x}_t) \cdot \tilde{\mathbf{x}}_t \right] \right\}. \quad (6.51)$$

We can now restore the continuous-time limit, $\Delta \rightarrow 0$. We define $\lim_{\Delta \rightarrow 0} \prod_t d\tilde{\mathbf{x}}_t / (2\pi)^d = \mathcal{D}\tilde{\mathbf{x}}$, and using $\prod_t \exp(\cdot) = \exp(\sum_t \cdot)$ we can write the path probability density as

$$\mathcal{P}[\{\mathbf{x}\}] = \int \mathcal{D}\tilde{\mathbf{x}} \exp \left\{ i \int_0^T \tilde{\mathbf{x}} \cdot \left[\dot{\mathbf{x}} - \mathbf{A}(\mathbf{x}) + \frac{i}{2} \varepsilon \mathbb{B}(\mathbf{x}) \cdot \tilde{\mathbf{x}} \right] dt \right\}. \quad (6.52)$$

We can now define the probability density that the system reaches a point $\mathbf{x}_T$ at time $t = T$, given that it started at $\mathbf{x}_0$ at time $t = 0$. This is achieved by summing over all possible paths that meet these boundary conditions. We write this path integral as [23]

$$\mathcal{P}[\mathbf{x}_T, T | \mathbf{x}_0, 0] = \int_{\mathbf{x}_0}^{\mathbf{x}_T} \mathcal{D}\mathbf{x} \mathcal{P}[\{\mathbf{x}\}]. \quad (6.53)$$

This is now a probability density in the space of endpoints, $\mathbf{x}_T$. Through a relabelling of the auxiliary field, $\tilde{\mathbf{x}} = i\mathbf{p}/\varepsilon$, this density is given by

$$\begin{aligned} \mathcal{P}[\mathbf{x}_T, T | \mathbf{x}_0, 0] &= \int_{\mathbf{x}_0}^{\mathbf{x}_T} \mathcal{D}\mathbf{x} \mathcal{D}\mathbf{p} \exp \left\{ -\varepsilon^{-1} \int_0^T \left[\mathbf{p} \cdot \dot{\mathbf{x}} - \mathbf{p} \cdot \mathbf{A}(\mathbf{x}) - \frac{1}{2} \mathbf{p} \cdot \mathbb{B} \cdot \mathbf{p} \right] dt \right\} \\ &= \int_{\mathbf{x}_0}^{\mathbf{x}_T} \mathcal{D}\mathbf{x} \mathcal{D}\mathbf{p} \exp \left\{ -\varepsilon^{-1} \int_0^T \left[\mathbf{p} \cdot \dot{\mathbf{x}} - H^{(\text{FPE})}(\mathbf{x}, \mathbf{p}) \right] dt \right\}. \end{aligned} \quad (6.54)$$

Jump Process:

For the master equation we can construct a similar path-based solution. By following the work of Ref. [41], we can construct the probability of observing a path by approximating the discrete dynamics as a series of Poisson-distributed jumps in discrete time. This is the process described by the tau-leaping stochastic simulation algorithm [42]. If the state of the discrete population at time t is given by $\mathbf{n}_t = (n_{1,t}, n_{2,t}, \dots, n_{d,t})^T$, then the state at time $t + \Delta$ is described by

$$\mathbf{n}_{t+\Delta} = \mathbf{n}_t + \sum_r \mathbf{v}_r k_{r,t}, \quad (6.55)$$

where $k_{r,t}$ is a Poisson-distributed random variable with mean $\lambda_{r,t} = T_{\mathbf{n}_t}^{\mathbf{v}_r} \Delta$. The $k_{r,t}$ represent the number of jumps of reaction r that occur in a time-step Δ . As the state-space of $\mathbf{n}_t$ is discrete, we can define the probability of observing the path $\{\mathbf{n}\} = \{\mathbf{n}_0, \mathbf{n}_\Delta, \dots, \mathbf{n}_T\}$. This probability is given by

$$\mathcal{P}[\{\mathbf{n}\}] = \prod_t \sum_{\mathbf{k}_t} \delta \left[\mathbf{n}_{t+\Delta} - \mathbf{n}_t - \sum_r \mathbf{v}_r k_{r,t} \right] \mathcal{P}[\mathbf{k}_t], \quad (6.56)$$

where $\mathcal{P}[\mathbf{k}_t]$ is the joint probability distribution of the Poisson random variables $\mathbf{k}_t$, which can be factorised as the $k_{r,t}$ are independent. Note that the vector $\mathbf{k}_t$ has elements labelled by the reaction index r , as opposed to the variable index i . The sum over $\mathbf{k}_t$ represents $\sum_{k_{1,t}=0}^{\infty} \sum_{k_{2,t}=0}^{\infty} \dots$.

We now switch to the continuous variable $\mathbf{x} = \mathbf{n}/\Omega$, and introduce the auxiliary field $\tilde{\mathbf{x}}$. We are still considering a discrete-time process, but the continuous state-space means we need to again consider the probability density of the path $\{\mathbf{x}\} = \{\mathbf{x}_0, \mathbf{x}_\Delta, \dots, \mathbf{x}_T\}$. We write this density as

$$\mathcal{P}[\{\mathbf{x}\}] = \prod_t \int \frac{d\tilde{\mathbf{x}}_t}{(2\pi)^d} \sum_{\mathbf{k}_t} \exp \left\{ i\tilde{\mathbf{x}}_t \cdot \left[\mathbf{x}_{t+\Delta} - \mathbf{x}_t - \sum_r \frac{\mathbf{v}_r k_{r,t}}{\Omega} \right] \right\} \mathcal{P}[\mathbf{k}_t]. \quad (6.57)$$

Now separating the terms which are dependent on the $k_{r,t}$, and inserting the Poisson probability distributions $\mathcal{P}[\mathbf{k}_t]$, we have

$$\begin{aligned} \sum_{\mathbf{k}_t} \exp \left\{ -i\tilde{\mathbf{x}}_t \cdot \left[\sum_r \frac{\mathbf{v}_r k_{r,t}}{\Omega} \right] \right\} \prod_r \frac{\lambda_{r,t}^{k_{r,t}} \exp[-\lambda_{r,t}]}{k_{r,t}!} \\ = \prod_r \exp[-\lambda_{r,t}] \sum_{k_{r,t}} \frac{(\exp[-i\tilde{\mathbf{x}}_t \cdot \mathbf{v}_r / \Omega] \lambda_{r,t})^{k_{r,t}}}{k_{r,t}!} \\ = \prod_r \exp \left\{ \lambda_{r,t} \left(\exp \left[\frac{-i\tilde{\mathbf{x}}_t \cdot \mathbf{v}_r}{\Omega} \right] - 1 \right) \right\}. \quad (6.58) \end{aligned}$$

Substituting this expression back into Eq. (6.57) we obtain

$$\begin{aligned} \mathcal{P}[\{\mathbf{x}\}] = \prod_t \int \frac{d\tilde{\mathbf{x}}_t}{(2\pi)^d} \left[\exp \left\{ i\tilde{\mathbf{x}}_t \cdot [\mathbf{x}_{t+\Delta} - \mathbf{x}_t] \right\} \times \right. \\ \left. \exp \left\{ \sum_r \lambda_{r,t} \left(\exp \left[\frac{-i\tilde{\mathbf{x}}_t \cdot \mathbf{v}_r}{\Omega} \right] - 1 \right) \right\} \right]. \quad (6.59) \end{aligned}$$

We can now restore the continuous-time limit. Again we define $\lim_{\Delta \rightarrow 0} \prod_t d\tilde{\mathbf{x}}_t / (2\pi)^d = \mathcal{D}\tilde{\mathbf{x}}$, such that we can write

$$\mathcal{P}[\{\mathbf{x}\}] = \int \mathcal{D}\tilde{\mathbf{x}} \exp \left\{ \int_0^T \left[i\tilde{\mathbf{x}} \cdot \dot{\mathbf{x}} + \sum_r \Omega f_r(\mathbf{x}) \left(\exp \left[\frac{-i\tilde{\mathbf{x}} \cdot \mathbf{v}_r}{\Omega} \right] - 1 \right) \right] dt \right\}, \quad (6.60)$$

where we have used $\lambda_{r,t} = T_{n_r}^{\mathbf{v}_r} \Delta \rightarrow \Omega f_r(\mathbf{x}_t) \Delta$.

We can now define the probability density in the space of end points. That is the probability that we find the system at state $\mathbf{x}_T$ at time $t = T$ given it was started at state $\mathbf{x}_0$ at time $t = 0$. This quantity is given by Eq. (6.53). Substituting Eq. (6.60) into Eq. (6.53), and relabelling the auxiliary field $\tilde{\mathbf{x}} = i\Omega \mathbf{p}$, gives

$$\begin{aligned} \mathcal{P}[\mathbf{x}_T, T | \mathbf{x}_0, 0] = \int_{\mathbf{x}_0}^{\mathbf{x}_T} \mathcal{D}\mathbf{x} \mathcal{D}\mathbf{p} \exp \left\{ -\Omega \int_0^T \left[\mathbf{p} \cdot \dot{\mathbf{x}} - \sum_r f_r(\mathbf{x}) (e^{\mathbf{p} \cdot \mathbf{v}_r} - 1) \right] dt \right\} \\ = \int_{\mathbf{x}_0}^{\mathbf{x}_T} \mathcal{D}\mathbf{x} \mathcal{D}\mathbf{p} \exp \left\{ -\Omega \int_0^T [\mathbf{p} \cdot \dot{\mathbf{x}} - H^{(\text{ME})}(\mathbf{x}, \mathbf{p})] dt \right\}, \quad (6.61) \end{aligned}$$

which is of the same form as Eq. (6.54).

The function $H^{(\text{ME})}(\mathbf{x}, \mathbf{p})$ (or its Fokker–Planck equivalent) characterises the probability of observing a path. Using the saddle-point approximation of the path-integral (6.61) [or Eq. (6.54)], we can define the most-likely path between the

points $\mathbf{x}_0$ and $\mathbf{x}_T$ as the path, $(\{\mathbf{x}\}, \{\mathbf{p}\})$, which minimises $S[(\{\mathbf{x}\}, \{\mathbf{p}\})] = \int_0^T [\mathbf{p} \cdot \dot{\mathbf{x}} - H(\mathbf{x}, \mathbf{p})] dt$. This path is embedded in space *and* time, such that its relation with the most-likely paths in Sects. 6.2 and 6.4 is not entirely obvious.

The function $S[(\{\mathbf{x}\}, \{\mathbf{p}\})]$ is referred to as the ‘action’ of the path, as this formalism is very similar to the description of classical mechanics [40]. For $(\{\mathbf{x}\}, \{\mathbf{p}\})$ to be the most-likely (minimum-action) path, it should satisfy

$$\delta S = S[(\{\mathbf{x} + \delta\mathbf{x}\}, \{\mathbf{p} + \delta\mathbf{p}\})] - S[(\{\mathbf{x}\}, \{\mathbf{p}\})] = 0, \quad (6.62)$$

where $\delta\mathbf{x}$ and $\delta\mathbf{p}$ are small, independent perturbations to the path. Evaluating the above function shows that the minimum-action paths are given by the solutions of Hamilton’s equations, Eq. (6.42) [40]. Hence the characteristic curves obtained from the WKB formalism are the most-likely paths between two states.

An alternative approach is to first extremise the action at each point along the path with respect to $\mathbf{p}$. This is a simple exercise in the Fokker–Planck framework, from which we obtain

$$\begin{aligned} 0 &= \frac{\partial}{\partial p_i} \left[\mathbf{p} \cdot \dot{\mathbf{x}} - \mathbf{p} \cdot \mathbf{A}(\mathbf{x}) - \frac{1}{2} \mathbf{p} \cdot \mathbb{B} \cdot \mathbf{p} \right] \\ \Rightarrow 0 &= \dot{x}_i - A_i(\mathbf{x}) - B_{i,j} p_j \\ \Rightarrow \mathbf{p} &= \mathbb{B}^{-1} \cdot [\dot{\mathbf{x}} - \mathbf{A}(\mathbf{x})]. \end{aligned} \quad (6.63)$$

Substituting this value back into the action gives

$$\int_0^T [\mathbf{p} \cdot \dot{\mathbf{x}} - H^{(\text{FPE})}(\mathbf{x}, \mathbf{p})] dt = \frac{1}{2} \int_0^T [\dot{\mathbf{x}} - \mathbf{A}(\mathbf{x})] \cdot \mathbb{B}^{-1} \cdot [\dot{\mathbf{x}} - \mathbf{A}(\mathbf{x})] dt, \quad (6.64)$$

where the integrand is the Onsager–Machlup⁹ functional [43]. From the SDE (6.45), we recognise this as an integral over the noise variables. Along the deterministic paths, where $\dot{\mathbf{x}} = \mathbf{A}(\mathbf{x})$, the action is zero. Such treatment is not possible analytically in the case of the master equation, but we will describe a numerical method which directly minimises this action by solving the associated Euler–Lagrange equations [23].

Large Deviations Theory

Large deviations theory is the field of mathematics concerned with the analysis of rare events and statistical outliers [23]. The application of this field to stochastic processes is described in Ref. [22]. The aim of this approach is to construct exponential estimates for the probability densities to observe paths between two states (not necessarily stable states). If we have a continuous path $\varphi(t)$, then the probability density to observe a simulation path $\mathbf{X}_\Omega(t)$ that is within the δ -tube of this path is given by

⁹Lars Onsager (1903–1976) and Stefan Machlup (1927–2008).

$$\Pr \left(\sup_{0 \leq t \leq T} |\mathbf{X}_\Omega(t) - \boldsymbol{\varphi}(t)| < \delta \right) \sim f(\delta) \exp(-\Omega S[\boldsymbol{\varphi}]). \quad (6.65)$$

Here the function $f(\delta)$ is often ignored, as we are only concerned with the exponential approximation, and $S[\boldsymbol{\varphi}]$ is known as the large deviations rate functional [23]. The Onsager–Machlup functional in Eq. (6.64) is the rate functional of the SDE (6.45) [22, 23]. From Eq. (6.65), the most-likely path from $\boldsymbol{\varphi}(0) = \mathbf{x}_0$ to $\boldsymbol{\varphi}(T) = \mathbf{x}_T$ must minimise $S[\boldsymbol{\varphi}]$.

In the WKB formulation we have not specified the length of the path in terms of the time, T . If we vary the length of the paths, we would expect the value of the action to change, and hence it should have a minimum at some value of T . The most likely path between two states, $\mathbf{x}_0$ and $\mathbf{y}$, then has action

$$V(\mathbf{x}_0, \mathbf{y}) = \inf_T \inf_{\boldsymbol{\varphi}} \{S[\boldsymbol{\varphi}] : \boldsymbol{\varphi}(0) = \mathbf{x}_0, \boldsymbol{\varphi}(T) = \mathbf{y}\}. \quad (6.66)$$

This quantity measures the difficulty in moving from the point $\mathbf{x}_0$ to the point $\mathbf{y}$, and is referred to as the quasi-potential [22]. This is equivalent to the action recovered from the WKB formulation, $S_0(\mathbf{x})$ [23, 44]. It is called a ‘quasi’-potential as this function is not additive; for distinct points $\mathbf{x}_0$, $\mathbf{x}_1$, and $\mathbf{x}_2$ we have $V(\mathbf{x}_0, \mathbf{x}_2) \leq V(\mathbf{x}_0, \mathbf{x}_1) + V(\mathbf{x}_1, \mathbf{x}_2)$. This inequality becomes an equality only if $\mathbf{x}_1$ lies on the most-likely path between $\mathbf{x}_0$ and $\mathbf{x}_2$.

The quasi-potential must be defined with respect to a reference state. If there exists a stable fixed point in the deterministic dynamics, $\mathbf{x}_a^*$, then we will often use this point as the reference state for the quasi-potential, i.e. $V^{(a)}(\mathbf{x}) = V(\mathbf{x}_a^*, \mathbf{x})$.

Numerical Methods

As described above, the quasi-stationary distribution about the stable state $\mathbf{x}_a^*$ is characterised by the action $S_0^{(a)}(\mathbf{x}) = \int \mathbf{p} \cdot d\mathbf{x}$, where $(\{\mathbf{p}\}, \{\mathbf{x}\})$ is the most-likely path between the stable fixed point $\mathbf{x}_a^*$ and the position $\mathbf{x}$, where t parametrises the curve. The paths are described by the characteristic equations (6.42). Obtaining $S_0(\mathbf{x})$ thus reduces to solving a two-boundary-value problem, as we have an initial position $[\mathbf{x}(0) = \mathbf{x}_a^*, \mathbf{p}(0) = \mathbf{0}]$ and a final position $\mathbf{x}(T) = \mathbf{x}$. Solutions of this problem can be obtained through the application of shooting methods or iterative schemes. A more general approach is to directly minimise the action of the path integral by solving the associated Euler–Lagrange equations. We now discuss these methods in turn.

Solutions of Characteristic Equations: Shooting methods are the conventional choice of numerical method for solving boundary-value problems. The procedure is to turn the boundary-value problem into an initial-value problem, which are much easier to solve. One must then find the initial condition that generates the correct final boundary condition [45]. The new initial condition is specified as a perturbation from the fixed point $\mathbf{x}_a^*$, i.e. we let $\mathbf{x}(\Delta) = \mathbf{x}_a^* + \delta\mathbf{x}$, where Δ is a small time-step. In the limit $\Delta \rightarrow 0$ this condition is equivalent to specifying a condition for $\dot{\mathbf{x}}(0)$, provided that $\delta\mathbf{x}$ scales suitably with Δ . A corresponding perturbation $\delta\mathbf{p}$ is specified and with these conditions Hamilton’s equations (6.42) are forward-integrated to generate a

trajectory [10, 26]. Varying the initial perturbation allows us to generate multiple trajectories across the domain. We can then find the action, $S_0^{(a)}(\mathbf{x}) = \int \mathbf{p} \cdot d\mathbf{x}$, and hence the quasi-stationary distribution about this stable state.

The downfall of the shooting method, however, is that convergence to the correct final boundary condition can be very slow. In fact, if the final boundary condition is a saddle point, then it is very unlikely that shooting to such a point will be possible [10]. Furthermore, as the dimensionality gets higher it becomes increasingly difficult and inefficient to find the target boundary condition.

An alternative approach to solving the boundary-value problem has been used in Refs. [16, 21, 46, 47]. Here the solution is found through an iterative scheme, where the characteristic equations (6.42) for $\mathbf{x}$ are integrated forwards in time and the equations for $\mathbf{p}$ are integrated backwards in time. This method is straightforward to implement and was described in Sect. 5.5. It does not suffer from the same problems as the shooting method. However, this method requires the final value of $\mathbf{p}$, which is not uniquely specified in $d \geq 2$ dimensions. If the final position is a stationary point of the characteristic equations (6.42), then these can be solved to determine the boundary condition, otherwise we are in trouble.

Numerical Minimisation: Rather than trying to solve the integral problem as is done in the shooting or iterative methods, a different approach is to solve the Euler–Lagrange equations that define the minimum-action path. This can be achieved through relaxation methods, such as the method of steepest descent [45]. A problem that is common to the above solution methods is the convergence of the equations of motion to a stationary point. By definition it takes infinitely long for deterministic equations to reach a stationary point, so for how long do we integrate Eqs. (6.42)? The geometric minimum action method (gMAM) overcomes this problem by reparametrising the path, in this case onto the interval $[0, 1]$ [44]. Importantly this method does not require any knowledge of the initial or final values of $\mathbf{p}$, only the position boundary conditions are needed. Details of this method are omitted from this thesis due to its complexity, however it is extensively described in Ref. [44].

6.6 Examples

Tunnelling Example

We return to the problem discussed in Chap. 5 of the fixation of two mutations in a fixed size population. We identified different dynamical regimes in that system. In region I, when there existed a single stable state located on the boundary of the domain, we were able to compute an analytical expression for the QSD. However, in regions II and III when a fixed point existed in the middle of the domain, we could determine the most-likely paths, but not the QSD across the state space. Using the techniques developed in this chapter, we are now in a position where we can compute this distribution numerically. To do this we use the gMAM algorithm mentioned above [44]. The results are shown in Fig. 6.9, where we plot the quasi-potential $S_0(\mathbf{x})$






Fig. 6.9 **a** The quasi-potential about the stable fixed point in region II for the cancer initiation model, along with the dominant escape path from the interior stable state to the 1–2 boundary. Parameters are $r_0 = 1.00$, $r_1 = 0.98$, $r_2 = 0.95$, and $u_1 = u_2 = 10^{-2}$. **b** The quasi-potential about the stable fixed point in region III, along with the dominant escape path from the interior metastable state to the absorbing state. Parameters are $r_0 = 1.00$, $r_1 = 0.95$, $r_2 = 0.98$, and $u_1 = u_2 = 10^{-2}$. In **(a and b)** the contours are a multiplicative factor of $\sqrt{2}$ apart, such that every two contours corresponds to a doubling of the quasi-potential. Equal contour levels are used in each panel

about the interior fixed point in regions II and III. The QSD, $\rho^*(\mathbf{x}) \sim \exp[-N S_0(\mathbf{x})]$, can easily be inferred from these results as having a peak at the fixed point and steeply dropping off as the distance from the stable state increases. The most-likely escape paths, which are the same as those plotted in Fig. 5.16 but obtained using the gMAM algorithm, are seen to be orthogonal to the equipotential contours, validating the use of the term ‘potential’. Numerical issues arise close to the boundaries of the concentration simplex, when the system jumps from being two-dimensional to one-dimensional.

Toggle Switch

To further demonstrate the applicability of this theory, we will consider a two-dimensional example which has multiple stable states and transitions between the different basins of attraction. The example is that of the genetic toggle switch, which is one of the archetypal problems used to demonstrate the WKB method and quasi-potential landscapes [15, 48–51]. We consider the minimal model which only describes the behaviour of the two types of protein, labelled 1 and 2. We assume that proteins degrade at a constant rate, and they are produced at a rate that is dependent on the presence of the other protein through mutual inhibition. This example is described by the following reaction scheme:



where γ_0 is the protein degradation rate, r is the basic production rate, and the Hill coefficient h determines the shape of the inhibition function. The system size is Ω , which determines the typical protein number. We set the parameter $\gamma_0 = 1$ throughout, which corresponds to a rescaling of time. The parameter α allows us to introduce some asymmetry into the problem. We choose the parameter r such that there exists two stable states in the deterministic dynamics. One of these stable states, $\mathbf{x}_1^*$, has a low concentration of protein 1 and a high concentration of protein 2, and vice-versa for the other stable state $\mathbf{x}_2^*$, where $\mathbf{x} = (n_1, n_2)^T / \Omega$. There is a separatrix between the basins of attraction, and there exists a saddle point on this boundary at $\mathbf{x}_s$.

To determine the quasi-potential landscape in this system, we first need to find the values of $S_0^{(a)}(\mathbf{x})$ for each of the stable states. For this we use the gMAM algorithm. We then match these actions at the saddle point to generate the quasi-potential $S_0(\mathbf{x})$, which satisfies [10]

$$S_0(\mathbf{x}) = \min \left\{ S_0^{(1)}(\mathbf{x}), S_0^{(2)}(\mathbf{x}) + S_0^{(1)}(\mathbf{x}_s) - S_0^{(2)}(\mathbf{x}_s) \right\}. \quad (6.68)$$

We plot examples of this quasi-potential in Fig. 6.10. In panels (a) and (b) of Fig. 6.10, we can see that the most-likely transition paths pass through the saddle point in between the basins of attraction. Once the path has crossed the saddle point, it follows




Fig. 6.10 **a** Quasi-potential in the symmetric toggle-switch ($\alpha = 1$), along with the most likely path from each stable state to the other. The quasi-potential is symmetric about the separatrix $x_1 = x_2$, so we only show one half of it. The *gray lines* in right-hand side are the streamlines of the deterministic flow. **b** Quasi-potential in an asymmetric toggle-switch ($\alpha = 0.85$), along with the most-likely transition paths. The *thick black line* is the separatrix between the basins of attraction. In **(a and b)** the contours are a multiplicative factor of $\sqrt{2}$ apart, such that every two contours corresponds to a doubling of the quasi-potential. Equal contour levels are used in each panel. Remaining model parameters are $\gamma_0 = 1, r = 2, h = 3$

the deterministic trajectory to the stable state. Again we see that the escape paths are orthogonal to the equipotential lines. In the asymmetric toggle-switch in Fig. 6.10b, we see that the stable state $\mathbf{x}_1^*$ to the left of the separatrix has a ‘deeper potential well’ than $\mathbf{x}_2^*$. Considering the stationary distribution, $\rho^{\text{st}}(\mathbf{x}) \sim \exp[-NS_0(\mathbf{x})]$, there is a much higher probability to be found close to the state $\mathbf{x}_1^*$, as opposed to the stable state $\mathbf{x}_2^*$.

6.7 Summary

The WKB method is a very powerful technique that has been used to describe fixation and equilibration across a wide-variety of stochastic models, in particular when describing the statistics of rare-events such as extinction. However, there has been a lack of consistency across multiple disciplines when describing this approach. This includes the work of this PhD student. In this chapter we have returned to the mathematical basis of the WKB method in the hope of clearing up the ambiguities in the terminology, and we have provided a tutorial-style walk-through of how we can apply the method to stochastic processes.

Through the use of toy models we have reintroduced the concepts of the quasi-stationary distribution and the most-likely path between two states. While the QSD has a rigorous mathematical definition, the most-likely path is a much more subtle concept. Along with the path, we must also specify the ensemble in which this path resides. In Sect. 6.2, we identify a hierarchy of three ensembles with varying complexing for a simple individual-based model. The first space of trajectories contains the full information about the transitions, including the time spent at each state. This is the most complex space with an uncountable set of trajectories where we can only define path densities. This object is useful in the path-integral or large-deviations frameworks, but does not help with the practitioner’s intuition. If we only consider the states of the system by marginalising over the time, then we can define the ensemble as all paths which start at stable point in the state-space and reach an absorbing boundary. However, there is large degeneracy in this space because paths often return to the stable state. By imposing the condition that the system does not return to the stable state, we can construct a much smaller space of paths and easily attribute weights to the associated trajectories. It is not wrong to define the most-likely path in the other spaces, but the definition is more clear in this smaller space.

To introduce the mechanics of the WKB method, we first focussed on processes with a single stochastic variable in Sect. 6.3. We discussed the approach used to analyse continuous stochastic processes described by a Fokker–Planck equation, and the familiar discrete jump processes described by a master equation. We considered both frameworks to make contact with the literature from both of these fields, and to show the similarities, and most importantly the differences, between these approaches to modelling a system. In these one-dimensional systems we are able to obtain explicit expressions for the quasi-stationary distributions.

In a one-dimensional system, it was easy to show that the most-likely path from a stable state to an absorbing boundary is the path which does not step ‘backwards’, in the sense that the system always moves away from the stable state (it is easy provided that we know what we are talking about!). In the scenario where there are multiple ‘forward-only’ paths to the absorbing boundary (or to another stable state), as is found in higher-dimensional systems, then we again need to define the most-likely path. In Sect. 6.4, we use a four-state toy model to illustrate this definition. We are also able to calculate the quasi-stationary distribution in this system.

For stochastic processes featuring a larger number of stochastic variables ($d \geq 2$), analytical progress in the WKB framework is often not possible. Even in the absence of tractability, the WKB method illuminates the problem and provides a large amount of information by mapping the problem from a stochastic process to classical mechanics. Applying the WKB method allows us to construct quasi-stationary distributions about the stable states of the underlying model. These distributions are described by an ‘action’, which can be derived from the path-integral framework. This approach also shows that characteristic equations obtained from the WKB method describe most-likely paths between two states. The action is also closely related to the quasi-potential described in the theory of large deviations [22]. A variety of numerical methods have been proposed to extract the quasi-potential and the most-likely paths, and examples of these quantities are shown in Sect. 6.6.

The WKB approach greatly improves on the ‘traditional’ methods of analysis of individual-based systems, which include direct numerical integration of the master equation or using diffusion approximations. Integration of the master equation can be very inefficient; each point in state-space has an associated equation, and the complete set of equations need to be integrated in parallel as they are coupled. Also, as we are often interested in rare events, this integration must be carried out for a long time. At the end of this integration we are left with a description of only a single point in parameter space, with no knowledge of how the system behaves in the vicinity of this point. Diffusion approximations can reduce the large set of master equations to a partial differential equation, which greatly improves the tractability. However, this approach only yields accurate results for the bulk of the probability distribution, and the Fokker–Planck equation fails to describe rare events which are characterised by the tails of the distribution. The WKB method captures these large deviations from the expected behaviour, providing insight into events such as extinction, fixation and equilibration.

References

1. G. Green, On the motion of waves in a variable canal of small depth and width. *Trans. Cambridge Phil. Soc.* **6**, 457 (1837)
2. J. Liouville, Troisième mémoire sur le développement des fonctions ou parties de fonctions en séries dont les divers termes sont assujettis à satisfaire à une même équation différentielle du second ordre, contenant un paramètre variable. *J. Math. Pure Appl.* **2**, 16 (1837)

3. H. Jeffreys, On certain approximate solutions of linear differential equations of the second order. *Proc. London Math. Soc.* **2**, 428 (1925)
4. G. Wentzel, Eine Verallgemeinerung der Quantenbedingungen für die Zwecke der Wellenmechanik. *Z. Phys.* **38**, 518 (1926)
5. H.A. Kramers, Wellenmechanik und halbzahlige Quantisierung. *Z. Phys.* **39**, 828 (1926)
6. L. Brillouin, La mécanique ondulatoire de Schrödinger: Une méthode générale de résolution par approximations successives. *C. R. Acad. Sci. Paris* **183**, 24 (1926)
7. R. Kubo, K. Matsuo, K. Kitahara, Fluctuation and relaxation of macrovariables. *J. Stat. Phys.* **9**, 51 (1973)
8. R. Graham, T. Tél, Existence of a potential for dissipative dynamical systems. *Phys. Rev. Lett.* **52**, 9 (1984)
9. H. Gang, Stationary solution of master equations in the large-system-size limit. *Phys. Rev. A* **36**, 5782 (1987)
10. M. Dykman, E. Mori, J. Ross, P. Hunt, Large fluctuations and optimal paths in chemical kinetics. *J. Chem. Phys.* **100**, 5735 (1994)
11. P. Hänggi, P. Talkner, M. Borkovec, Reaction-rate theory: fifty years after Kramers. *Rev. Mod. Phys.* **62**, 251 (1990)
12. H.A. Kramers, Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica* **7**, 284 (1940)
13. B. Gaveau, M. Moreau, J. Tóth, *Master Equations and Path-integral Formulation of Variational Principles for Reactions*, eds. by S. Sieniutycz, H. Farkas. Variational and Extremum Principles in Macroscopic Systems (Elsevier, Amsterdam, 2005)
14. C. Escudero, A. Kamenev, Switching rates of multi-step reactions. *Phys. Rev. E* **79**, 041149 (2009)
15. M. Assaf, E. Roberts, Z. Luthey-Schulten, Determining the stability of genetic switches: explicitly accounting for mRNA noise. *Phys. Rev. Lett.* **106**, 248102 (2011)
16. V. Elgart, A. Kamenev, Rare event statistics in reaction-diffusion systems. *Phys. Rev. E* **70**, 041106 (2004)
17. D.A. Kessler, N.M. Shnerb, Extinction rates for fluctuation-induced metastabilities: a real-space WKB approach. *J. Stat. Phys.* **127**, 861 (2007)
18. A. Kamenev, B. Meerson, Extinction of an infectious disease: a large fluctuation in a nonequilibrium system. *Phys. Rev. E* **77**, 061107 (2008)
19. M. Assaf, B. Meerson, Extinction of metastable stochastic populations. *Phys. Rev. E* **81**, 021116 (2010)
20. L. Billings, L. Mier-Y-Teran-Romero, B. Lindley, I.B. Schwartz, Intervention-based stochastic disease eradication. *PLoS ONE* **8**, e70211 (2013)
21. P. Ashcroft, F. Michor, T. Galla, Stochastic tunneling and metastable states during the somatic evolution of cancer. *Genetics* **199**, 1213 (2015)
22. M.I. Freidlin, A.D. Wentzell, *Random Perturbations of Dynamical Systems* (Springer, New York, 1984)
23. H. Touchette, The large deviation approach to statistical mechanics. *Phys. Rep.* **478**, 1 (2009)
24. B. Meerson, P.V. Sasorov, Noise-driven unlimited population growth. *Phys. Rev. E* **78**, 060103 (2008)
25. O. Ovaskainen, B. Meerson, Stochastic models of population extinction. *Trends Ecol. Evol.* **25**, 643 (2010)
26. A.J. Black, A.J. McKane, WKB calculation of an epidemic outbreak distribution. *J. Stat. Mech.* **2011**, P12006 (2011)
27. P. Collet, S. Martínez, J. San Martín, *Quasi-Stationary Distributions* (Springer, Berlin, 2013)
28. H. Risken, *The Fokker-Planck Equation* (Springer, Berlin, 1989)
29. N.G. van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, Amsterdam, 2007)
30. C.M. Bender, S.A. Orszag, *Advanced Mathematical Methods for Scientists and Engineers* (Springer, New York, 1999)
31. C.W. Gardiner, *Handbook of Stochastic Methods* (Springer, New York, 2009)
32. N. Goel, N. Richter-Dyn, *Stochastic Models in Biology* (Academic Press, New York, 1974)

33. P. Hänggi, H. Grabert, P. Talkner, H. Thomas, Bistable systems: master equation versus Fokker-Planck modeling. *Phys. Rev. A* **29**, 371 (1984)
34. V. Elgart, A. Kamenev, Classification of phase transitions in reaction-diffusion models. *Phys. Rev. E* **74**, 041101 (2006)
35. M. Assaf, A. Kamenev, B. Meerson, Population extinction in a time-modulated environment. *Phys. Rev. E* **78**, 041123 (2008)
36. M. Assaf, A. Kamenev, B. Meerson, Population extinction risk in the aftermath of a catastrophic event. *Phys. Rev. E* **79**, 011127 (2009)
37. V. Méndez, M. Assaf, D. Campos, W. Horsthemke, Stochastic dynamics and logistic population growth. *Phys. Rev. E* **91**, 062133 (2015)
38. P.F. Verhulst, Notice sur la loi que la population suit dans son accroissement. *Correspondance Mathématique et Physique* **10**, 113 (1838)
39. T. Antal, P. Krapivsky, Exact solution of a two-type branching process: models of tumor progression. *J. Stat. Mech.* **2011**, P08018 (2011)
40. L.D. Landau, E.M. Lifshitz, *Mechanics*, vol. 1 (Pergamon Press, Oxford, 1976)
41. T. Brett, T. Galla, Stochastic processes with distributed delays: chemical Langevin equation and linear-noise approximation. *Phys. Rev. Lett.* **110**, 250601 (2013)
42. D.T. Gillespie, Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.* **115**, 1716 (2001)
43. L. Onsager, S. Machlup, Fluctuations and irreversible processes. *Phys. Rev.* **91**, 1505 (1953)
44. M. Heymann, E. Vanden-Eijnden, The geometric minimum action method: a least action principle on the space of curves. *Comm. Pure Appl. Math.* **61**, 1052 (2008)
45. W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, Cambridge, 2007)
46. A.I. Chernykh, M.G. Stepanov, Large negative velocity gradients in Burgers turbulence. *Phys. Rev. E* **64**, 026306 (2001)
47. I. Lohmar, B. Meerson, Switching between phenotypes and population extinction. *Phys. Rev. E* **84**, 051901 (2011)
48. D.M. Roma, R.A. O’Flanagan, A.E. Ruckenstein, A.M. Sengupta, R. Mukhopadhyay, Optimal path to epigenetic switching. *Phys. Rev. E* **71**, 011902 (2005)
49. S. Bhattacharya, Q. Zhang, M.E. Andersen, A deterministic map of Waddington’s epigenetic landscape for cell fate specification. *BMC Syst. Biol.* **5**, 85 (2011)
50. C. Lv, X. Li, F. Li, T. Li, Constructing the energy landscape for genetic switching system driven by intrinsic noise. *PLoS ONE* **9**, e88167 (2014)
51. M. Lu, J. Onuchic, E. Ben-Jacob, Construction of an effective landscape for multistate genetic switches. *Phys. Rev. Lett.* **113**, 078102 (2014)

Chapter 7

Conclusions

7.1 General Discussion

This thesis has been concerned with the phenomena of fixation and equilibration in stochastic, individual-based processes. Fixation, which is often a consequence of the extinction of one or more types of individual within a population, is an important aspect of biological systems, particularly in the field of population genetics from where the term originates. Many systems exhibit fixation (or extinction), such as the disappearance of a disease, the spreading of an opinion, or the propagation of mutated cells through a tissue. Equilibration, on the other hand, describes the process of reaching a steady state. These stationary systems often arise when types of individuals are constantly reintroduced through mutations, for example, or through adopting novel strategies. These two features can be closely related. In systems where fixation takes a long time, the population will first relax to a so-called metastable state. In this thesis these features have been analysed in a range of applications: From simple one-dimensional birth–death processes describing the interaction of two strategies in an evolutionary game, to a two-dimensional model describing the accumulation of mutations during the initiation of cancer. The impact that a switching environment has on the evolutionary dynamics of a population has also been investigated.

The analysis of these systems is carried out using the tools and techniques of statistical mechanics. Here knowledge of the underlying microscopic dynamics on the level of the individual has been used to make predictions about the macroscopic outcomes at the population level. Through the analytical characterisation of the evolutionary dynamics that are observed in stochastic processes, we have obtained a greater understanding of fixation and equilibration. For the majority of this thesis we have not focused on a specific application, instead we have been concerned with constructing generic mathematical frameworks to analyse the evolutionary processes. The application of these techniques is then illustrated through some popular examples. Often the focus has been on evolutionary game dynamics, which describe the interaction of competing strategies. These scenarios demonstrate the richness of

dynamical behaviours that can be seen in these models, such as coexistence and bi-stability. We have also explored examples of individual-based models that describe some specific biological processes, such as a genetic toggle-switch or the initiation of cancer. The accompanying analytical treatment of these models has been presented in as general a framework as possible, such that these techniques are readily transferable to different systems.

7.2 Summary of Results

Finite Populations in Switching Environments

Models, and the underlying system that they represent, can contain additional levels of complexity that go beyond just the dynamics of the individuals. In Chap. 3 we investigated the impact of a stochastically switching environment on the evolutionary dynamics of a two-species population. Such systems are relevant when environmental changes occur on roughly the same timescale as the population dynamics, as is the case when bacteria are repeatedly exposed to antibiotics [1]. We were able to extend the existing methods used to describe fixation and extinction in birth–death processes to account for this extra stochasticity. Applying these results to evolutionary games, where different environments favour different strategies, we observed combinations of switching rates for which the mutant was more successful in the switching environment than in either of the fixed evolutionary games. This non-trivial result can be explained by considering the interplay between the selection effect of the environment and the external noise. Equilibration was also investigated in this model. We introduced mutations in the dynamics, such that the individuals can randomly switch their strategy, and successfully predicted the resulting stationary states. This work provides a first mathematical characterisation of the effects one may expect in systems that are subject to selection, mutation, demographic stochasticity and external randomness.

Fixation Time Distributions in Birth–Death Processes

In Chap. 4 we questioned whether the mean fixation time is a good description of the overall fixation statistics in a birth–death process. To answer this we computed the exact fixation time distributions. This was achieved by evaluating the spectrum of the master equation describing the process, and we considered the dynamics in eigenspace. Applying this method to some typical scenarios from evolutionary game theory, we observed that fixation time distributions can be broad and skewed, especially in the coexistence game where a heterogeneous population is favoured by selection. Along with the exact representations of the fixation time distributions, we produced exponential processes that permit efficient sampling from these distributions. These have the potential to be used as a very effective model-reduction tool. When rare mutations were introduced into the dynamics, we were able to establish

a relation between the mixing time to stationarity and the *median* time to fixation in the limit of vanishing mutation rates (in the right circumstances). No such relation had previously been investigated.

Metastable States in a Model of Cancer Initiation

In Chap. 5 we turned our attention to a well-studied model that describes the accumulation of two mutations in a tissue during the initiation of cancer. Previous studies had been restricted to the regions of parameter space which permitted a coarse-graining approach to reduce the complexity of the problem. Through analysis of the deterministic equations of motion, we identified parameter regimes where this approach is not valid because of the existence of quasi-equilibria, or metastable states. It is in these regimes that we applied the WKB method, a technique from mathematical physics that relies on the interplay of fixation and equilibration. In these systems with metastable states, there exists a separation of timescales. The system quickly relaxes to a quasi-stationary distribution, but fixation only occurs on a longer timescale. We were able to exploit this separation to compute the quasi-stationary distribution about the metastable state, and then we could use this to compute the expected fixation time. When the system reduced to one dimension (a consequence of having absorbing boundaries at the edge of our state space) we obtained closed-form expressions for the quasi-stationary distribution and the mean fixation time. If the system does not reduce to one dimension, we rely on numerical methods to extract the fixation time statistics.

Through our analysis we identified that the phenomenon of stochastic tunnelling is, in fact, a deterministic effect. Tunnelling is the process that describes how a homogeneous population of wild-type cells can evolve into a homogeneous population of cells with two mutations without visiting the state in which all cells harbour only one mutation [2]. Furthermore, our analysis identified the escape from the metastable states as the key bottleneck to fixation of cells with two mutations. For parameter values for which there are no metastable states (i.e. when cells with two mutations have the highest fitness, as would be expected from the inactivation of an oncogene [3]), the fixation dynamics is largely governed by the deterministic flow. The rate-limiting steps are then the appearance of successful mutant lineages [4], and the subsequent fixation of cells with two mutations is a zero-hit process for large population sizes. As such the progression from healthy tissue (no mutations) to susceptible tissue (two mutations, corresponding to an inactivated tumour suppressor gene) will be fast relative to the cases in which a metastable state exists. If there is one stable fixed point in the deterministic dynamics, the process becomes a one-hit phenomenon limited by the escape from the corresponding metastable state. In regions with two fixed points one observes a two-hit process; the population becomes trapped in a first metastable state, escapes to a second metastable state, and then reaches full fixation. Our analysis allowed us to classify how changes to the fitness landscape, mutation rates, and population size affect the fixation time of cells harbouring two mutations, as well as the probability of tunnelling.

In terms of the development of tumours, our analysis shows that the path to accumulating mutations is not simply limited by the mutation rates, but also by the escape from metastable states. Populations can exist in a heterogeneous state for very long periods of time before fluctuations eventually drive the second mutation to fixation. The probability with which stochastic tunnelling occurs is, in part, determined by the location of these metastable states. If they are located close to the homogeneous state with all cells harbouring one mutation, then the probability of tunnelling is low. This work has filled the gap left by the existing literature and leads to a more comprehensive understanding of mutation acquisition and stochastic tunnelling in evolving populations.

The WKB Method

We used the WKB method as an ‘off the shelf’ tool to analyse this model of mutation acquisition. In conducting this study we found inconsistencies and confusion among the existing literature, and this prompted further investigation into the origins of this method. In Chap. 6 we took a closer look at the WKB method. We explored the mathematical basis of this approach, and illustrated the different constructs that arise when it is applied to stochastic systems. Through toy models and some well-studied systems, we demonstrated the meaning of the multitude of terms that come from this field, as well as making connections with the related path-integral formulations and the theory of large deviations. We discussed the different numerical methods that can be used to solve these problems, and gave examples of the quasi-potential landscapes that can be computed.

The WKB approach improves on the ‘traditional’ methods of analysis of individual-based systems; numerical integration of the master equation can be very inefficient and diffusion approximations only yield accurate results for the bulk of the probability distribution, failing to capture the rare events which are characterised by the tails. By mapping the problem from the intractable stochastic process to classical mechanics, the WKB method provides some much-needed intuition about the population dynamics in systems described by a large number of random variables. This approach captures the large deviations from the expected behaviour, which play a crucial role in describing events such as extinction or switching between metastable states.

7.3 Outlook

Through these investigations we hope to have promoted the benefits of analytical and semi-analytical methods. In this age of burgeoning computer power, the temptation to rely on simulations alone is huge. But analytical approaches can greatly complement this data, allowing us to identify the crucial components of a model that give rise to observed phenomena. They also allow us to extrapolate to inaccessible parameter regimes, such as large system-sizes, and highlight possible parameter combinations where interesting effects may be observed. They can even tell us what we need to look

for in the experimental or computational data; measuring a most-likely path from simulation data would be very tricky without prior knowledge of what that object is. Analytical treatments can also lead us to more efficient simulation procedures. For example, in this thesis we have documented how mathematical manipulation of the master equation describing the birth–death process allows us to sample arrival times much faster than with direct simulations alone.

On a more general level, constructing a mathematical theory of evolutionary dynamics is very much work in progress. Nature is inherently discrete and stochastic, and fluctuations must be taken into account when trying to understand these systems. An integral part of the evolution of microbes and higher organisms alike is frequency-dependent selection [5–7], as found in evolutionary games. At the same time external factors determining the detailed mechanics of selection may vary throughout the duration of the process. In this thesis we have combined frequency-dependent selection, fluctuating environments, and stochastic dynamics in discrete populations into a single model, and we have provided the analytical tools for its analysis. This, we hope, is a contribution toward a more complete understanding of evolutionary processes.

The work presented in this thesis can be further developed in many ways. Individual-based processes can be used to describe systems from a broad range of disciplines, and the list of possible applications is endless. The methods that we have developed, which include tools that can handle additional sources of stochasticity and a solid framework for applying the WKB approach, now allows us to analyse more models than ever before.

Evolution in a changing environment is a current topic of great interest. The field of evolutionary rescue is dedicated to describing how populations adapt to sudden change [8], as observed when the environment switches its state. We expect that the theoretical framework discussed in Chap. 3 could be used to describe these scenarios, which arise, for example, when drugs are administered to fight diseases or infections, or when ecological niches are disturbed through anthropogenic effects.

We are currently in the era of ‘big data’, and as such we are surrounded by the results of experiments, surveys, and numerical simulations. Studies often report only the mean statistics of these data, and maybe the associated standard deviation. As the volume of data grows, analysing the full distribution becomes more interesting and meaningful. The method discussed in Chap. 4 describes the first derivation of exact distributions in scenarios from evolutionary game theory. We see this as a leap forward in analytical predictions, and we expect to see these techniques become increasingly popular when describing evolutionary processes.

In some situations, it takes somebody from outside a field to provide new insights. The influence of physicists is felt across a range disciplines, mostly for the right reasons. Sometimes we must depart from the traditional line of thinking and consider new approaches. Our investigation of the well-studied cancer initiation model is an example of this. The role of physics in the investigation of cancer is stronger than ever, and is likely to increase as we seek a more complete understanding of this disease. Extending our model to investigate the dynamics of the system beyond the

second mutation, or the emergence of a cancerous phenotype prior to fixation, are likely directions for the future of our work.

On a personal level the work presented in this thesis, and in particular the modelling of cancer in Chap. 5, has fuelled my passion to study biologically-motivated models. Using the techniques of statistical physics, I hope to obtain further insights into these fundamental processes of evolution and population dynamics. While doing so I will be flying the flag for the analytical methods of maths and physics.

References

1. M. Acar, J.T. Mettetal, A. van Oudenaarden, Stochastic switching as a survival strategy in fluctuating environments. *Nat. Genet.* **40**, 471 (2008)
2. N.L. Komarova, A. Sengupta, M.A. Nowak, Mutation-selection networks of cancer initiation: tumor suppressor genes and chromosomal instability. *J. Theor. Biol.* **223**, 433 (2003)
3. R.A. Weinberg, *The Biology of Cancer* (Garland Science, New York, 2013)
4. M.A. Nowak, F. Michor, N.L. Komarova, Y. Iwasa, Evolutionary dynamics of tumor suppressor gene inactivation. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 10635 (2004)
5. R.C. MacLean, A. Fuentes-Hernandez, D. Greig, L.D. Hurst, I. Gudelj, A mixture of “cheats” and “co-operators” can enable maximal group benefit. *PLoS Biol.* **8**, e1000486 (2010)
6. J. Gore, H. Youk, A. van Oudenaarden, Snowdrift game dynamics and facultative cheating in yeast. *Nature* **459**, 253 (2009)
7. R.C. Maclean, I. Gudelj, Resource competition and social conflict in experimental populations of yeast. *Nature* **441**, 498 (2006)
8. S.M. Carlson, C.J. Cunningham, P.A. Westley, Evolutionary rescue in a changing world. *Trends Ecol. Evol.* **29**, 521 (2014)