

Differential Equations & Dynamical Systems
Volume 5

**Modelling
Disease Ecology
with Mathematics**

2nd edition

Stacey Smith?



American Institute of Mathematical Sciences

Stacey Smith?

The University of Ottawa

Modelling Disease Ecology with Mathematics

Free online edition



American Institute *of* Mathematical Sciences

EDITORIAL COMMITTEE

Editor in Chief: Alain Miranville (France)
Jerry L. Bona (USA), Tomás Caraballo (Spain)
Zhaosheng Feng (USA), Maurizio Grasselli (Italy)

Stacey Smith?

Department of Mathematics and Statistics
The University of Ottawa
150 Louis-Pasteur Pvt
Ottawa, ON K1N 6N5, Canada
E-mail: stacey.smith@uottawa.ca

AMS 2000 subject classifications: 37N25, 46N60, 47N60, 62P10, 92B05, 34-01,
35-01, 65D05, 65D07, 37G10, 37G15, 15-01, 65F15, 41A58.

ISBN-10: 1-60133-020-0; ISBN-13: 978-1-60133-020-8

© 2023 by the author after reversion of rights. Original publisher: the American Institute of Mathematical Sciences (2017).

Made available for free online so that students everywhere can benefit.

aims sciences.org

Contents

1	Preface	1
1.1	Origins	2
1.2	What's new in the second edition	3
1.3	The free online edition	3
1.4	Acknowledgements	4
1.5	A note about zombies	4
<hr/>		
I	Introductory material	
2	Introduction	9
3	Fitting curves to data	13
3.1	Model fitting vs. interpolation	13
3.2	Sources of error in the modelling process	17
3.3	Visual fitting with the original data	18
3.4	Transforming the data	20
3.4.1	Regression coefficients can be misleading	23
3.4.2	Transformations can also be misleading	23
3.5	Lab work	28
3.5.1	Exercises	30
4	Splines	33
4.1	Spline interpolation	33
4.2	Linear splines	34
4.3	Cubic splines	36
4.3.1	An example	38
4.3.2	Generalising	41
4.4	Lab work	41
4.4.1	Exercises	41

5	Simple epidemic models	45
5.1	SIS epidemic	45
5.1.1	Solving directly	48
5.1.2	Phase portraits	50
5.2	SIR epidemics	54
5.3	Lab work	56
5.3.1	Exercises	59
6	Calculating R_0	63
6.1	Threshold parameters	63
6.2	The method of Anderson & May	64
6.3	The Jacobian	65
6.3.1	The SIS model threshold	67
6.3.2	The endemic equilibrium for the SIS model	68
6.3.3	The SIR model threshold	69
6.3.4	The endemic equilibrium for the SIR model	70
6.4	Other methods for calculating R_0	71
6.4.1	Survival function	71
6.4.2	Next-generation method	72
6.4.3	Average age at infection	75
6.4.4	The final-size equation	76
6.4.5	Calculation from the intrinsic growth rate	76
6.5	Vectorial capacity	77
6.6	Lab work	78
6.6.1	Exercises	78
7	A vector-borne disease with lifelong immunity	81
7.1	Building a vector-borne disease model	81
7.2	Finding eigenvalues	83
7.3	Deriving a threshold condition	84
7.3.1	The Routh–Hurwitz criterion	85
7.4	Total population behaviour	86
7.5	Lab work	87
7.5.1	Exercises	89
8	The spread of measles	93
8.1	The Conservation Law	93
8.1.1	Diffusion	96
8.2	Measles in a corridor	97
8.3	Lab work	99
8.3.1	Exercises	106

9	Solving Partial Differential Equations	109
9.1	Separation of variables	109
9.2	Boundary conditions	112
9.3	Initial conditions	113
9.4	Lab work	114
9.4.1	Exercises	114
10	The discrete logistic equation	117
10.1	Developing a nonlinear model	117
10.1.1	Spatial limitations	117
10.1.2	Difference equations	118
10.2	Cobwebbing	121
10.2.1	Stability in the discrete logistic equation	122
10.2.2	More complex behaviour	124
10.3	Lab work	126
10.3.1	Exercises	128
11	Bifurcations	131
11.1	Stability in difference equations	131
11.1.1	The range $0 < r < 3$	132
11.1.2	The range $3 < r < 3.45$	133
11.1.3	The range $r > 3.45$	136
11.2	The <i>Doctor Who</i> effect	137
11.3	Lab work	137
11.3.1	Exercises	138

II Advanced material

12	More advanced epidemic models	140
12.1	The infection rate	140
12.2	Demography	142
12.3	The effects of media	145
12.4	Hospital bed capacity	146
12.5	Lab work	147
12.5.1	Exercises	151
13	Measles with vaccination	154
13.1	The model	154
13.2	Finding equilibria	155
13.3	Stability in a discrete model	156
13.3.1	Stability in one dimension	157
13.3.2	Stability in two dimensions	157
13.4	Lab work	160
13.4.1	Exercises	162

14 A disease with an asymptomatic class	166
14.1 Asymptomatic infection	166
14.1.1 Using the characteristic polynomial	167
14.1.2 Using the next-generation method	170
14.2 Lab work	171
14.2.1 Exercises	171
15 Impulsive Differential Equations	174
15.1 Introduction	174
15.1.1 Fixed-time impulse example	174
15.1.2 Nonfixed-time example	177
15.2 Motivation	179
15.3 Impulsive semidynamical systems	180
15.3.1 Existence of solutions	184
15.3.2 Definitions of stability	184
15.3.3 Autonomous systems with impulsive effect	185
15.4 Floquet Theory in \mathbb{R}^2	187
15.4.1 Floquet Theory for continuous systems	187
15.4.2 Floquet theory	187
15.4.3 Orbital stability in \mathbb{R}^2	190
15.5 Lab work	194

III Case studies

16 Application: AIDS and end-stage renal disease	199
16.1 Determining AIDS prevalence	200
16.2 Determining end-stage renal disease prevalence	203
16.3 Model fitting	207
16.4 Using the model to predict future outcomes	208
16.5 Lab work	210
16.5.1 Exercises	210
17 Application: Malaria with a time delay	213
17.1 A delay differential equation model of malaria	213
17.2 Lab work	217
17.2.1 Exercises	217
18 Application: Guinea-worm disease	221
18.1 Introduction	221
18.2 The model	222
18.3 The system without impulses	223
18.4 The system with impulses	225
18.4.1 Fixed chlorination	226
18.4.2 Nonfixed chlorination	227

18.5	Numerical simulations.....	228
18.6	Discussion	232
18.7	Lab work	233
19	Application: Zombies! (Aargh!)	237
19.1	Introduction — of peril!	237
19.2	The model — of doom!	238
19.3	Analysis — of anguish!	239
19.4	Incubation — of destiny!	241
19.5	Quarantine — of terror!	243
19.6	A cure — of fear!.....	246
19.7	Using our braaaiiinnssss! (impulsively)	248
19.8	Discussion — of dread!	249
19.9	Lab work — of horror!	250

IV Appendices

Solving the time series directly	254
Partial fractions	257
Eigenvalues	258
The R_0 sleight of hand	260
Finding eigenvalues for the case of permanent immunity	262
Integrating factors	264
Trivial solutions for nonnegative constants	266
Taylor's Theorem	269
Stability of periodic orbits in the logistic equation	271
One-dimensional discrete stability condition	273
The lower bound for λ_5	275
Solving the end-stage renal disease equations	277
Matlab overview	279
M.1 Variables and operators	279
M.2 M-Files	280
M.3 Figures.....	281
M.4 Loops	282

M.5 Solving ODEs	283
M.6 Solving delay differential equations	283
M.7 Glossary of other terms	284
References	285
Index	289

List of Figures

2.1	The mathematical modelling cycle.	11
3.1	Data from the 1918 influenza pandemic.	15
3.2	Interpolating the data using a smooth curve.	16
3.3	Fitting a parabola to the data points.	16
3.4	Each data point is thought of as an interval of confidence.	18
3.5	Minimising the sum of the absolute deviations	19
3.6	Minimising the largest absolute deviation from the fitted line. . .	20
3.7	Plot of collected data for new HIV infections in 1981.	21
3.8	Plot of y versus e^x for the original data.	21
3.9	Plot of $\ln y$ versus x for the transformed data.	22
3.10	Avian influenza data.	24
3.11	Log transformation of avian influenza data	24
3.12	Exponential vs. original curve	25
3.13	Biannual outbreaks of influenza.	26
3.14	The transformed data points and a line of best fit.	26
3.15	A plot of the curve $y = Ce^{1/x}$ based on the value $\ln C \approx -1.25$. .	27
4.1	A linear spline model	34
4.2	A cubic spline model.	37
4.3	The natural cubic spline.	41
5.1	An SI model.	46
5.2	An SIS model.	47
5.3	The time series of the infection.	50
5.4	The linear relationship between S and I	51
5.5	Equilibria for the two cases.	52
5.6	The stability of the equilibrium points.	53
5.7	An SIR model.	54
5.8	An SIR model with vaccination.	55
5.9	An SIR model with mutation.	55

5.10	Phase portrait for $b = 10$	58
5.11	The epidemic diagram for the SIMD model.	60
6.1	A example of a two-stage basic reproductive ratio.	75
7.1	A vector-host model for a disease with permanent immunity.	82
7.2	The roots of a quadratic.	84
7.3	Infected mosquitos vs. infected humans.	90
8.1	The total flow of droplets in a segment of space of length Δx	94
8.2	The flow of droplets within a small length of space.	94
8.3	Approximating a thin area by a rectangle	96
8.4	The trajectory of an airborn measles droplet	98
8.5	Average trajectory of droplets over time.	99
10.1	The previous value vs. the updated value	120
10.2	Iterating the dynamics.	121
10.3	Cobwebbing	122
10.4	Logistic map for $r = 0.5$	123
10.5	Logistic map for $r = 2$	124
10.6	Logistic map for $r = 3.2$	125
10.7	Logistic map for $r = 4$	125
11.1	Bifurcation diagram for $0 < r < 3$	133
11.2	Equilibrium points of $g(x)$	134
11.3	Bifurcation diagram for the appearance of Period 2 points.	135
12.1	Constant growth, linear death.	143
12.2	Logistic growth.	144
12.3	The media-affected transmission function	146
12.4	The media-affected transmission function after an initial delay .	146
12.5	The recovery rate $\mu(b, I)$ with $\mu_0 = 1$ and $\mu_1 = 5$	148
12.6	The recovery rate as a function of I	148
12.7	Standard incidence versus mass action	150
13.1	The progression of measles.	154
13.2	The addition of vaccination.	155
13.3	Solutions oscillate around the endemic equilibrium.	162
13.4	Solutions oscillate towards the endemic equilibrium	163
13.5	Solutions move directly towards the disease-free equilibrium .	163
14.1	An influenza model with an asymptomatic class	167
15.1	Drug concentrations with impulses at fixed times	176
15.2	Drug concentrations with impulses at nonfixed times	179
15.3	When $\alpha < 1$, the impulsive orbits approach part of the x -axis. .	183

15.4 When $\alpha > 1$, the impulsive orbits converge to a fixed point.	183
15.5 The periodic orbit of first order.	195
15.6 The periodic orbit of third order.	196
16.1 AIDS deaths among African Americans in the US	201
16.2 AIDS prevalence among African Americans in the US	202
16.3 Linear fit to pre- and post-ART AIDS data	202
16.4 Linear fit to the entire AIDS data set.	203
16.5 End-stage renal disease mortality	204
16.6 End-stage renal disease prevalence	205
16.7 Linear fit to pre- and post-ART end-stage renal disease data	206
16.8 Linear fit to the entire end-stage renal disease data.	206
16.9 ART-blocking effects on disease progression.	209
17.1 A example of a two-stage basic reproductive ratio.	215
17.2 Infected humans and mosquitos.	217
18.1 The Guinea-worm disease model	223
18.2 The overestimate versus the actual parasite value.	225
18.3 Eradication threshold surface	229
18.4 Persistence of the disease under annual chlorination.	231
18.5 Eradication when the parasite birth rate is decreased	231
19.1 The basic zombie model.	239
19.2 Basic zombie outbreak scenario	240
19.3 The basic model with latent infection.	242
19.4 An outbreak with latent infection.	244
19.5 The model with quarantine.	245
19.6 With quarantine, the entire population collapses, and we all lose.	246
19.7 The model with treatment.	246
19.8 With treatment, humans eventually survive	247
19.9 Impulsive attacks can control the epidemic.	248
K.1 A bounded, concave-down function with negative initial value . .	276

List of Tables

3.1	Collected data for new HIV infections in 1981.	20
3.2	Transformed data	22
3.3	Initial infections and cull size for an avian influenza outbreak...	23
3.4	Diameter of a ponderosa pine versus volume.....	28
3.5	Mean number of bugs per infected site	30
4.1	Data for MRSA, which increases sharply one year.....	35
4.2	A linear spline model.....	36
4.3	A natural cubic spline model.	40
17.1	Variable and parameters for the malaria model with delay	214
18.1	Guinea-worm disease parameters	232

1

Preface

The primary purpose of these lecture notes is to use disease modelling as a framework in which to explore the usefulness that mathematical modelling can have. These notes are mainly intended for students and researchers in the biological sciences who may have limited mathematical background, but are also accessible to a wider audience who may have an interest in mathematical biology, but not necessarily expertise in either.

This monograph is intended to bridge the gap, at least in part, between biology and mathematics. While mathematical biology has long been an important aspect of epidemiology, many students and researchers in the biological sciences have stated that they went into biology “to get away from math”. Increasingly, there has been an awareness that it is no longer possible to be a biologist and still avoid aspects of mathematical modelling, curve fitting and understanding what computer transformations are actually doing, rather than just applying them. This trend will only continue in the future.

This monograph assumes a basic level of calculus, although not necessarily a level that is at the reader’s fingertips. Thus, most topics are introduced or refreshed when they occur and several may be found in the appendices. However, while differential equations and matrices are introduced here, a complete overview is beyond the scope of these notes.

Each chapter comes complete with computer labs that explore the usefulness of Matlab. While it is expected that Matlab is unlikely to be the biologist’s natural choice of computer program, it was chosen for its accessibility and ease of adaptation. Recent innovations in the graphical user interface in Matlab 7 and beyond have made this program an invaluable tool to most biomathematicians. A working knowledge of the program would assist future collaborations between biologists and applied mathematicians.

However, it is also understood that the reader may not be interested in learning much about a program that he or she may never use in practice. Thus, the labs have been designed so that little memorisation of the program’s ability is needed; most of the computer codes are provided, so that the emphasis is on

using the tool, as an interface between the theory and application. Problems and solutions are worked through in their entirety to get a feel for Matlab before the actual exercises are embarked upon. The focus is primarily on applications, rather than theory.

This monograph is structured into three parts: 1) introductory material that the general reader can access 2) advanced material for the more mathematically inclined and 3) several case studies. The style is an accessible one; although much of the content is mathematical, there has been extensive proof-reading by epidemiologists, the exercises worked through by graduate students and refined for accessibility and educational insight. To aid in this accessibility, each chapter is preceded by a flowchart that outlines the relevant points and connections between ideas.

Note also that the spelling throughout conforms to the British, rather than American, style. And yes, the author's name does have a question mark in it, so please don't send her emails about it.

Chapters 5, 8, 10 and 11 are adapted from LaPeare *et al.* (1998). Chapters 3 and 4 are adapted from Giordano & Weir (1985). Much of chapter 6 is adapted from Heffernan *et al.* (2005). Chapter 15 is adapted from Bainov & Simeonov (1989, 1993). Chapter 16 is adapted from Schwartz *et al.* (2005), Chapter 17 is adapted from Ruan *et al.* (2008), Chapter 18 is adapted from Smith? *et al.* (2012) and Chapter 19 is adapted and updated from Munz *et al.* (2009). However, in each case, the chapters have been extensively rewritten in order to promote the accessibility that is so desperately needed to bring the disparate fields of mathematics and biology closer together.

1.1 Origins

The material for these notes was originally developed for a workshop on modelling disease ecology with mathematics. The workshop was hosted by Ricardo Görtler and held at the Faculty of Exact and Natural Sciences, University of Buenos Aires, Argentina, July 11–15, 2005.

The development of the workshop was funded by an award from the NIH/NSF Ecology of Infectious Disease program (Award RO1 TW05836 funded by the Fogarty International Center and the National Institute of Environmental Health Sciences (NIEHS) to Uriel Kitron and Ricardo Görtler)

This workshop was conducted with support from the NIH/NSF Ecology of Infectious Disease program (Award RO1 TW05836 funded by the Fogarty International Center and the National Institute of Environmental Health Sciences (NIEHS) to Uriel Kitron and Ricardo Görtler) and an award from the U.S. National Aeronautics and Space Administration (NASA) and the National Institute of Allergy and Infectious Diseases (NIAID), a component of the National Institutes of Health, U.S. Department of Health and Human Services to Charlie King, Uriel Kitron and Ricardo Görtler.

1.2 What's new in the second edition

The first edition was a popular textbook with a diverse audience. Although it is a math textbook, it is written in a friendly and accessible style, with mathematical details relegated to various appendices, in order that the non-mathematician be able to access the book.

Consequently, the book has been used as a textbook in a number of courses, as well as for workshops for non-mathematicians (e.g., the Public Health Agency of Canada, which is the Canadian government's disease department). Students have often reported that they learned everything from it and that it is their "bible", often just calling it "the green book" as a shorthand.

One of the most useful aspects has been the inclusion of Matlab codes, meaning students can get started on programming straight away, using hands-on learning rather than figuring everything out from scratch. These codes were workshopped in advance and are thus error-free.

This edition is no small update of the original. More than a third of the material is new, while every chapter has been re-edited and re-worked, in small or large ways. Errors have been fixed and new takes on existing material have been offered after feedback from students. The material has been re-ordered to begin with data fitting after this structure was found to be an ideal opening in the associated course. Mathematical details in Chapter 5 have been moved to an appendix, as this was offputting for introductory readers.

The material has been divided into several parts. The bulk of the original version constitutes the first part, which is accessible to a general audience. The second part consists of new chapters on advanced model construction, detailed discrete-time modelling, a disease with an asymptomatic class and a chapter on impulsive differential equations. The latter is an exciting new branch of mathematics that has useful and grabbable applications for the non-mathematician, but books on the subject are not very accessible. While the introductory chapter here gets somewhat dense, it nevertheless contains motivating examples to explain the process to the general reader.

Several new case studies have been added that illustrate new theoretical chapters. All chapters come with Matlab codes and exercises, with appendices where applicable. Some smaller additions include a discussion of the Routh–Hurwitz criterion in Chapter 7 and a list of the top ten diseases of all time in Chapter 16. The latter

1.3 The free online edition

This textbook has been made freely available for use by anyone who may benefit from it. Various typos have been corrected, but if you find others, feel free to let me know, and I'll update them.

1.4 Acknowledgements

The author is grateful to Elissa Schwartz, Shigui Ruan, Shoshana Magnet, Rachelle Miron, Kristina Donato, Lilian Tu, Jane Heffernan, Marc Beauparlant and Jianhong Wu for advice and assistance. The original lecture notes were developed with assistance from Uriel Kitron, Gideon Wasserman, Riccardo Gürtler, Anna Schotthoefer and Charlie King. Assistance by members of Ricardo Gürtler's laboratory is gratefully acknowledged. The author is supported by an NSERC Discovery Grant, an Ontario Early Researcher Award and funding from MITACS.

In the first part of this book, the more dense mathematics has been relegated to the various appendices. Again, this makes the text more accessible, without the reader getting bogged down in mathematical details. These appendices are in no way intended to be a complete overview of their mathematical topic; rather, they supplement the relevant chapters and provide an introduction to the usefulness of certain topics. The interested reader is encouraged to seek out mathematical textbooks on the relevant subject for further information. It is hoped that these notes will provide a stimulating introduction to the power of mathematics to understand the world around us, while still maintaining that, despite the math-phobia that exists in our society, mathematics can nevertheless be as accessible and as enjoyable as any other science.

1.5 A note about zombies

A final word on the inclusion of a chapter on zombies in this book, which may raise some eyebrows. Its inclusion here is not simply a reproduction of an immensely popular article (it was, for a time, the #1 pdf on Google and was, for 24 hours, the most popular news story in the world), but a reworking with a revised model and some new conclusions.

Since I first published an academic article on zombies (which went on to win a Guinness World Record), I have been amazed by the power that mixing science and media has had in inspiring people. Those who would never otherwise look twice at an equation read an entire academic article because they were so transfixed by zombies. Many people had never even heard of the discipline of disease modelling until the zombie article came out.

By using both real-life problems and also pop culture (such as zombies), mathematics can be made both fun and accessible to non-mathematicians. Where physics has a popular following — concepts such as the Higgs Boson, black holes or quarks excite the popular imagination, even among people who may never understand the details — mathematics struggles to make its concepts similarly appealing beyond a niche audience. Zombies break this trend, allowing people who will never solve an equation to be engaged and excited by mathematical modelling.

I make no apologies for this. My efforts to promote mathematics to the wider world (including both zombies and the first edition of this textbook) won me a prestigious Partners in Research award as mathematics ambassador, and they have inspired people to go into disease modelling, raising its awareness as a discipline. Rather than keeping mathematics in the ivory tower, by taking it to where people actually are, it's possible to make the subject engaging, fresh and amusing. If we can harness this in a broader sense, we stand a chance of breaking down the barriers that hold mathematics back from being accessible to all and widely accepted in the greater community. This can only be a good thing. It is definitely something that is worth striving for.

Stacey Smith?
Ottawa, Canada
October 2023

Part I

Introductory material

Introduction

Infectious diseases have been a part of the human condition since time immemorial. Some, such as herpes or chicken pox, usually have mild symptoms and vanish of their own accord. Others, such as HIV, tuberculosis or malaria, are responsible for millions of deaths each year. Diseases have been the source of fear and superstition throughout the ages; sociological upheavals following the black plague led to the demise of the church as an all-powerful institution, the destruction of the serf system and subsequent creation of labour movements, and the creating of colonialism. Malaria is likely the greatest killer of human beings of all time, being thought to have killed one in two humans who ever lived (see Page 199). The 1918 Spanish influenza pandemic killed 50–100 million people in the space of six months (about 5.5% of the entire world; this would be equivalent to 350 million deaths today).

Mathematical epidemiology has its roots in 1760, when Daniel Bernoulli formulated and solved a model for smallpox. In 1906, Hamer used a discrete-time model of measles to understand recurrent epidemics. In 1911, Sir Ronald Ross used mathematical models to help demonstrate that malaria was spread by mosquitos (later winning a Nobel prize for his work). This work was later followed up by public-health physician William Kermack and biologist Anderson McKendrick in 1927, who showed that mortality rates were much more dependent on the year of birth than the year of death, as previously thought. This was particularly impressive, since Kermack had been blinded three years earlier and did all the mathematics in his head.

Since the middle of the twentieth century, mathematical epidemiology has grown exponentially. The appearance of HIV/AIDS saw a further explosion in the growth of models, as ways of predicting the spread of the disease became crucial. Mathematical models were used to design the triple-drug cocktail, responsible for saving the lives of millions of people. In recent years, mathematics has been used to determine intervention strategies for the foot and mouth outbreak in Britain, the SARS epidemic, human papillomavirus vaccination and West Nile virus. Mathematics can determine the critical proportion of

individuals who must be vaccinated against a disease, the minimum amount of drug required to treat an infection, the frequency of use of an insecticide to control mosquitos or the number of doses of a drug regimen that may be missed before drug resistance emerges.

Mathematical modelling has become a tool that is now widely used to study problems and questions in public health. Such models have focussed on the spread of an infectious agent in a population, the pathogenesis of an infectious agent in a host and its effect on the immune system, the growth and spread of tumour cells in a body, and even the economics of vaccine and drug delivery at the population or individual level. However, public-health workers rarely have the mathematical skills to develop informative models for public-health policy.

In recent years, the rise of detailed electronic surveillance of infectious disease and improvements in computing power, electronic data management, rapid diagnostic tests, genetic sequence analysis and the ability to share and deposit data over the internet has given rise to a vastly improved ability to monitor the spread of diseases. Statistical methods have been integrated into the modelling process, allowing estimation of key parameters and the testing of hypotheses using available data.

Mathematical modelling is analogous to map-making. Without a decent map, even finding our way home is difficult, let alone exploring new avenues. Like any good map, we aren't trying to create a perfect replica of reality; rather, we hope to elucidate key features and ignore what isn't relevant. Often, what we leave out is as important as what we include. To solve biological problems, first we translate the biology into a mathematical model. (Despite appearances, usually borne out of math phobia, this is often the most challenging part of the process.) Once we are satisfied with the model, we perform mathematical analysis. Assuming we do this correctly, we can be confident that what happens here is a series of entirely logical steps. Our mathematical conclusion then needs to be translated back to a biological conclusion. Since the process of mathematical analysis is rigorous, it follows that the strength of our conclusions are wholly dependent upon the formulation of the mathematical model. See Figure 2.1. Thus it is vital to understand the mechanisms, as well as the limitations, of this modelling cycle.

More formally, mathematical modelling of infectious disease is a framework designed to convey ideas about the components of host-pathogen interactions. It requires a clear understanding of the interaction between the infectious agent and the host within an individual, the mode and rate of transmission between individuals and host-population characteristics, such as demography and behaviour. Models can be used to determine the dominant factors that generate observed patterns and phenomena. They can aid data collection, interpretation and parameter estimation. They can identify possible approaches for disease control and for assessing the potential impact of different intervention measures. Importantly, they can do all this without costly and time-consuming experiments.

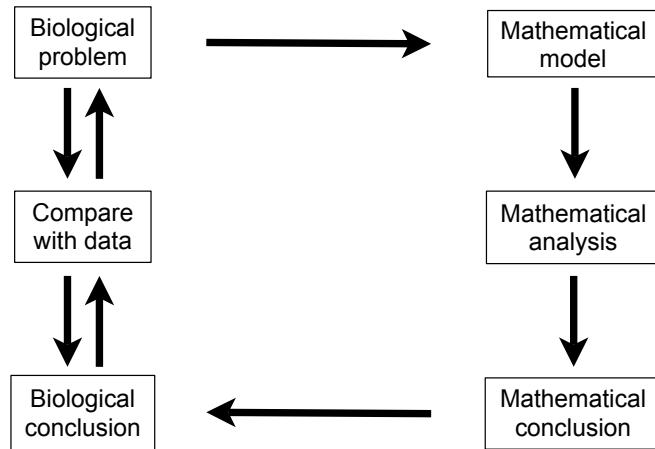
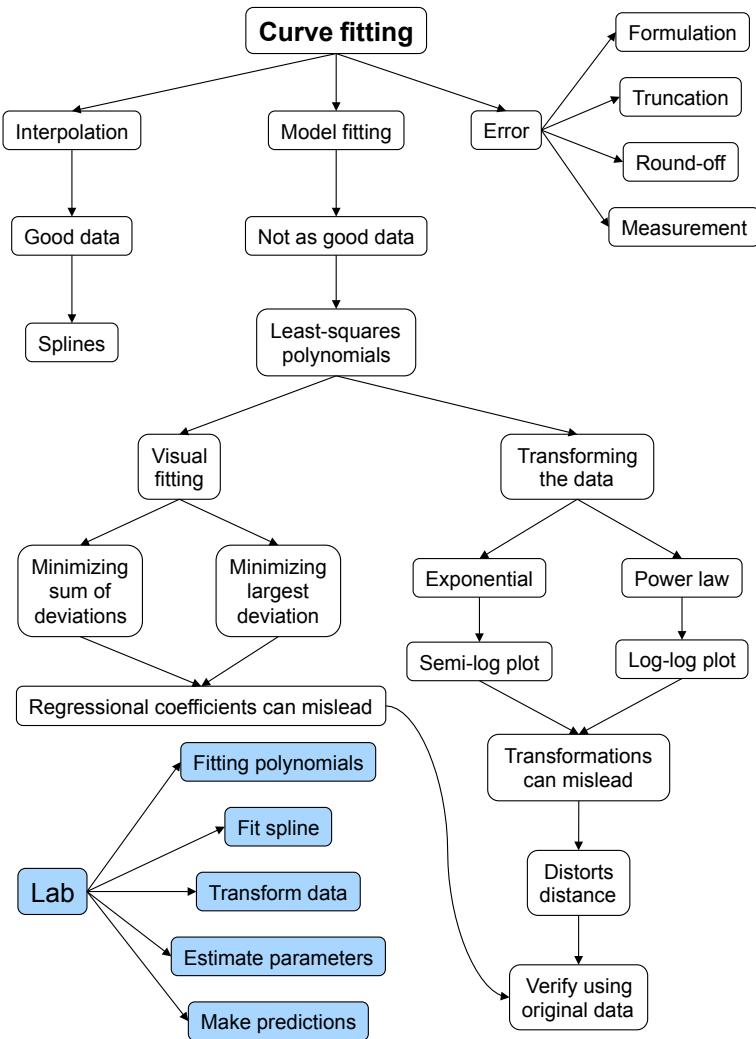


Fig. 2.1. The mathematical modelling cycle.

This monograph facilitates the learning of mathematical biology for public-health researchers and mathematical modellers, who may have a range of skills and knowledge drawn from mathematics, clinical medicine, sociology, economics, decision-making, risk analysis, psychology, industrial production capacity and politics. There needs to be a symbiosis between these areas; people working in these disciplines need to communicate to each other effectively so that key questions in health can be addressed.

This monograph specifically introduces mathematical modelling to those in the public-health sector who have had no formal training in mathematics. It allows people working in a range of disciplines to understand the strengths and weaknesses of models. Mathematical models are formulated that represent current diseases; the tools needed to analyse these models are introduced as necessary. Through Matlab exercises, numerical simulations are developed that provide the reader with the ability to formulate control strategies, test hypothetical interventions and explore disease-management options. The exposition of the material is mainly addressed to researchers and graduate students interested in the application of mathematics to biological problems.

One of the objectives is to familiarise public-health workers with questions and outcomes that can be produced by simple mathematical models. Ultimately, policy-makers need to understand the importance and limitations of models, in order to be confident in basing decisions on them. Mathematicians need to be able to connect the equations to the biology. And, crucially, both groups need to be able to talk to each other. Only then will we make real progress in solving some of the world's most difficult problems.



3

Fitting curves to data

Fitting curves to data is one of the most common things we can do when analysing our data. There are all manner of ways in which various curves can be fitted, leading to the question of which is the “best”, under a variety of circumstances. Curves are fit to data all the time in science, especially using linear regression. However, as we’ll see, it’s not always as straightforward as it seems.

By the end of this chapter, you should know how to use Matlab to fit basic curves, how selecting a model depends on the confidence you have in your data, some sources of common error and some of the potential pitfalls associated with computers fitting data blindly.

3.1 Model fitting vs. interpolation

When analysing data, we can use information that data implies to formulate mathematical models. These models rely on assumptions about the data, or about the data we have not collected. We may encounter situations in which there are different assumptions leading to different models. For example, as an influenza pandemic moves through a population, we could make assumptions about the heterogeneous mixing of the population being proportional to the local population density or the urban vs. rural environment, or we could ignore the heterogeneity of the population altogether.

We may be faced with using collected data to determine unknown parameters in our model in a way that selects the curve from each model that “best fits” the data and then choose whichever resultant model is most appropriate for the particular situation under investigation.

A different case arises when the problem is so complex as to prevent the formulation of a model explaining the situation. For instance, a three-dimensional model for the spread of measles might involve partial differential equations for the movement of infectious droplets in three spatial dimensions, plus one temporal dimension. This will not only be enormously complicated, but the

equations may not even be solvable, so there is little hope for constructing a master model that can be solved and analysed analytically. Or there may be so many variables that one would not even attempt to construct an explicit model.

In such cases, experiments may have to be conducted to investigate the behaviour of the system. Then the experimental data can be used to predict the outcome, but only within the range of the data points collected.

The preceding discussion identifies three possible tasks when analysing a collection of data points:

1. Fitting a selected model type or types to the data. For example, applying a line of best fit to known data points.
2. Choosing the most appropriate model from competing types that have been fitted. For example, we may need to determine whether the best-fitting exponential model is a better model than the best-fitting least-squares model.
3. Making predictions from the collected data. This may involve interpolation (predicting in between known data points) or extrapolation (predicting a point outside the range of collected data).

In Task 1, the precise meaning of “best” model must be identified and the resulting mathematical problem resolved. In Task 2, a criterion is needed for computing models of different types. In Task 3, criteria must be established for determining how to make predictions in between the observed data points or outside the ranges of what we know.

In Tasks 1 and 2, a relationship of a particular type is strongly suspected, and the modeller is willing to accept some deviation between the model and the collected data points in order to have a model that satisfactorily *explains* the situation under investigation.

However, in the third task, when interpolating, the modeller is strongly guided by the data that have been carefully collected and analysed, and a curve is sought that captures the trend of the data in order to *predict* in between the data points. Less reliably, we may want to predict outside the range of known data points. However, as we’ll see in the lab, curves of best fit are very good for interpolation but not nearly so good for extrapolation, so we have to be careful.

In all situations, the modeller may ultimately want to make predictions from the model. However, the modeller tends to emphasize the proposed models over the data when model fitting, whereas she places greater confidence in the collected data when interpolating and attaches less significance to the form of the model.

For example, suppose we have data from the 1918 influenza pandemic in Philadelphia shown in Figure 3.1. The x axis represents time in days, and the y axis represents the number of fatal cases. To make predictions based solely upon this data, we could use a technique such as *spline interpolation* (which we will study in the next chapter) in order to pass a smooth curve through the

points (see Figure 3.2). In this case, the interpolating curve passes through the data points and captures the trend of the behaviour over the range of observations.

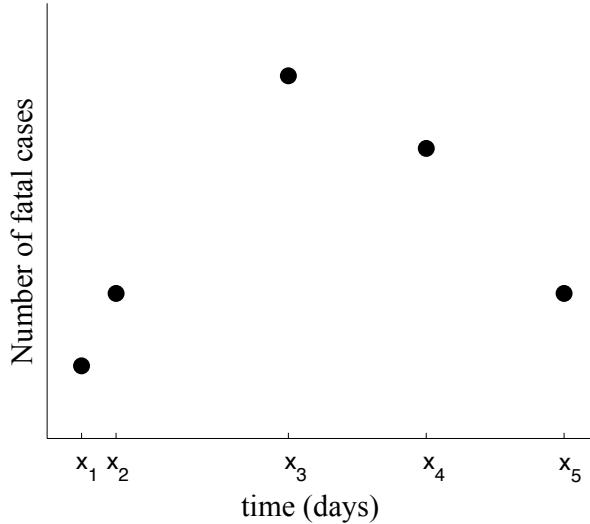


Fig. 3.1. Data from the 1918 influenza pandemic.

However, we know that data isn't always perfect, so a curve that passes precisely through every data point may actually be less useful than a curve that misses them all, but captures the “trend”.

Suppose that, in studying the data, the modeller makes assumptions leading to the expectation of a quadratic model, or parabola, of the form $y = C_1x^2 + C_2x + C_3$. In this case, the data of Figure 3.1 would be used to determine the arbitrary constants C_1 , C_2 and C_3 in order to select the “best” parabola. See Figure 3.3. The fact that the parabola may deviate from some or all of the data points would be of no concern. Outside the range of data points, the curves may vary significantly; e.g., in the vicinity of x_5 , the predictions made by the curves in Figures 3.2 and 3.3 are quite different.

Of course, we may find it necessary to both fit a model and to interpolate in the same problem. The best-fitting model of a given type may prove to be unwieldy or even impossible for subsequent analysis involving operations like integration or differentiation. In such situations, the model may have to be replaced with an interpolating curve (such as a polynomial) that is more readily differentiated or integrated.

For example, a step function used to model the sudden onset of a pandemic might be replaced by a trigonometric approximation to facilitate subsequent

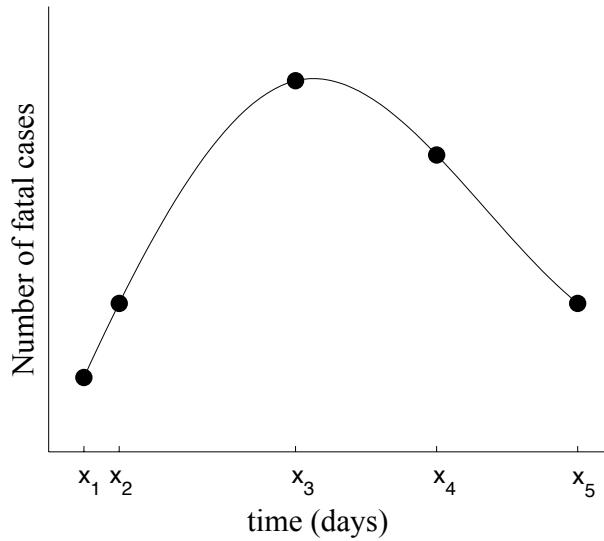


Fig. 3.2. Interpolating the data using a smooth curve.

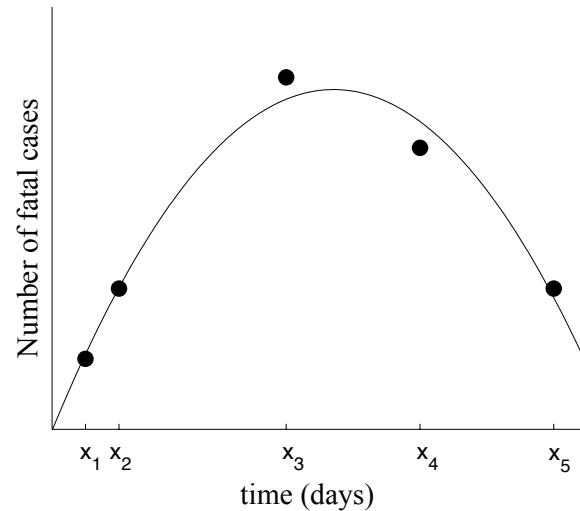


Fig. 3.3. Fitting a parabola to the data points.

analysis. In these instances, we want the interpolating curve to closely approximate the essential characteristics of the function it replaces.

3.2 Sources of error in the modelling process

When fitting models to data or interpolating, we need to examine the modelling process in order to ascertain where errors can arise. If error considerations are neglected, undue confidence may be placed in intermediate results, causing faulty decisions in subsequent steps. Our goal is to ensure that all parts of the modelling process are computationally compatible and to consider the effects of cumulative errors likely to exist from previous steps.

We can classify errors under the following category scheme:

1. Formulation error
2. Truncation error
3. Round-off error
4. Measurement error.

A *formulation error* is an error resulting from the assumption that certain variables are negligible or from simplifications arising in describing interrelationships among the variables in the various submodels. For example, if we ignore the spatial heterogeneity of a population as an influenza pandemic sweeps through, we may be neglecting important relationships among individuals that facilitate or prevent the transmission of the disease. Formulation errors are present in even the best models. As George Box famously said, “All models are wrong, but some are useful.”

Truncation errors are those errors attributable to the numerical method used to solve a mathematical problem. For example, we may find it necessary to approximate e^x with a polynomial representation obtained from the power series

$$e^x = 1 + x + \frac{x^2}{2} + \frac{x^3}{6} + \frac{x^4}{24} + \dots$$

An error will be introduced when the series is truncated (i.e., only a finite number of terms are included) to produce the polynomial. In fact, every time your calculator or computer calculates an exponential (or a sine or cosine or a logarithm), it is using just such a finite polynomial representation, complete with truncation error.

Round-off error refers to any error caused by using a finite-digit machine for computation. Since all numbers cannot be represented exactly using only finite representations, we must always expect round-off errors to be present. For example, consider a calculator or computer that uses 8-digit arithmetic. Then the number $1/3$ is represented by 0.33333333 so that $3 \times 1/3 = 0.99999999$, rather than the actual value 1. The error of 10^{-8} is due to round-off. The ideal real number $1/3$ is an *infinite* string of decimal digits 0.33333... , but any calculator or computer can do arithmetic only with numbers having finite precision. When many arithmetic operations are performed in succession, each with its own round-off, the accumulated effect of round-off can significantly alter the numbers that are supposed to be the answer. Round-off is just one of

the things we have to live with — and be aware of — when we use computing machines.

Measurement errors are caused by imprecision in the data collection. This imprecision may include such diverse things as human errors in recording or reporting the data or the actual physical limitations of the laboratory equipment. For example, considerable measurement error would be expected in the data reflecting the spread of influenza through a population, since we can't measure everyone, fatalities caused by influenza are often attributed to other factors and the speed of spread may outpace the ability to measure symptoms.

3.3 Visual fitting with the original data

Suppose we want to fit the model $y = ax + b$ to the data shown in Figure 3.4. How might the constants a and b be chosen to determine the line that “best fits” the data? Generally, when more than two data points exist, all of them cannot be expected to lie exactly along a straight line, even if such a line accurately models the relationship between the two variables x and y . That is, ordinarily there will be some vertical discrepancy (residuals) between some of the data points and any particular line being considered. We refer to these vertical discrepancies as *absolute deviations*.

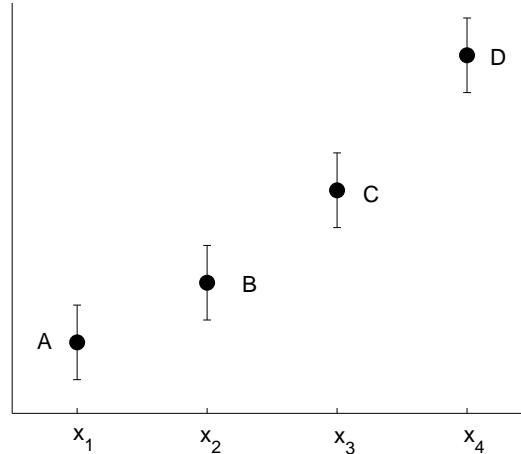


Fig. 3.4. Each data point is thought of as an interval of confidence.

For the “best fitting” line, we might want to try to minimise the sum of these absolute deviations leading to the model depicted in Figure 3.5. While

we may be successful at minimising the sum of the absolute deviations, the absolute deviation from individual points may be quite large. For example, while points A, B and C are quite close to the fitted line, point D is some considerable distance from it. If we have confidence in the accuracy of this data point, there would be concern for the predictions made from the fitted line near this point.

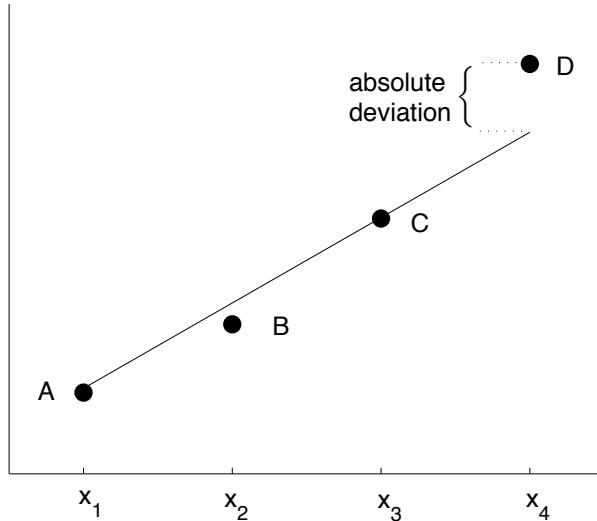


Fig. 3.5. Minimising the sum of the absolute deviations from the fitted line.

As an alternative, suppose a line is selected that minimises the largest deviation from any point. In this case, we would have the line shown in Figure 3.6. In this case, no point is exactly on the line, but no point is too far from it either.

Although these visual methods for fitting a line to data points may appear imprecise, the methods are often quite compatible with the accuracy of the modelling process itself. The grossness of the assumptions and the imprecision involved in the data collection may not warrant a more sophisticated analysis. In such situations, the blind application of more analytical methods may lead to models far less appropriate than one obtained graphically.

Furthermore, a visual inspection of the model fitted graphically to the data immediately gives an impression of how good the fit is and where it appears to fit well. Unfortunately, these important considerations are often overlooked in problems with large amounts of data analytically fitted via computer codes. Since the model-fitting portion of the modelling process appears to be more

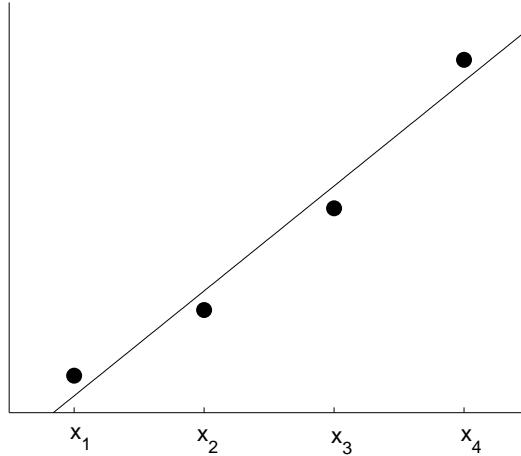


Fig. 3.6. Minimising the largest absolute deviation from the fitted line.

precise and analytic than some of the other steps, there is a tendency to place undue faith in the numerical computations.

3.4 Transforming the data

Most of us are limited visually to fitting only lines. So to graphically fit other curves as models, we have to transform the data. For example, consider the data shown in Table 3.1 of new cases of HIV infections detected in 1981. The data is plotted in Figure 3.7.

Month	Jan	Feb	Mar	Apr
New Cases	51	179	370	1207

Table 3.1. Collected data for new HIV infections in 1981.

We may suspect that the relationship is exponential; i.e., of the form $y = Ce^x$, where x is the time in months since the beginning of the survey and y is the number of cases. Thus, if we plot y versus e^x , we should obtain approximately a straight line. See Figure 3.8. Since the plotted data points do lie approximately along a line that projects through the origin, we conclude that the assumed proportionality is reasonable. From the figure, the slope of the line is approximated as $C = \frac{1207-51}{54.6-2.7} \approx 22.3$.

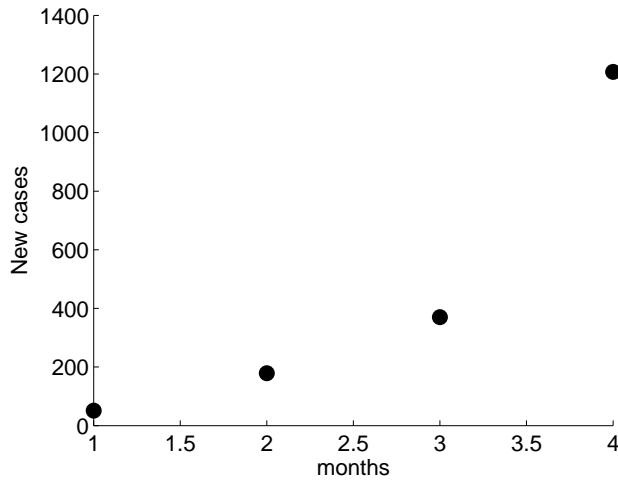


Fig. 3.7. Plot of collected data for new HIV infections in 1981.

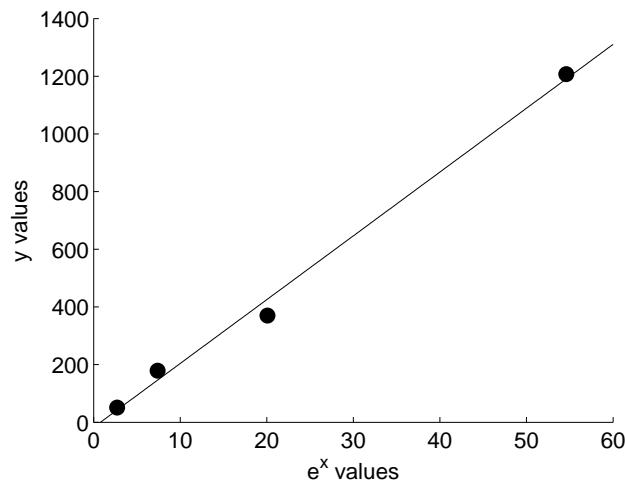


Fig. 3.8. Plot of y versus e^x for the original data.

An alternative technique involves taking the logarithm of each side of the equation $y = Ce^x$ to obtain

$$\begin{aligned}
 \ln y &= \ln[Ce^x] \\
 \ln y &= \ln C + \ln e^x && \text{(since } \ln(ab) = \ln(a) + \ln(b)\text{)} \\
 \ln y &= \ln C + x && \text{(remember } \ln \text{ and } e \text{ are inverses).}
 \end{aligned}$$

Note that this expression is an equation of a line in the variables $\ln y$ and x . The number $\ln C$ is the intercept when $x = 0$. The transformed data are shown in Table 3.2 and plotted in a “semi-log” plot, Figure 3.9.

x	1	2	3	4
$\ln y$	3.932	5.167	5.914	7.096

Table 3.2. Transformed data

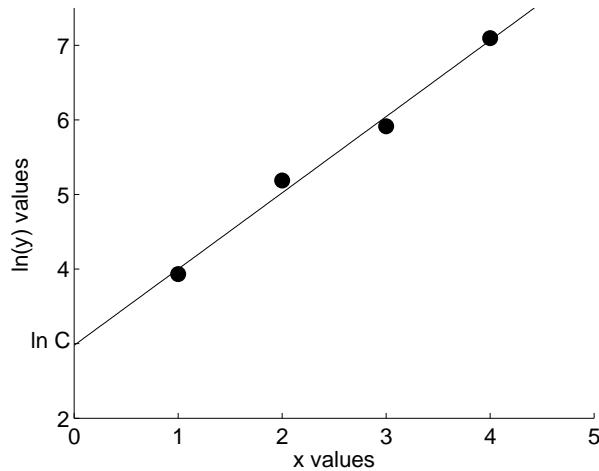


Fig. 3.9. Plot of $\ln y$ versus x for the transformed data.

From Figure 3.9, we can determine that the intercept $\ln C$ is 2.9776, giving $C = e^{2.9776} \approx 19.6$. So which C is the right one? Answer: probably neither. Which C do we have more faith in? Answer: definitely the second one! It uses more data points (using all four to determine the line of best fit and hence the intercept), whereas the first C only uses two data points to determine the slope. And we can see from Figure 3.8 that this isn’t going to be the exact slope of the fitted line anyway. Of course, a smarter approach to this would be to use linear regression to calculate the slope of the best-fit line.

So does that mean we should always transform our data into something where we can fit a line, if that’s possible? Well... not necessarily. As we’ll see in the next section, this too has the potential to mislead. It’s a tough curve-fitting world out there.

3.4.1 Regression coefficients can be misleading

A similar transformation can be performed on a variety of other curves to produce linear relationships among the resulting transformed variables. For example, if $y = x^a$, then

$$\begin{aligned}\ln y &= \ln(x^a) \\ \ln y &= a \ln x \quad (\text{remember } \ln(b^c) = c \ln(b))\end{aligned}$$

is a linear relationship in the transformed variables $\ln y$ and $\ln x$, giving us a “log-log” transformation with slope a .

For example, consider the data in Table 3.3, where x represents the number of avian influenza infections detected and y represents the number of birds that must be culled to contain the disease.

x	3	7	20	148
y	8	65	549	36300

Table 3.3. Initial infections and cull size for an avian influenza outbreak.

If you use any basic linear regression package (a simple calculator will do), you’ll find that a line of best fit could be applied, with $r = 0.9956$. Seems pretty good, right? But if we plot the data, as in Figure 3.10, you can see that the fit isn’t nearly so good as we might suspect. (Indeed, the inset shows just how far off the line is for the first three data points, due to the scaling!) This is one of the dangers of relying on computers to do the work for us: we can easily be fooled into believing results that are misleading or outright wrong.

If instead we suspect a relationship of the form $y = x^a$ and plot $\ln y$ versus $\ln x$, as in Figure 3.11, then we find $r = 0.9995$ and the slope of the line of best fit is 2.1496, suggesting that $a \approx 2.1$.

Finally, we can check this curve against the original data, by plotting the curve $y = x^{2.1}$ against the original data, as in Figure 3.12. This looks a lot better than Figure 3.10. It’s at this point, by fitting the curve to the *original* data, that we make our decisions about which curve is the best.

Note that the linear regression r values weren’t terribly helpful here. Be careful: there’s a tendency for scientists to put all their faith in the r value they’ve calculated (even for very small data sets like the ones we used) and many are willing to accept r values a lot lower than the ones seen here.

3.4.2 Transformations can also be misleading

At this point, we must make an important observation. Suppose we do invoke a transformation and plot $\ln y$ versus x , as in Figure 3.12 and find the line

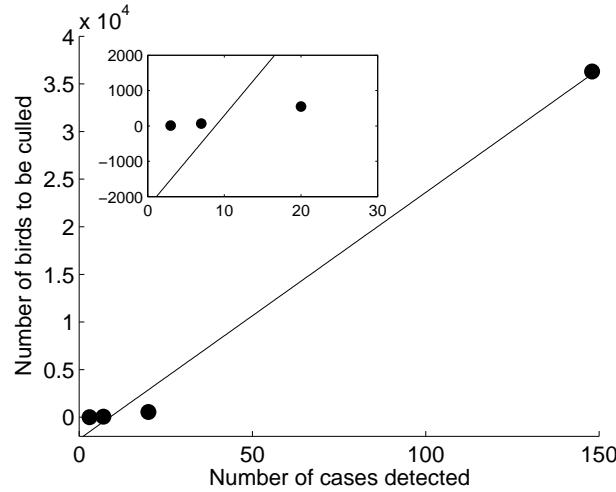


Fig. 3.10. Plot of y versus x and a line of best fit for the avian influenza data. Inset: a rescaling around the first three data points.

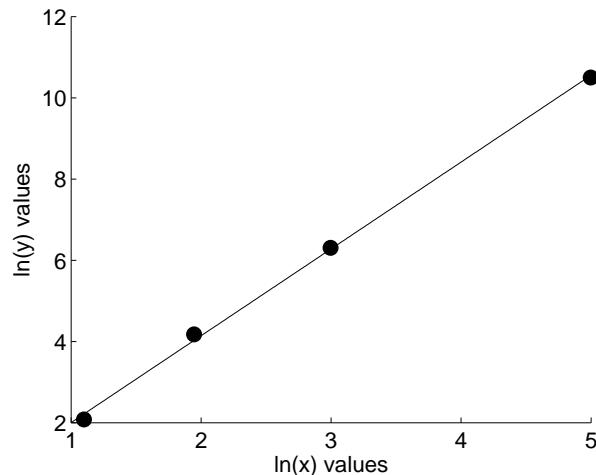


Fig. 3.11. Plot of $\ln y$ versus $\ln x$ and a line of best fit for the avian influenza data.

that successfully minimises the sum of the absolute deviations of the transformed data points. The line then determines $\ln C$, which in turn produces the proportionality constant C . Although it is not obvious, the resulting model $y = Ce^x$ is *not* the member of the family of exponential curves of the form

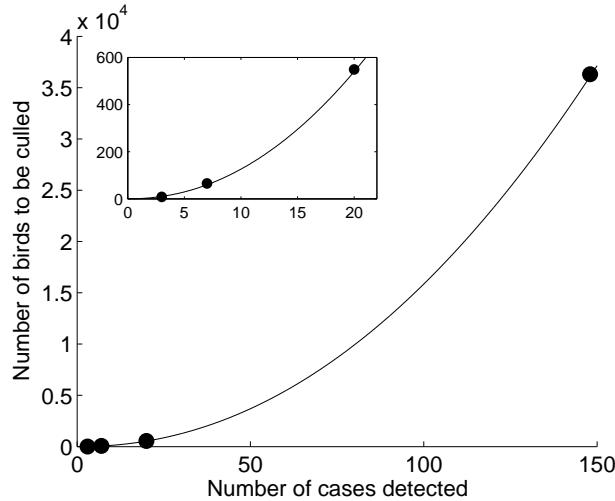


Fig. 3.12. Comparing the exponential curve with the original avian influenza data. Inset: a rescaling of the first three data points.

ke^x that minimises the sum of the absolute deviations from the original data points (when we plot y versus x).

That is, the line may be the best fit in the transformed data, but it doesn't follow that the corresponding curve is necessarily the best fit in the original.

When transformations of the form $y = \ln x$ are made, the distance concept is distorted. While a fit that is compatible with the inherent limitations of a graphical analysis may be obtained, we must be aware of this distortion and, crucially, verify the model using the graph from which it is intended to make predictions: namely the y versus x graph in the original data, rather than the graph of the transformed variables.

For example, consider the data plotted in Figure 3.13, which might represent biannual outbreaks of influenza. Suppose we have reason to believe the data are expected to fit a model of the form $y = Ce^{1/x}$. We want to choose the “best” C that fits this. Using a logarithmic transformation as before, we find

$$\begin{aligned}\ln y &= \ln(Ce^{1/x}) \\ \ln y &= \ln C + \ln(e^{1/x}) \quad (\text{since } \ln(ab) = \ln a + \ln b) \\ \ln y &= \frac{1}{x} + \ln C \quad (\text{since } \ln \text{ and } e \text{ are inverses}).\end{aligned}$$

A plot of the points $\ln y$ versus $1/x$ based on the original data is shown in Figure 3.14. Note from the figure how the transformation distorts the distances between the original data points and squeezes them all together. Consequently,

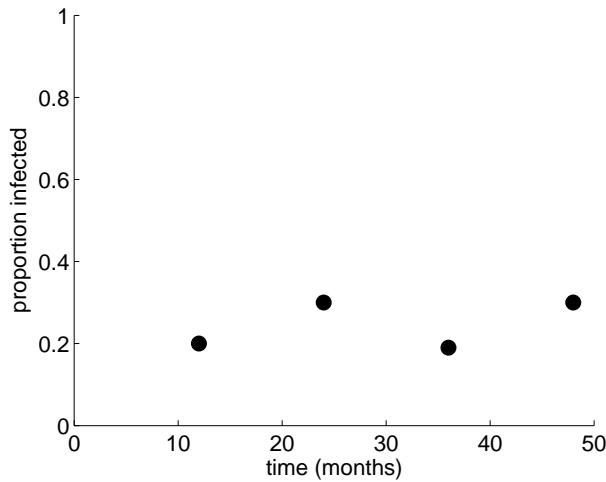


Fig. 3.13. Biannual outbreaks of influenza.

if a straight line is made to fit the transformed data plotted in Figure 3.14, the absolute deviations appear relatively small (that is, small on the Figure 3.14 scale rather than on the Figure 3.13 scale).

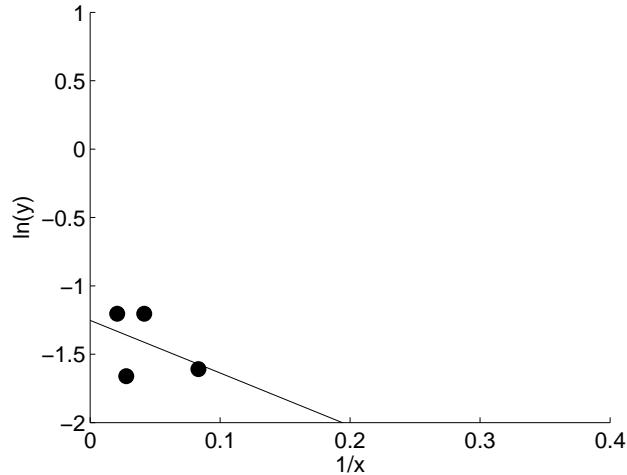


Fig. 3.14. The transformed data points and a line of best fit.

This means that the model is a reasonably good fit for the transformed data, with $\ln C \approx -1.25$. It might not look so great to our eye, but the deviations are quite small, so a computer would tell us it's a very good fit. We could then solve for C and assume we've got the best C (because we had the best $\ln C$).

But let's test this out on the original data. If we plot the fitted model $y = Ce^{1/x}$ to the data in Figure 3.13, you would see that it fits the data relatively poorly, as shown in Figure 3.15. There are no biannual peaks to the fitted line, and, worse, the line would behave quite badly (heading up to infinity) in the vicinity of 0.

What's gone wrong here? Answer: The data were never supposed to fit a model of the form $y = Ce^{1/x}$. But we wouldn't know that from the transformation, which tells us that the fit is actually pretty good.

From this example, we can see that if we are not careful when using transformations, we can be tricked into selecting a relatively poor model. This realisation becomes especially important when comparing alternative models. Very serious errors can be introduced when selecting the best model unless all the comparisons are made with the original data. Otherwise, the choice of "best" model may be determined by a peculiarity of the transformation rather than on the merits of the model and how well it fits the original data. While the danger of making transformations is evident in this example, it is easy to be fooled if we are not especially observant, since many computer codes fit models by first making a transformation.

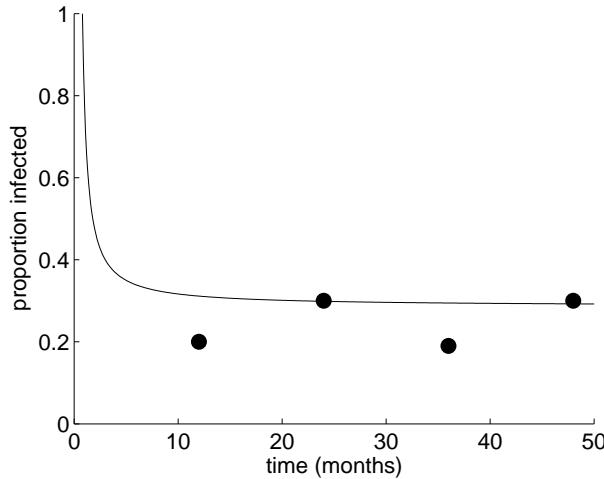


Fig. 3.15. A plot of the curve $y = Ce^{1/x}$ based on the value $\ln C \approx -1.25$.

3.5 Lab work

The problem

In the data in Table 3.4, x is the diameter of a ponderosa pine in inches measured at chest height and y is a measure of volume: the number of board feet divided by 10. (Most lumber is sold by the board foot, which is equal to a board that is one foot long, one foot wide and one inch thick.)

x	17	19	20	22	23	25	28	31	32	33	36	37	38	39	41
y	19	25	32	51	57	71	113	141	123	187	192	205	252	259	294

Table 3.4. Diameter of a ponderosa pine versus volume.

- Use the `polyfit` command to fit polynomials of degree 1, 2, 3, 4, 6, 9 and 15 to the data. Which of these seem reasonable?
- Use the Basic Data fitting tool to fit same polynomials (except for 15) to the data. Which seems reasonable now?
- Use Matlab to plot the points and plot the natural cubic spline joining them.
- Make an appropriate transformation to fit the model $y = ax^b$.
- Estimate the parameters a and b of the model.
- Plot your model against the original data. Which of the earlier polynomials does this most closely approximate?

The solution

First we need to enter the data. We'll then use `polyfit` to evaluate a range of possible best-fit polynomials.

```

x=[17 19 20 22 23 25 28 31 32 33 36 37 38 39 41];
y=[19 25 32 51 57 71 113 141 123 187 192 205 252 259 294];
n=input('Enter the degree of the polynomial ');
p=polyfit(x,y,n);
f=polyval(p,x);
figure(1)
plot(x,y,'o') %This plots a 'clean' set of data for later use
figure(2)
plot(x,y,'o',x,f,'-')

```

Matlab is great for fitting data. If you keep Figure 1 open, then choose “Tools → Basic Fitting”, you have the option of fitting all manner of things to the data. Try fitting linear, quadratic, cubic etc polynomials to the data.

How does this look? (You may want to type `axis([15 45 0 350])` into the Matlab command window in order to rescale the axes.)

Fitting low-order polynomials isn't too bad, but higher-order polynomials don't seem to work very well. They're a reasonable fit to the *actual* data, but outside the region we have data for, they're an extremely bad fit.

While we're here, let's choose the "spline interpolant" option (again, you may want to rescale the axes). This fits a cubic spline to the data. We'll learn more about the theory of cubic splines in the next chapter, but basically they're a way to fit a smooth curve to the data that actually passes through all the data points, should we so desire it.

To transform the data, since we have exponents, the obvious transformation is to take the logarithm of both sides. Thus

$$\begin{aligned}\ln y &= \ln(ax^b) \\ \ln y &= \ln a + \ln x^b \\ \ln y &= \ln a + b \ln x.\end{aligned}$$

Thus, if we plot $\ln y$ versus $\ln x$, then the intercept should be $\ln a$ and the slope should be b .

In Matlab, the function $\ln x$ is represented as `log(x)` (for historical reasons, "log" and "ln" are the same thing in mathematics, although confusingly they aren't for most scientists), so you can type `plot(log(x),log(y), 'o')` directly into the Matlab command window. The '`o`' means "plot the data as circles". You could also use '`*`', '`+`', etc, but Matlab's default is to join the data with lines, which we don't want.

To apply a line of best fit, simply choose "Tools → Basic Fitting" and choose "linear". Click the  button and you'll see the coefficients of the linear polynomial. Thus,

$$\begin{aligned}\ln a &= -5.7427 \\ \Rightarrow \quad a &= e^{-5.7427} = 0.0032061 \\ b &= 3.0919.\end{aligned}$$

We can thus plot our model versus the data by typing the following code into the command window:

```
plot(x,y, 'o', x, 0.0032061.*x.^ (3.0919))
```

Matlab can plot multiple things on top of each other. Every two entries (plus an optional descriptor) is a data set.

We can also see that the cubic polynomial was the closest to the best fit, assuming this model is accurate.

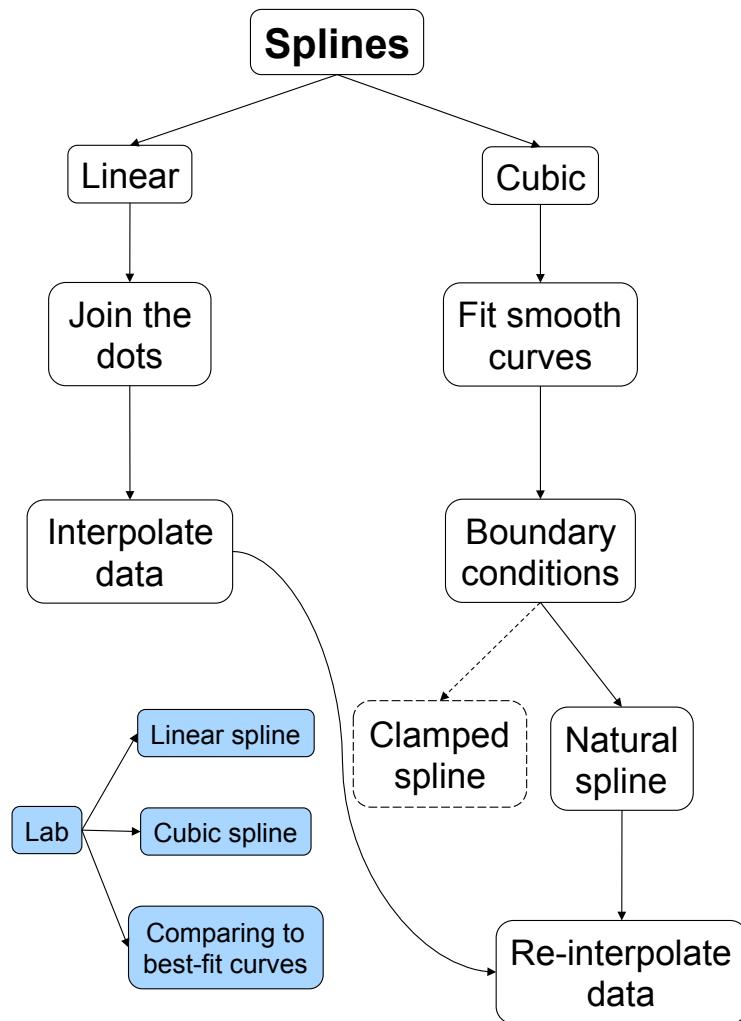
3.5.1 Exercises

Table 3.5 shows the mean number of *T. infestans* bugs collected per infected site, after community-wide insecticide spraying in 1992 to eliminate Chagas' disease in two rural areas in Argentina.

Month	Amamá	Mercedes/Trinidad
Oct-93	2.7	1
Nov-94	1	5
May-95	2	1
Nov 95	2.8	4.8
May-96	3.4	3.7
Nov-96	3.3	2.8
May-97	8.7	6.2
Nov-97	7.4	11.2
May-98	6.4	5.4
Nov-98	8.4	3.4
May-99	7.9	4.7

Table 3.5. Mean number of bugs per infected site. (Data courtesy Ricardo Görtler.)

1. Plot the data for each site on two separate figures. (Watch out for the scaling on the x -axis; the October entry means you'll have to be careful.)
2. Fit some polynomials to each data set by modifying your polyfit program. Which polynomial seems the “best” to you?
3. Use the “Basic Fitting” tools to fit your polynomials again. What does the behaviour around the edges tell you?
4. Fit a cubic spline to the data. Do you think this would be a better fit than a polynomial?
5. Do you have a reason to suspect the data might need to be transformed? If so, make an appropriate transformation and determine the parameters.
6. Is there a qualitative difference between the mean number of bugs per infested site in the two villages?
7. What would you predict would be the mean number of bugs per infected site in Nov 99?



Splines

In the previous chapter, we saw that an overreliance on what the computer tells us can lead to false confidence in our choice of models. Here, we explore in more detail the precise manner in which the computer fits splines to data. We use splines as an illustrative example, because they are a type of model fitting that assumes the data is excellent and it is the duty of the model to adapt to this data.

By the end of this chapter, you should have some understanding of how the computer uses splines. It should be emphasised that this chapter focusses on the theory behind splines, but in practice the computer will do all the work. However, it's vital to understand the theory behind what the computer is doing in order to use it.

4.1 Spline interpolation

We want to construct models that capture the trend of the data, and one very good way to do this is to use polynomials. Polynomials are especially appealing because they're so easy to integrate and differentiate. This allows us to do further mathematical analysis (which we won't do here) that can't be done if the curves have sharp corners or discontinuities in the slopes.

However, high-order polynomials tend to oscillate near the endpoints of the data interval, and the coefficients can be quite sensitive to small changes in the data. Smoothing with a low-order polynomial lessens these effects. However, unless the data are essentially quadratic or cubic in nature, smoothing with a low-order polynomial may yield a relatively poor fit somewhere over the range of the data.

Instead we can use a technique called *spline interpolation* to fit different polynomials between successive pairs of data points. This allows us to capture the trend of the data, regardless of the nature of the underlying relationship, while simultaneously reducing the tendency toward oscillation and the sensitivity to changes in the data.

We'll look at two applications: linear and cubic splines. The former because it's so straightforward (and shows up in unexpected places) and the latter because that's the usual spline used. A cubic has enough freedom to match first and second derivatives but is a low-order polynomial, which won't give unnecessarily oscillations.

4.2 Linear splines

The simplest method of connecting data points would simply be to "join the dots", as in Figure 4.1. When x is in the interval $x_1 \leq x < x_2$, the model that is used is the *linear spline* $S_1(x)$ passing through the data points (x_1, y_1) and (x_2, y_2) :

$$S_1(x) = a_1 + b_1 x \quad \text{for } x_1 \leq x < x_2. \quad (4.1)$$

Similarly, when $x_2 \leq x \leq x_3$, the linear spline $S_2(x)$ passing through (x_2, y_2) and (x_3, y_3) is

$$S_2(x) = a_2 + b_2 x \quad \text{for } x_2 \leq x \leq x_3. \quad (4.2)$$

Note that both spline segments meet at the point (x_2, y_2) .

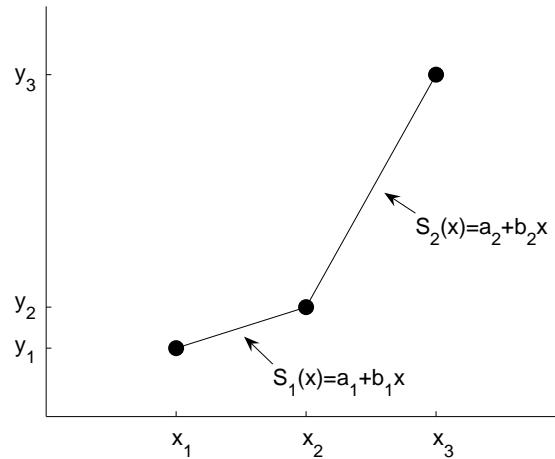


Fig. 4.1. A linear spline model is a continuous function consisting of line segments.

Methicillin-resistant *staphylococcus aureus* (MRSA) was detected in the Los Angeles jails at low levels for a number of years before suddenly increasingly sharply in the following year. The data is shown in Table 4.1, where x_i is the time (in years) and $y(x_i)$ is the corresponding number of cases detected.

x_i	1	2	3
$y(x_i)$	5	8	25

Table 4.1. Data for MRSA, which increases sharply one year.

The spline $S_1(x)$ must pass through the points $(1, 5)$ and $(2, 8)$. Substituting into (4.1)–(4.2) gives us

$$\begin{aligned} a_1 + 1b_1 &= 5 \\ a_1 + 2b_1 &= 8. \end{aligned}$$

We could write this in matrix form

$$\begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} a_1 \\ b_1 \end{pmatrix} = \begin{pmatrix} 5 \\ 8 \end{pmatrix}$$

and use Matlab to solve... except that we're not going to quite yet.

Similarly, the spline $S_2(x)$ must pass through the points $(2, 8)$ and $(3, 25)$ yielding

$$\begin{aligned} a_2 + 2b_2 &= 8 \\ a_2 + 3b_2 &= 25. \end{aligned}$$

Again, we could write this in matrix form

$$\begin{pmatrix} 1 & 2 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} a_2 \\ b_2 \end{pmatrix} = \begin{pmatrix} 8 \\ 25 \end{pmatrix}$$

and use Matlab to solve... but why solve two equations when we could be smarter and just solve one?

Let's write *both* systems as one big matrix equation:

$$\begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 2 & 0 & 0 \\ 0 & 0 & 1 & 2 \\ 0 & 0 & 1 & 3 \end{pmatrix} \begin{pmatrix} a_1 \\ b_1 \\ a_2 \\ b_2 \end{pmatrix} = \begin{pmatrix} 5 \\ 8 \\ 8 \\ 25 \end{pmatrix}.$$

Matlab is excellent for solving matrix equations. Indeed, that's probably the thing it's "best" at, if such a term can be applied to a computer program. Remember that if we have the linear system $Ac = b$ (where A is a matrix and b and c are vectors), then — so long as $\det(A) \neq 0$, but Matlab will give you an error anyway if it isn't — the solution is $c = A^{-1}b$. The following code will do all the work for us:

```
A=[1 1 0 0;
 1 2 0 0;
```

```

0 0 1 2;
0 0 1 3];
b=[5;8;8;25];
c=A^(-1)*b

```

The solution of these linear systems of equations yields $a_1 = 2$, $b_1 = 3$, $a_2 = -26$ and $b_2 = 17$. Thus, the linear spline model for the data of Table 4.1 is summarised in Table 4.2.

Interval	Spline model
$1 \leq x \leq 2$	$S_1(x) = 2 + 3x$
$2 < x \leq 3$	$S_2(x) = -26 + 17x$

Table 4.2. A linear spline model.

To illustrate how the linear spline model is used, let's predict $y(1.67)$ and $y(2.33)$. Since $1 \leq 1.67 < 2$, we need to use $S_1(x)$ for this x value. Hence

$$\begin{aligned} S_1(1.67) &= 2 + 3(1.67) \\ &\approx 7.01. \end{aligned}$$

Likewise, since $2 \leq 2.33 \leq 3$, we use $S_2(x)$ for this x value. Thus

$$\begin{aligned} S_2(2.33) &= -26 + 17(2.33) \\ &\approx 13.61. \end{aligned}$$

While the linear spline method is sufficient for many applications, it fails to capture the trend of the data. Furthermore, if you examine Figure 4.1, you see that the linear spline model does not appear “smooth”. That is, in the interval $1 \leq x < 2$, $S_1(x)$ has constant slope 3, whereas in the interval $2 \leq x \leq 3$, $S_2(x)$ has constant slope 17. At $x = 2$, there is an abrupt change in the slope of the model from 3 to 17 so that the first derivatives $S'_1(x)$ and $S'_2(x)$ fail to agree at $x = 2$.

4.3 Cubic splines

Consider now Figure 4.2. In a manner analogous to linear splines, we define a separate spline function for the intervals $x_1 \leq x < x_2$ and $x_2 \leq x < x_3$ as follows:

$$\begin{aligned} S_1(x) &= a_1 + b_1x + c_1x^2 + d_1x^3 && \text{for } x_1 \leq x < x_2 \\ S_2(x) &= a_2 + b_2x + c_2x^2 + d_2x^3 && \text{for } x_2 \leq x \leq x_3. \end{aligned}$$

Since we will want to refer to the first and second derivatives, let's find them as well:

$$\begin{aligned} S'_1(x) &= b_1 + 2c_1x + 3d_1x^2 & \text{for } x_1 \leq x < x_2 \\ S''_1(x) &= 2c_1 + 6d_1x & \text{for } x_1 \leq x < x_2 \\ S'_2(x) &= b_2 + 2c_2x + 3d_2x^2 & \text{for } x_2 \leq x \leq x_3 \\ S''_2(x) &= 2c_2 + 6d_2x & \text{for } x_2 \leq x \leq x_3. \end{aligned}$$

The model is presented geometrically in Figure 4.2.

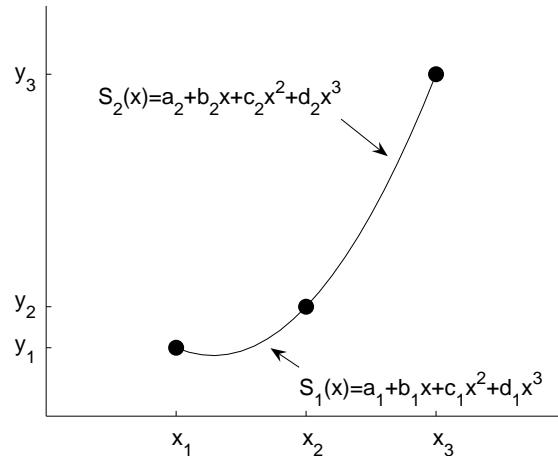


Fig. 4.2. A cubic spline model is a continuous function with continuous first and second derivatives consisting of cubic polynomial segments.

Cubic splines offer the possibility of matching up not only slopes, but also the curvatures at each interior data point. To determine the constants defining each cubic spline segment, we appeal to the requirement that each spline pass through the two data points specified by the interval over which the spline is defined. For the spline model depicted in Figure 4.2, this requirement yields the equations

$$y_1 = S_1(x_1) = a_1 + b_1x_1 + c_1x_1^2 + d_1x_1^3 \quad (4.3)$$

$$y_2 = S_1(x_2) = a_1 + b_1x_2 + c_1x_2^2 + d_1x_2^3 \quad (4.4)$$

$$y_2 = S_2(x_2) = a_2 + b_2x_2 + c_2x_2^2 + d_2x_2^3 \quad (4.5)$$

$$y_3 = S_2(x_3) = a_2 + b_2x_3 + c_2x_3^2 + d_2x_3^3. \quad (4.6)$$

Note that there are eight unknowns $(a_1, b_1, c_1, d_1, a_2, b_2, c_2, d_2)$ and only four equations. In order to determine the constants, we need the same number of equations as unknowns, so we need four more equations.

Since smoothness of the spline system is also required, adjacent first derivatives must match at each interior data point (in our case when $x = x_2$). We thus require

$$\begin{aligned} S'_1(x_2) &= S'_2(x_2) \\ b_1 + 2c_1x_2 + 3d_1x_2^2 &= b_2 + 2c_2x_2 + 3d_2x_2^2. \end{aligned} \quad (4.7)$$

We can also require that adjacent second derivatives match at each interior data point:

$$\begin{aligned} S''_1(x_2) &= S''_2(x_2) \\ 2c_1 + 6d_1x_2 &= 2c_2 + 6d_2x_2. \end{aligned} \quad (4.8)$$

We still need two additional independent equations. While conditions on the derivatives at interior data points have been applied, nothing has been said about the derivative at the exterior endpoints (x_1 and x_3 in our case).

If there were more points to be joined, we'd use information on the other splines to construct further equations for x_1 and x_3 . However, in this case, these are our endpoints, so we need to decide what to do about nailing them down.

A condition we'll impose, called a *natural spline*, is to require that there be no change in the first derivative at the exterior endpoints¹. Mathematically, since the first derivative is constant, the second derivative must then be zero (because the derivative of any constant is zero). Applying this condition at x_1 and x_3 yields

$$S''_1(x_1) = 2c_1 + 6d_1x_1 = 0 \quad (4.9)$$

$$S''_2(x_3) = 2c_2 + 6d_2x_3 = 0. \quad (4.10)$$

We now have eight equations (4.3)–(4.10) and eight unknowns, so this system is solvable.

4.3.1 An example

Let's construct the natural cubic spline model using the data from Table 4.1. The procedure readily extends to problems with more data points. (Though fortunately we can let the computer do the work for us in such cases.)

Requiring the spline segment $S_1(x)$ to pass through the two endpoints $(1, 5)$ and $(2, 8)$ of its interval requires that $S_1(1) = 5$ and $S_1(2) = 8$. Applying (4.3) and (4.4), we have

¹ Other conditions could be imposed here, if we had more information about the values of the first derivatives at the exterior endpoints, for example. This sort of cubic spline is called a *clamped spline*.

$$\begin{aligned} a_1 + 1b_1 + 1^2c_1 + 1^3d_1 &= 5 \\ a_1 + 2b_1 + 2^2c_1 + 2^3d_1 &= 8. \end{aligned}$$

Similarly, $S_2(x)$ must pass through the endpoints of the second interval so that $S_2(2) = 8$ and $S_2(3) = 25$. Applying (4.5) and (4.6), we have

$$\begin{aligned} a_2 + 2b_2 + 2^2c_2 + 2^3d_2 &= 8 \\ a_2 + 3b_2 + 3^2c_2 + 3^3d_2 &= 25. \end{aligned}$$

Next, the first derivatives of $S_1(x)$ and $S_2(x)$ are forced to match at the interior data point $x_2 = 2$, ie $S'_1(2) = S'_2(2)$. Applying (4.7), we have

$$b_1 + 2c_1(2) + 3d_1(2)^2 = b_2 + 2c_2(2) + 3d_2(2)^2.$$

Forcing the second derivatives of $S_1(x)$ and $S_2(x)$ to match at $x_2 = 2$ requires $S''_1(2) = S''_2(2)$. Applying (4.8), we have

$$2c_1 + 6d_1(2) = 2c_2 + 6d_2(2).$$

Finally, a natural spline is built by requiring that the second derivatives at the two endpoints be zero; ie., $S''_1(1) = S''_2(3) = 0$. Applying (4.9) and (4.10), we have

$$\begin{aligned} 2c_1 + 6d_1(1) &= 0 \\ 2c_2 + 6d_2(3) &= 0. \end{aligned}$$

Thus the procedure has yielded a linear algebraic system of eight equations in eight unknowns that can be solved uniquely. We could do it by hand, with a (fair) bit of linear algebra or we could ask Matlab to do it for us. If we summarize the equations in the form

$$\begin{aligned} a_1 + b_1 + c_1 + d_1 &= 5 \\ a_1 + 2b_1 + 4c_1 + 8d_1 &= 8 \\ a_2 + 2b_2 + 4c_2 + 8d_2 &= 8 \\ a_2 + 3b_2 + 9c_2 + 27d_2 &= 25 \\ b_1 + 4c_1 + 12d_1 - b_2 - 4c_2 - 12d_2 &= 0 \\ 2c_1 + 12d_1 - 2c_2 - 12d_2 &= 0 \\ 2c_1 + 6d_1 &= 0 \\ 2c_2 + 18d_2 &= 0, \end{aligned}$$

then what we really have is the matrix system

$$\begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 2 & 4 & 8 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 2 & 4 & 8 \\ 0 & 0 & 0 & 0 & 1 & 3 & 9 & 27 \\ 0 & 1 & 4 & 12 & 0 & -1 & -4 & -12 \\ 0 & 0 & 2 & 12 & 0 & 0 & -2 & -12 \\ 0 & 0 & 2 & 6 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & 18 \end{bmatrix} \begin{bmatrix} a_1 \\ b_1 \\ c_1 \\ d_1 \\ a_2 \\ b_2 \\ c_2 \\ d_2 \end{bmatrix} = \begin{bmatrix} 5 \\ 8 \\ 8 \\ 25 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

We can calculate the solution in Matlab using the following code.

```
A=[1 1 1 1 0 0 0 0;
 1 2 4 8 0 0 0 0;
 0 0 0 0 1 2 4 8;
 0 0 0 0 1 3 9 27;
 0 1 4 12 0 -1 -4 -12;
 0 0 2 12 0 0 -2 -12;
 0 0 2 6 0 0 0 0;
 0 0 0 0 0 0 2 18];
b=[5;8;8;25;0;0;0;0];
c=A^(-1)*b
```

The results are summarized in Table 4.3 and illustrated in Figure 4.3.

Interval	Spline model
$1 \leq x \leq 2$	$S_1(x) = 2 + 10x - 10.5x^2 + 3.5x^3$
$2 < x \leq 3$	$S_2(x) = 58 - 74x + 31.5x^2 - 3.5x^3$

Table 4.3. A natural cubic spline model.

Let's illustrate the use of the model by again predicting $y(1.67)$ and $y(2.33)$:

$$\begin{aligned} S_1(1.67) &= 2 + 10(1.67) - 10.5(1.67)^2 + 3.5(1.67)^3 \\ &\approx 5.72 \\ S_2(2.33) &= 58 - 74(2.33) + 31.5(2.33)^2 - 3.5(2.33)^3 \\ &\approx 12.32. \end{aligned}$$

Compare these values with the values predicted by the linear spline. In which prediction values do you have the most confidence?

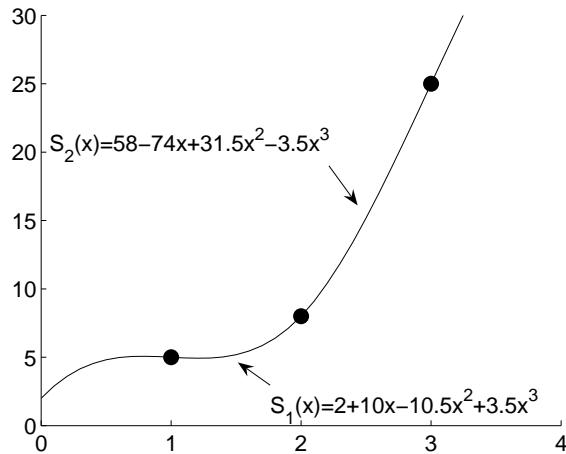


Fig. 4.3. The natural cubic spline model is a smooth curve that is easily integrated and differentiated.

4.3.2 Generalising

The construction of cubic splines for more data points proceeds in the same manner. That is, each spline is forced to pass through the endpoints of the interval over which it is defined, the first and second derivatives of adjacent splines are forced to match at the interior data points, and the natural (or possibly clamped) condition is applied at the two exterior data points.

Of course, for computational reasons, it would be necessary to implement the procedure on a computer. The procedure described here does not give rise to a computationally or numerically efficient computer algorithm. Matlab has a built-in spline interpolant, as we saw in the lab. So fitting a cubic spline to data is actually easy in Matlab if we use the “Tools → Basic Fitting” command once we’ve plotted our data. Hopefully, this section facilitates your understanding of the basic concepts underlying cubic spline interpolation, which is important to understand when using the computer tools.

4.4 Lab work

4.4.1 Exercises

1. Consider the data set

x	0	5	10
y	1	2	7

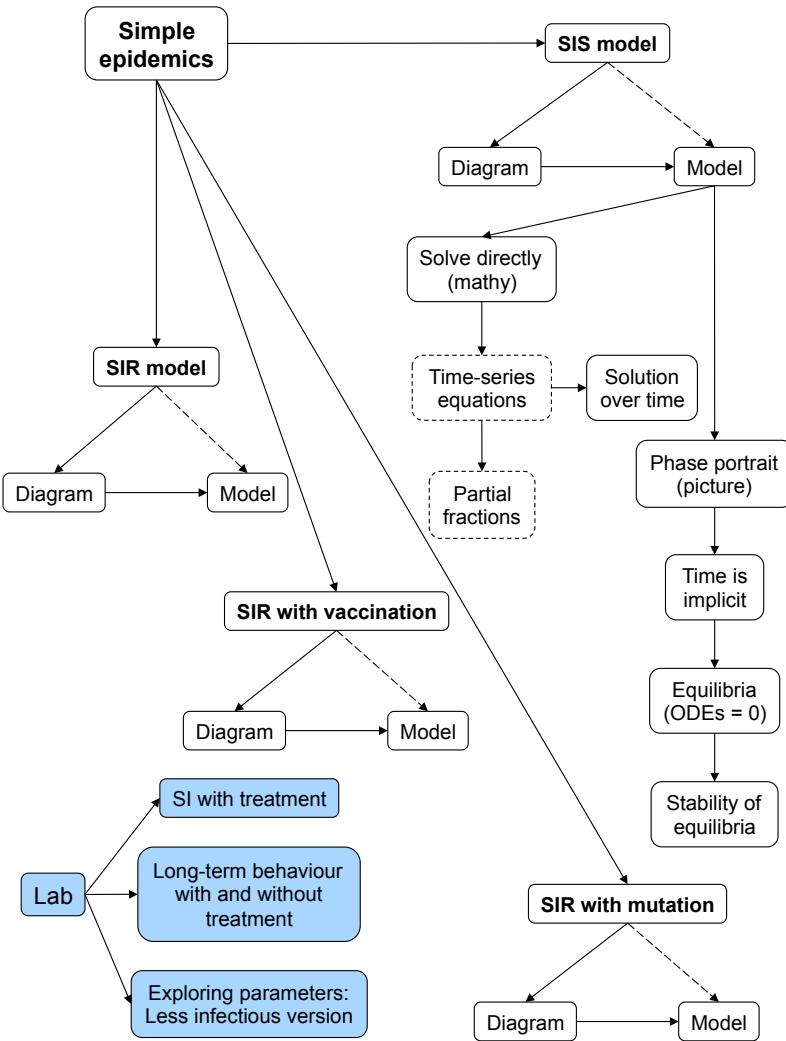
- a) Plot the data in Matlab using `plot(x,y, 'o')`.
- b) Calculate the linear spline coefficients for these data. (Don't you just wish Matlab had a linear spline tool?)
- c) Plot the linear spline on the same graph.
- d) On a separate figure, plot the data using `plot(x,y)`.
- e) What's going on here?
- f) Now fit a cubic spline to the data. In which spline do you have more faith?

2. Consider the data set

x	7	55	401	2985
y	1.3	2	2.7	4

How can we determine the curve of best fit for these data?

- a) Plot the data as circles.
- b) Apply a line of best fit to the data and calculate the regression coefficient r . (You can get Matlab to do this for you using `corrcoef`, or you can enter the data on a calculator.) Does this seem like a good fit?
- c) Plot your line of best fit on the same graph as the data. How good is the fit now?
- d) Now add an extra point (22026,5.3). How does this change your line of best fit?
- e) Make an appropriate transformation in order to fit curves of the form $y = C \ln x$ and $y = x^a$. Plot the lines of best fit along with the transformed data.
- f) Find the best C and the best a in each transformation (using all five points).
- g) Calculate r in each case. Is there a clear "winner"?
- h) Plot each curve against the original data. Is there a clear winner now?
- i) Now fit a cubic spline to the original four data points and also the revised five data points (on separate graphs). How much faith do you have in the cubic spline now?



Simple epidemic models

This chapter deals with the basics of constructing a model. We construct Ordinary Differential Equation (ODE) models for classic disease models and show how to alter these models to include other factors, such as vaccination. We also demonstrate how the “diagram” of the model is often as useful as the equations and how to get from one to the other.

By the end of this chapter, you should understand how to “read” a model, how to draw the diagram of a model and how the ODEs and the diagram are related.

5.1 SIS epidemic

If your town is suffering an epidemic, the chance of your catching the disease is proportional to the probability that you will meet a carrier of the disease. That is, the probability of catching the disease is proportional to the probability that you will be in a given place at a given time multiplied by the probability that a carrier will also be in that given place at a given time. If we represent the probability of an encounter simply by the sheer number of susceptible individuals (S) times the number of infected individuals (I), then we can represent this mathematically as

$$\frac{dS}{dt} \propto -SI \quad (5.1)$$

$$\frac{dI}{dt} \propto SI, \quad (5.2)$$

where I increases due to the encounters (more people get sick) and S decreases by the same amount (the susceptible people have now become infected).

Of course, these aren’t really probabilities; for one thing, they’re always larger than 1. But we can scale them out by replacing the proportionality with a constant multiple. That is, we replace the “ \propto ” with an “ $= \beta$ ”, where

β is some parameter that reflects how easy the disease is to transmit and the average person's ability to resist the disease. Each equation must use the same β , since βSI represents the same increase in infected people as the loss of susceptible people.

What's fundamentally happening here? Answer: people start off susceptible and then they leave the susceptible class and move into the infected class. We might represent this visually by Figure 5.1.

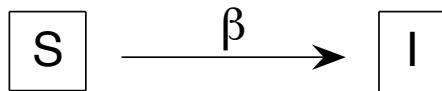


Fig. 5.1. An SI model.

This is known as an SI model, partly because it only has Susceptible and Infected classes but also because individuals move from S to I, and that's it. The corresponding equations are

$$\begin{aligned}\frac{dS}{dt} &= -\beta SI \\ \frac{dI}{dt} &= \beta SI.\end{aligned}$$

We have thus constructed a system of Ordinary Differential Equations (ODEs) to represent the overall system. Although we adapted these equations from equations (5.1)–(5.2), can you see how we might get these equations from Figure 5.1?

Of course, there's no birth or death rates, no recovery etc. Eventually, everyone gets infected, and that's the end of the story. In a moment, we'll add in other variations and conditions (such as death or vaccination), but for now at least try to make sure you can see how the diagram and the equations are linked.

One thing to note is that there's no infection if either $S = 0$ or $I = 0$ (or both, of course). That is, if there's no one infected, then no one can do any infecting, so there won't be any transmission of the disease. Likewise, if there are no susceptibles, then there's no one to get infected, so no transmission occurs here either. Another thing to note is that this isn't the only possible formulation of the infection term, although it is the simplest; we'll deal with others in Chapter 12.

This kind of transmission is called *mass-action transmission*. It assumes that populations are well-mixed: that is, every infected individual had an

equally likely chance of contacting every susceptible individual. So we are ignoring spatial heterogeneity in the population. Such an assumption is probably never true in reality, but it's more true in some circumstances than others. An airborne disease like influenza is well-approximated by mass-action transmission, so long as the population isn't too large. Conversely, an STD like HIV isn't suited to this kind of assumption at all (unless you're at an orgy, I suppose), so we'd have to think more carefully about how to deal with this kind of transmission.

ODEs generally deal with compartments, with some kind of flow into the compartment and some kind of flow out (e.g., lakes joined by rivers). These compartments need not be physical entities. We generally keep track of certain objects, keeping track of how much is going in and how much is going out. ODEs represent the way in which things change. We can model multiple compartments quite easily as a chain of compartments (e.g., towns joined by freeways) whose outflow is the inflow of another compartment.

The next epidemic we'll examine is the SIS epidemic, which stands for Susceptible \rightarrow Infected \rightarrow Susceptible. These epidemics represent diseases such as the schistosomiasis or brucellosis, which generally cause a person to be sick for a time and then recover but without immunity to the disease. Graphically, this can be represented in Figure 5.2.

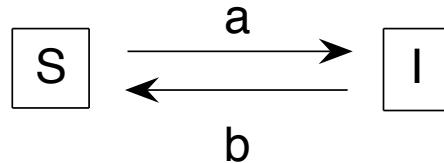


Fig. 5.2. An SIS model.

Here a is the parameter associated with catching the disease, and b is the parameter associated with recovery from the disease. Recovery is something that occurs naturally without the need for an “encounter” of some kind (i.e., you don't need to meet someone well to get better), so the number of individuals recovering is simply proportional to the number of sick individuals; hence b represents the average fraction of individuals that are able to recover at a given time. Our system of ODEs follows quite nicely from Figure 5.2:

$$\frac{dS}{dt} = bI - aSI \quad (5.3)$$

$$\frac{dI}{dt} = aSI - bI. \quad (5.4)$$

Can you see what's happening, from these equations? Individuals move into the Susceptible compartment only when they recover (bI) and move out when they become infected ($-aSI$). Individuals move into the Infected compartment when they become infected (aSI) and move out when they recover ($-bI$).

Let $N = S + I$ be the total population. Note that if we add the equations together, then $\frac{dS}{dt} + \frac{dI}{dt} = \frac{dN}{dt} = 0$. Thus, the total population is constant (since the derivative of a constant is zero).

To analyse these equations, we have two options. We can attempt to solve them directly and find the *time series*: the time-course of the disease, where the time variable is explicit. Or we can make time an *implicit* variable and discuss solutions in state space consisting of only S and I . We'll look at both methods in the following sections.

5.1.1 Solving directly

Our first instinct is to solve the differential equations. We'll try that, but it turns out that our first instinct might be misleading. Most differential equations can't be solved explicitly (only numerically). Even for those that can be, it requires a lot of work for a payoff that might not be as useful as you think.

System (5.3)–(5.4) is solved in Appendix A. As with all the appendices, follow along if you can, but don't sweat the details if you can't. We can find the entire solution, which is

$$S(t) = N - \frac{(aN - b)I_0 e^{(aN-b)t}}{(aN - b) + aI_0[e^{(aN-b)t} - 1]}$$

$$I(t) = \frac{(aN - b)I_0 e^{(aN-b)t}}{(aN - b) + aI_0[e^{(aN-b)t} - 1]}.$$

Bear in mind that we're explicitly noting that S and I depend on time... but of course they do, since they had derivatives in (5.3)–(5.4). Otherwise the derivative would automatically be zero. This is how you can tell the difference between variables and parameters: variables are in the left-hand side of the differential equations and are things that will always vary; we're not solving them for a single number but rather for a function. Conversely, parameters (like a , b , N and I_0) are numbers that we happen not to know.

Hooray! We have a solution! It was a lot of work, but we have one. So now what? Problem solved and we can go home?

Well... maybe. Sure, we have a solution, but we still need to interpret it. And, let's be honest, this probably isn't the most grabholdable thing you've ever seen. It's possible to see from these equations what's happening... but it isn't exactly easy. And remember, we got lucky this time around; most differential equations aren't so easy.

What's really happening here is that we weren't asking the right question. We were asking "What's the solution to the differential equations?" but what

we really should have been asking was “What’s going on?” or “What happens in the long run?” Mathematics is good for answering very specific questions, but it’s important not to lose sight of the big picture.

We’ll get to these questions in the next subsection, but for now we could plot the solutions numerically. Except... why did we need to solve the equations analytically if we could just solve them numerically? (Answer: we didn’t.)

There is one thing to notice from the solutions, however. If you look at the exponentials that appear in every term, they look like $e^{(aN-b)t}$. There will be very different outcomes if $aN - b < 0$ versus $aN - b > 0$. In the former case, the exponentials will go to zero; in the latter, they’ll blow up. It turns out that either option is okay, but it’s important to know which one we’re dealing with. So we need to i) notice this in the first place and ii) deal with two cases.

What if $aN - b = 0$? In theory, no problem: the exponentials will be constants, and we have the solution. However, in practice, this is a knife-edge case. Tiny fluctuations in a , N and b could tip the balance one way or the other, making the results unpredictable. So we’re actually not that interested in these kinds of knife-edge cases. These sorts of issues are going to pop up in all kinds of iterations throughout this book.

The solutions are plotted in Figure 5.3. In both graphs, $N = 1$, $I_0 = 0.25$ and $a = 0.6$. In the left-hand graph, $b = 0.05$; in the right-hand graph, $b = 0.7$. (In reality, you probably wouldn’t start with a quarter of the population infected, but we’ve chosen this to clearly illustrate what happens in the second case.)

Thus we see that sometimes the infected individuals die out and everyone becomes susceptible again (the right-hand graph). In this case, there’s no epidemic, and the disease disappears on its own. However, at other times, the population reaches an equilibrium (steady state) in both infected and susceptible individuals (the left-hand graph). In this case, the disease has taken hold within the population and has become endemic.

This is pretty crucial. For one set of parameters, we don’t have to worry about this disease; even with a million people infected, it’ll die out in time. For the other, we really care. Even a few infected individuals could lead to a pandemic. The big question is: how do we know when such a split occurs?

The time series is great, if we can do it, because it gives us the complete solution, at every point in time. So if we want to know what happens at the beginning, middle or end of an epidemic, we simply have to pick a time we’re interested in and plug it in to the equations. But most of the time we’re not interested in the beginning or middle of an epidemic, just the end. And usually that means: will the disease die out on its own, or will it become endemic?

Deriving these equations was pretty hard, even for such a simple model. So maybe we’ve done more work than we needed, if we only want to know what happens eventually. And, in general, we can’t solve an arbitrary system of ODEs. So we need another option.

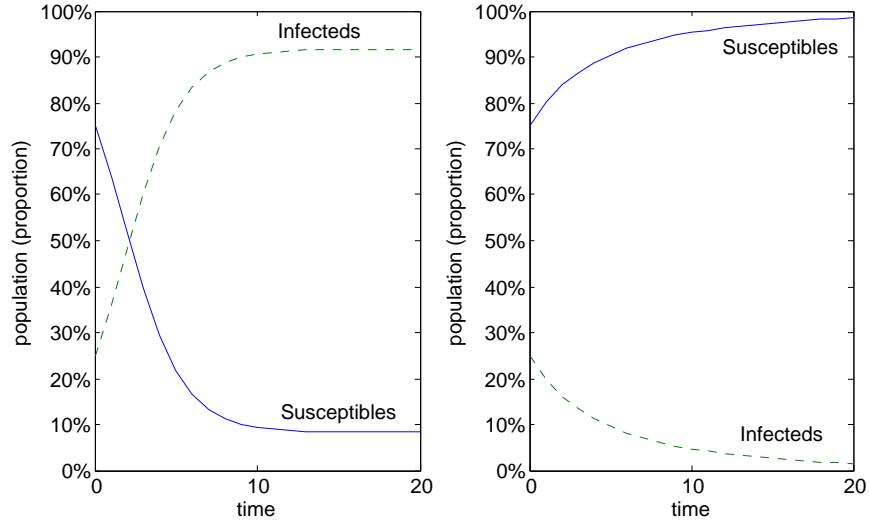


Fig. 5.3. The time series of the infection with $N = 1$ and $a = 0.6$. Left: $b = 0.05$. Right: $b = 0.7$.

5.1.2 Phase portraits

Instead of doing all that math every time, we can take another, less analytical, approach to find information about our SIS system. Phase portraits involve only the spatial variables, making time an implicit part of the solution. This gives us less overall information, but still tells us what happens at the beginning and end of an epidemic.

Since $N = S + I$, it follows that $I = -S + N$, which, when graphed in the S - I plane, is a straight line. See Figure 5.4. We have thus parameterised our equations so they are independent of time. The population must travel along the line — although, since we eliminated time, we have no idea how it does so.

Next, we look for equilibrium values. Equilibrium values of ODEs occur when there is no change in the system; thus, when the derivatives are zero; i.e., $\frac{dS}{dt} = 0$ and $\frac{dI}{dt} = 0$. Hence, we set equations (5.3)–(5.4) to zero:

$$(b - aS)I = 0 \quad (5.5)$$

$$(aS - b)I = 0. \quad (5.6)$$

Since (5.5) and (5.6) are identical equations (one is the negative of the other), they'll hold when either i) $S = \frac{b}{a}$ and $I = N - \frac{b}{a}$ (since $S + I = N$); or ii) $I = 0$ and $S = N$ (since $S + I = N$).

For simplicity, let's call $\frac{b}{a}$ “ p ”. Our equilibrium points are thus

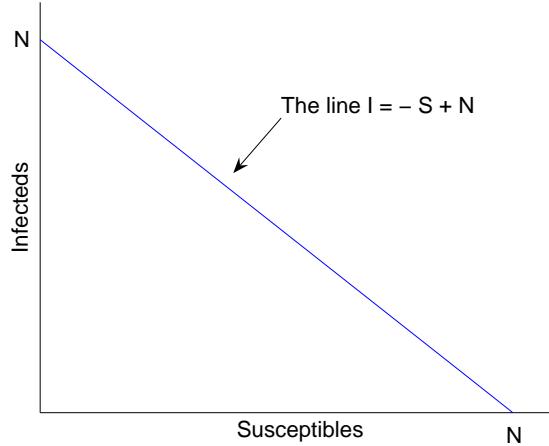


Fig. 5.4. The linear relationship between S and I .

$$(\bar{S}, \bar{I}) = (p, N - p)$$

or

$$(\bar{S}, \bar{I}) = (N, 0).$$

In terms of realistic solutions, the equilibrium $(N, 0)$ will always exist, but the equilibrium $(p, N - p)$ only exists when $p < N$ (if $p > N$, then the “second” equilibrium would have $S > N$, so this equilibrium it isn’t biologically realistic). We thus have two cases:

- Case I: $p < N$.
- Case II: $p > N$

In Case I, there will be two equilibrium points, whereas in Case II there will only be one equilibrium in the positive plane (which is all we’re interested in). See Figure 5.5. Such a time-independent diagram is known as a *phase portrait*. (Don’t forget that $p = \frac{b}{a}$.)

These two cases correspond exactly to the time series from Figure 5.3. In the case $p < N$ (which was when $b = 0.1$ in Figure 5.3), there’s an interior equilibrium, with both susceptibles and infecteds. In the case $p > N$ (which was when $b = 0.7$ in Figure 5.3), there’s only one equilibrium, which is when the infection has died out and only susceptibles remain.

Our major question now is “are these equilibrium points stable?” (You may be able to guess the answer from the time series!) We can rewrite equations (5.3)–(5.4) as

$$\begin{aligned} \frac{dS}{dt} &= a \left(\frac{b}{a} - S \right) I \\ &= a(p - S)I \end{aligned}$$

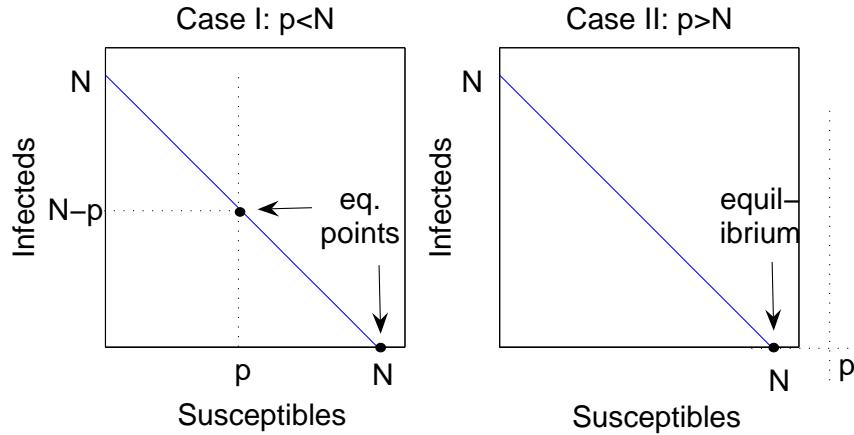


Fig. 5.5. Equilibria for the two cases.

and

$$\begin{aligned}\frac{dI}{dt} &= a \left(S - \frac{b}{a} \right) I \\ &= a(S - p)I.\end{aligned}$$

If $S < p$, then clearly $\frac{dS}{dt} > 0$ and $\frac{dI}{dt} < 0$ (since $p - S > 0$). Conversely, if $S > p$, then $\frac{dS}{dt} < 0$ and $\frac{dI}{dt} > 0$ (since $p - S < 0$).

Thus, in Case I, for the portion of the line to the left of the equilibrium $(p, N - p)$, S is increasing and I is decreasing. Trajectories must therefore be moving down the line. For the portion of the line to the right of the equilibrium $(p, N - p)$, S is decreasing and I is increasing. Trajectories must therefore be moving up the line. It follows that the equilibrium $(p, N - p)$ is stable (since trajectories approach it) while the equilibrium $(N, 0)$ is unstable (since trajectories move away from it).

For Case II, since $S < p$, along the line S is increasing and I is decreasing. Thus the equilibrium $(0, N)$ is stable. For an unstable equilibrium, any slight change will drive the dynamics away from this point and towards the stable point. (Note that, in Figure 5.3, we “see” the stable equilibrium, but of course we can’t “see” an unstable one.)

We can thus redraw the phase portraits with arrows indicating the direction along the line the population travels in Figure 5.6.

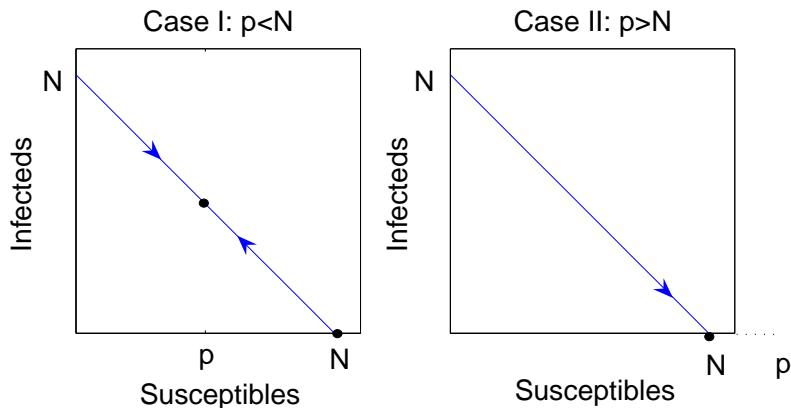


Fig. 5.6. The stability of the equilibrium points.

Thus, when $p > N$, the number of infected people will decrease, approaching zero. When $p < N$, there is a stable equilibrium where there are p susceptible and $N - p$ infected people. There is also an unstable equilibrium at $(N, 0)$, where even one infected person leads to an increase in the number of infected people until the stable equilibrium $(p, N - p)$ is reached.

The take-home message of all this is that the stability depends on $p = \frac{b}{a}$. So if the recovery rate b is very high compared to the infection rate a , then the infecteds recover quickly and the population moves to a population of susceptibles. If the infection rate a is very high compared to the recovery rate, then the infection stabilises at an endemic equilibrium.

Phase portraits are the way stability is usually determined, since the analysis is a lot easier and we only need to determine the *direction* of the arrows, not the entire solution. You might be familiar with the Lotka–Volterra predator–prey systems, which involve annual cycling of solutions. Our example here is analogous, except that we have used very basic assumptions (e.g., no birth or death rates).

5.2 SIR epidemics

SIR stands for Susceptible → Infected → Removed. The term “removed” is a general term that allows for infected individuals to be no longer infected yet not susceptible either. Practically, this could mean that the person gets better (either through treatment or natural immunity) or that the person dies (the ultimate removal). See Figure 5.7.

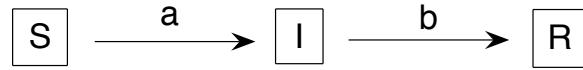


Fig. 5.7. An SIR model.

The equations are

$$\begin{aligned}\frac{dS}{dt} &= -aSI \\ \frac{dI}{dt} &= aSI - bI \\ \frac{dR}{dt} &= bI,\end{aligned}$$

with $N = S + I + R$ (constant). We assume that the rate of recovery/death is proportional to the number of sick people.

One modification we might make could be the administration of a vaccine. Susceptible individuals who are vaccinated would then go directly to the ‘Immune’ category (we can think of “recovered” people as immune). The number of vaccinations given to people would likely be a certain number of shots per day or time period. Thus, this term in the differential equations will not involve any of the S , I or R variables; rather, it will just be some parameter c that represents the number of vaccinations given per time period. See Figure 5.8.

The differential equations are

$$\begin{aligned}\frac{dS}{dt} &= -aSI - c \\ \frac{dI}{dt} &= aSI - bI \\ \frac{dR}{dt} &= bI + c,\end{aligned}$$

with $N = S + I + R$ as before (the total number of individuals hasn’t changed).

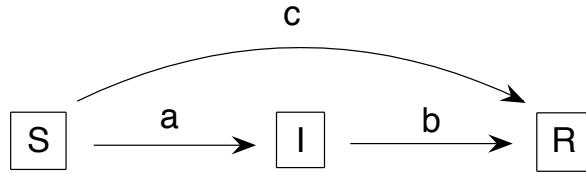


Fig. 5.8. An SIR model with vaccination.

In this case, many of the susceptibles become immune before they encounter the disease, so spatial considerations will come into place as the number of available susceptibles shrink. It's a complex relationship, because not everyone rushes out to get the vaccine. Some people will rely on "herd immunity" to protect them (the fact that if most people around you are vaccinated, then you're unlikely to catch the disease). Remember that we assumed our populations were well-mixed, but this mixing may not be appropriate. If so, the ODE model wouldn't apply any more and we'd need to use partial differential equations (PDEs) to deal with the spatial component. See Chapters 8 and 9 for more on this.

Another modification to the SIR model is to include a mutation. Some viruses can mutate over a given time T then come back with a vengeance. This time factor can be built in to allow recovered people a "grace period" before the disease mutates and they are once again susceptible. See Figure 5.9.

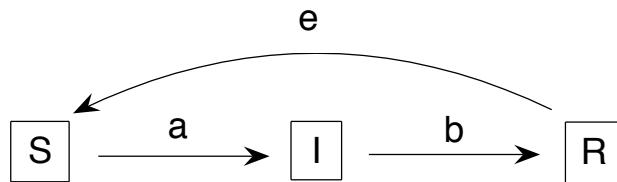


Fig. 5.9. An SIR model with mutation.

The differential equations are

$$\begin{aligned}\frac{dS(t)}{dt} &= -aS(t)I(t) + eR(t - T) \\ \frac{dI(t)}{dt} &= aS(t)I(t) - bI(t) \\ \frac{dR(t)}{dt} &= bI(t) - eR(t - T),\end{aligned}$$

with $N = S + I + R$, as before. In this case there is a delay of time T , so we need to explicitly include time t , unlike the previous two models. Events at time t depend on what happened at an earlier time $t - T$. Such equations are called *delay differential equations*. See Chapter 17 for a detailed application of this.

5.3 Lab work

Not every person can fit in one of three convenient categories. Here we explore an epidemic-type problem that requires more than three compartments.

The problem

Suppose I_0 people in a population of N people have been infected with a lethal disease. Assume also that the death rate of the infected population $I(t)$ is proportional to its population. A large shipment of medicine will take 9 days to get there. In the meantime, you can only treat 10 people per day for a total of 9 days. People who receive treatment can no longer be infected and are not contagious. The rate at which the disease spreads, k , is equal to 6.90675×10^{-5} people $^{-1} \cdot$ day $^{-1}$ (the units must match the units on the left-hand side of the differential equation), the death rate of infected individuals is 0.1 per day and $N = 10000$ people. At $t = 0$, $I_0 = 20$. We'll work through this in detail, but see if you can guess or approximate the answers to the first three questions before reading the solution:

- Write down the system of ODEs that describe this system (without the numbers).
- How many people will survive after 9 days?
- If there was no medicine, how many people would survive after 9 days?
- Obtain a phase portrait of S vs. I for $0 \leq t \leq 9$, with the rate of treatment $b = 10$ people/day.

The solution

We first create our epidemic diagram. Notice that while we still have our S and I compartments, we no longer have only a third “removed” compartment;

we have two others, which we will call M (treated people) and D (dead). We have the “flow parameters” k for $S \rightarrow I$, b for $I \rightarrow M$ and r for $I \rightarrow D$.

To become infected requires interaction between S and I . Death is proportional simply to the number of sick people I , and a constant number of shots are being given, regardless of any other number (assuming $I > 0$, since only sick people are treated). See if you can draw the diagram yourself. The answer is in Figure 5.11 on Page 60, but try and work it out for yourself first.

From this we can read off our ODEs:

$$\begin{aligned}\frac{dS}{dt} &= -kSI \\ \frac{dI}{dt} &= kSI - rI - b \\ \frac{dM}{dt} &= b \\ \frac{dD}{dt} &= rI.\end{aligned}$$

We must first define our ODE as a function M-File:

```
function pdot=epidemf(t,p)
%This is the ODE for the epidemic problem
k=6.90675e-5;
r=0.1;
b=10;
pdot(1,:)=-k.*p(1).*p(2);
pdot(2,:)=k.*p(1).*p(2)-r.*p(2)-b;
pdot(3,:)=b;
pdot(4,:)=r.*p(2);
```

This program must be saved as “epidemf.m”. You can’t run a function file on its own; you can only call it within another program that you can actually run. So we now need a program that makes use of this ODE.

```
%This is a program to use the epidemic ODE epidemf
clear all
t0=0;
tf=9;
p0=[9980,20,0,0];
tspan=[t0 tf];
[t,p]=ode23(@epidemf,tspan,p0);
plot(t,p)
```

Here “p0” contains the original number of S , I , M and D people, respectively. After we run this program, we will have the information necessary to determine how many people are still alive at $t = 9$: $S(\text{length}(t)) +$

$I(\text{length}(t)) + M(\text{length}(t))$ or $N - D(\text{length}(t))$ (since $D(\text{length}(t))$ is the number of dead people at the final time $tf = t(\text{length}(t))$). While we still have b set to 10, we might as well plot the phase portrait of S vs. I . In our program, this corresponds to $p(:, 1)$ vs. $p(:, 2)$, but don't forget that we put our horizontal component first within the *plot* command.

```
%This is a program to use the epidemic ODE epidemf
clear all
N=10000;
t0=0;
tf=9;
p0=[9980,20,0,0];
tspan=[t0 tf];
[t,p]=ode23(@epidemf,tspan,p0);
Total_still_alive=N-p(length(t),4)
pause
plot(p(:,2),p(:,1))
title('Phase Portrait - Susceptible vs. Infected for b=10')
xlabel('Infected Individuals')
ylabel('Susceptible Individuals')
```

The term $p(:, 4)$ is our D column and $p(\text{length}(t), 4)$ is the value of D at $t = tf$. Running our program, we find that 9890.3 people will still be alive and we get Figure 5.10.

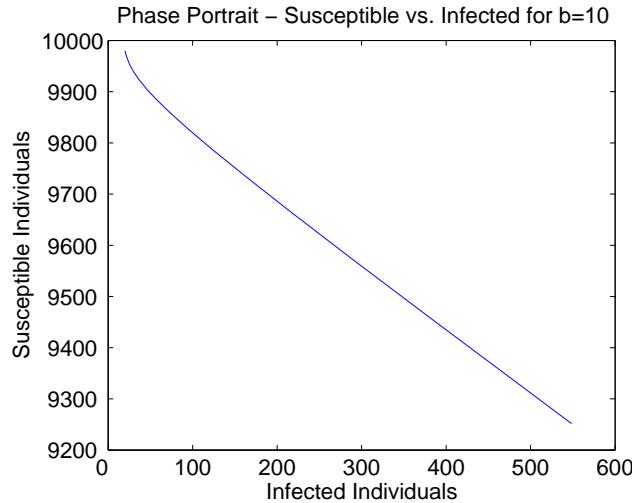


Fig. 5.10. Phase portrait for $b = 10$.

Setting $b = 0$ in our function M-File, we rerun the program and find that 9455.8 people will still be alive.

Analysis of the results

Of course, we can't have 0.3 or 0.8 of a person. In a case like this it is best to round down no matter what: the decimal place indicates that a person is in the process of dying, so we might as well ignore them. It is interesting to note that even though $10 \times 9 = 90$ shots were given, the difference in the remaining population for $b = 10$ and $b = 0$ is 435 and not 90. The relationship between b and the number of people living $N - D(\text{length}(t))$ is not linear.

5.3.1 Exercises

Let's see what happens with this disease in the long run, with and without treatment.

1. First, with $b = 0$, plot S, I, D vs. t over a long period, say $t = 200$. What do you think of this disease now?
2. Now do the same with $b = 10$. Can you explain what goes wrong when you do this?
3. Try putting in an `if p(2)>0 ... else ... end` statement and rerunning (you should be able to figure out what goes in the “...”s for yourself). Now how effective is the campaign to save everyone with treatment? Can you describe what has happened?
4. Suppose we repeat the same plots with a less infectious disease, say with k one half of what it was. How would you describe the progress of this disease, with and without treatment? Is there something qualitatively different happening here?
5. Suppose you have a disease with a latent period of infection. When you're exposed to the disease you're infected, but not infectious. You can be treated when infectious and become susceptible again, or recover naturally and have temporary immunity. Once your immunity wears off, you become susceptible again. Draw an epidemic diagram of this disease and construct a corresponding system of ODEs.

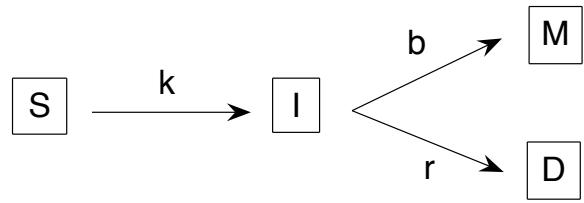
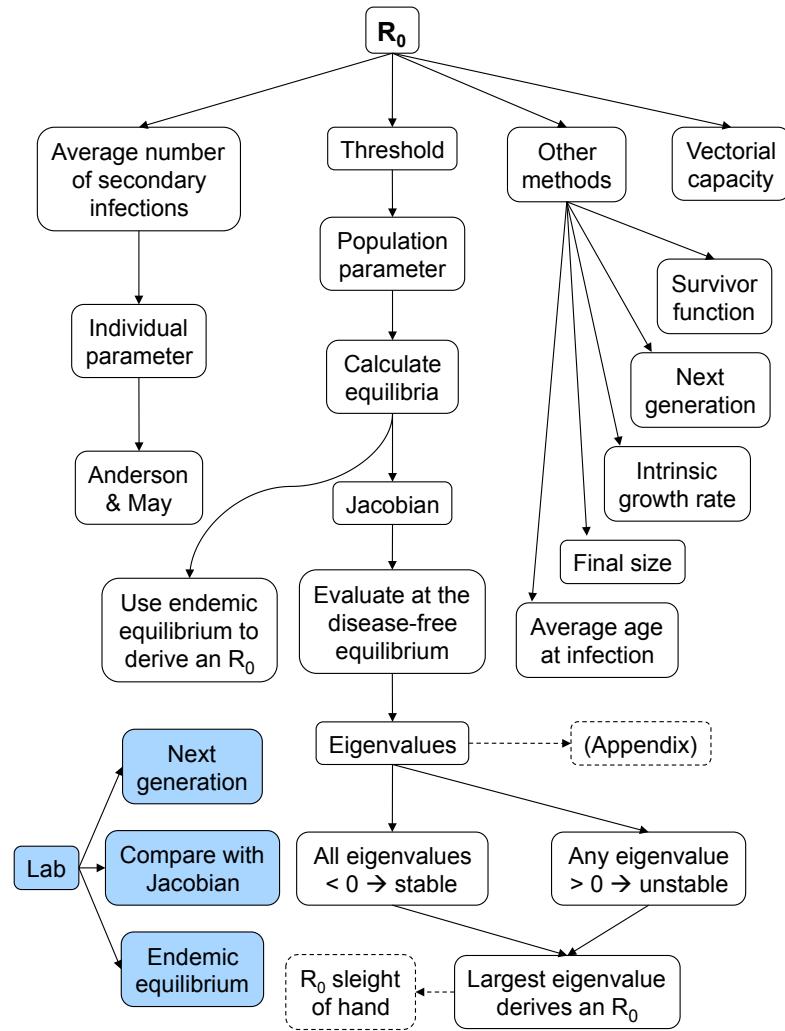


Fig. 5.11. The epidemic diagram for the SIMD model.



6

Calculating R_0

This chapter deals with the basic reproductive ratio, one of the fundamental concepts in mathematical biology. Originally developed for the study of demographics, it was independently studied for vector-borne diseases such as malaria and directly-transmitted human infections. It is now widely used in the study of infectious disease, and more recently in models of in-host population dynamics.

We also introduce the Jacobian matrix method, one of the most useful things you can learn in disease modelling. We discuss how to derive the basic reproductive ratio from the Jacobian, as well as a few other methods. This is the most theoretical of all the chapters. Many of the ideas here will be used again in subsequent chapters.

By the end of this chapter, you should be able to distinguish between R_0 and a threshold parameter, calculate the Jacobian matrix for a system of ODEs, apply the next-generation matrix and derive a threshold from an endemic equilibrium.

6.1 Threshold parameters

The basic reproductive ratio, R_0 , is defined as “the average number of secondary infections caused by a single infectious individual during their entire infectious lifetime.” It’s a measure of how quickly a disease spreads in its initial phase and can predict whether a disease will become endemic or whether it will die out. There is a threshold when $R_0 = 1$ such that if $R_0 > 1$, then, on average, each individual is causing more than one infection, so we expect the disease to take hold within a susceptible community. On the other hand, if $R_0 < 1$, then each infectious individual is leaving the infectious compartment with fewer (on average) infections than when they entered.

One thing to notice immediately is that R_0 is a parameter involving individuals, whereas SIS/SIR dynamics deal with compartments (populations). There’s some complicated matching required here, because it’s not so trivial

to match individual behaviour with compartments. Or rather, we can match them, but it's not a very neat matching, since a variety of different individual dynamics can lead to the same population model.

When measuring data in the field, we may occasionally have the chance to measure the average number of secondary infections caused by a single individual. If so, great. But it's rare that we get to study a disease in its very earliest stages of infection; many diseases are often not noticed until they're well underway. When this happens, we won't be able to measure R_0 precisely, because spatial considerations get in the way. A single infectious individual may be able to infect any susceptible individual, but multiple infectious individuals may "use up" their infectivity on those already infected. If 90% of a village is infected, then a single infected individual may have a fairly remote chance of infecting a susceptible individual, even if $R_0 > 1$, because they may never come into contact with them.

However, it's also possible that we may not actually want the true value of R_0 . Often what we'd really like to know are two things: 1) Will the disease become endemic, or will it die out of its own accord? 2) If the disease is endemic, will our control strategies be sufficient? Both of these questions are really concerned with the *threshold*, not the true value of R_0 . Of course, we know that R_0 is a threshold, but it turns out that there are many surrogate thresholds that are not R_0 .

These values will also have the property that if they are less than 1, then the disease will die out, whereas if they are greater than 1, the disease will take hold. What they can't do is allow comparison of different diseases; if SARS has a threshold value of 6 and HIV has a threshold value of 4, we can't actually conclude that SARS is a worse disease than HIV unless the same method was used to calculate each. It turns out that almost everybody has been confused about this over the years. If you've ever seen an R_0 in the literature, take it with a grain of salt. Unless you know the precise method used to calculate it (especially if it was derived from a mathematical model), chances are it's probably wrong. It will tell you whether your disease is endemic or not, but it may not be a true indicator of the average number of secondary infections and it shouldn't be used to compare different diseases.

6.2 The method of Anderson & May

The most well-known version of R_0 comes from Anderson & May (1991), who basically posit that

$$R_0 = \beta c D,$$

where β is the transmissibility, c is the number of contacts and D is the average time spent infectious (and if the rate of infection is b , then $D = \frac{1}{b}$). It turns out that this is quite useful and applies quite widely, under a very important

qualifier: namely, that there is no background death rate. It also applies only to SIR models (so, for example, no latent period) and assumes homogeneity and exponentially distributed times. Anderson & May included these assumptions in their work, but, unfortunately, a lot of people have applied their results to diseases where the background death rate is important.

For instance, HIV is a disease that eventually kills you if you're infected with it, but not for a long time, especially if you're on antiretroviral drugs. So if you're infected with HIV, the chances that you'll get killed in a car accident rather than die from AIDS are non-negligible, because you have a small but nonzero chance of getting killed in a car accident in the intervening ten (or more) years. So the Anderson & May definition wouldn't apply here.

A disease like Ebola, on the other hand, occurs on such a short timescale that the background death rate is negligible. If your village is infected with Ebola, the disease will rage through and quickly kill everyone it's going to kill in a matter of weeks. The chance of you dying from a car accident, influenza or old age during this time is pretty small. So, in this case, it would be reasonable to use the Anderson & May definition.

Unfortunately, most diseases operate on a fairly long timescale, so the Anderson & May definition isn't as useful as it might seem. However, it reinforces the importance of being careful about the assumptions under which one can apply a particular theory, because there have been a lot of cases of misapplication of this definition by biologists over the years.

6.3 The Jacobian

For simple two-dimensional systems, we can create a phase portrait, as we did in Section 5.1.2. For higher dimensional systems, we can't draw such a portrait (or we could only draw two of the dimensions), so we need a more general method. This method involves creating a matrix called the *Jacobian* matrix. This is a matrix of partial derivatives, created by differentiating every equation with respect to every variable. So if there are six equations and six variables, you'll get a 6×6 matrix.

The method of determining stability of equilibrium points for ODEs is pretty straightforward, in that the same method can be applied to most systems, so long as there's an actual equilibrium and the mathematics is tractable. The steps are:

1. Calculate the equilibrium, usually the disease-free equilibrium (which is often the easiest, fortunately).
2. Create the Jacobian matrix by differentiating every equation with respect to every variable.
3. Evaluate this matrix at the equilibrium of interest. (Usually this simplifies things significantly.)

4. Find the eigenvalues. (Eigenvalues are numbers that “represent” a matrix; if we have a matrix A and can find a number λ and a nonzero vector x such that $Ax = \lambda x$, then λ is an eigenvalue and x is an eigenvector. See Appendix C.)
5. If all eigenvalues are negative, the equilibrium is stable. If even one eigenvalue is positive, the equilibrium is unstable. (If there’s a zero eigenvalue, then we have no conclusion; we need to use higher-order techniques to deal with this.)
6. If the eigenvalues are complex, then the previous condition applies only to the real part of the eigenvalues. So if the real parts are all negative, the equilibrium is stable; if any eigenvalue has positive real part, the equilibrium is unstable. (That is, we basically ignore the imaginary part altogether. The contribution of the imaginary part is to induce oscillations.)
7. Use the largest eigenvalue to calculate an R_0 -like threshold parameter.

The Jacobian matrix is a very useful matrix for analysing stability. It’s a square matrix, where the first column is the partial derivatives of every equation with respect the first variable, the second column is the partial derivatives of every equation with respect to the second variable, and so on.

For example, you might be familiar with the Leslie matrix. The Leslie model is a discrete-time model of an age-structured population which describes development, mortality and reproduction of organisms. This model is mostly used to answer the following two questions: What is the rate of exponential growth (intrinsic rate of increase)? What is the proportion of each age class in the stable age distribution?

The three dimensional Leslie model is

$$\begin{aligned} x_1^{(k)} &= F_1 x_1^{(k-1)} + F_2 x_2^{(k-1)} + F_3 x_3^{(k-1)} \\ x_2^{(k)} &= P_1 x_1^{(k-1)} \\ x_3^{(k)} &= P_2 x_2^{(k-1)}, \end{aligned}$$

which can be rewritten as

$$\begin{pmatrix} x_1^{(k)} \\ x_2^{(k)} \\ x_3^{(k)} \end{pmatrix} = \begin{pmatrix} F_1 & F_2 & F_3 \\ P_1 & 0 & 0 \\ 0 & P_2 & 0 \end{pmatrix} \begin{pmatrix} x_1^{(k-1)} \\ x_2^{(k-1)} \\ x_3^{(k-1)} \end{pmatrix}.$$

(It can only be done like this because it’s a linear system.) The Leslie matrix is

$$\begin{pmatrix} F_1 & F_2 & F_3 \\ P_1 & 0 & 0 \\ 0 & P_2 & 0 \end{pmatrix},$$

but this would also be the Jacobian matrix for this model, if you took the partial derivatives (ignoring the fact that the Leslie model is a model of difference equations, while we are dealing with differential equations). The dominant eigenvalue of the Leslie matrix is the growth rate.

6.3.1 The SIS model threshold

For the SIS model, we have two equations,

$$\begin{aligned} S' &= bI - aSI \\ I' &= aSI - bI, \end{aligned}$$

and two variables (S and I). The Jacobian matrix will be

$$\begin{aligned} J &= \begin{bmatrix} \frac{\partial S'}{\partial S} & \frac{\partial S'}{\partial I} \\ \frac{\partial I'}{\partial S} & \frac{\partial I'}{\partial I} \end{bmatrix} \\ &= \begin{bmatrix} -aI & b - aS \\ aI & aS - b \end{bmatrix}. \end{aligned}$$

We've already calculated the equilibrium values in the previous chapter. The disease-free equilibrium is clearly $(N, 0)$. This is because everyone is susceptible and no one is infected; thus, disease-free.

Evaluating the Jacobian matrix at this equilibrium, we have

$$J|_{(N,0)} = \begin{bmatrix} 0 & b - aN \\ 0 & aN - b \end{bmatrix}.$$

The eigenvalues for this matrix are

$$\lambda = 0, aN - b$$

(see Appendix C). One of these eigenvalues ($\lambda = 0$) is unchanged no matter what, so the stability will be determined by the sign of the other eigenvalue ($\lambda = aN - b$). Specifically, if $aN - b > 0$ then there's a positive eigenvalue, so the disease-free equilibrium will be unstable. What if $aN - b < 0$? In this case, we don't know the stability of the disease-free equilibrium since one of the eigenvalues is zero. (The reason for the zero eigenvalue in this simple example is because the system is overdetermined; it's really one equation, not two, but we've kept it simple here for illustrative purposes.) However, in this case, it actually is stable if $aN - b < 0$, so let's ignore the zero eigenvalue and use the nontrivial eigenvalue as our threshold.

Everything is well and good to this point. The next step is the dodgy one. See if you can spot the flaw. We have stability if

$$aN - b < 0. \tag{6.1}$$

Moving the negative to the other side, we'll have stability if

$$aN < b.$$

Dividing both sides by b , we'll have stability if

$$\frac{aN}{b} < 1.$$

Thus we could define R_0^{SIS} to be the threshold

$$R_0^{\text{SIS}} = \frac{aN}{b},$$

since if $R_0^{\text{SIS}} < 1$, then we'll have stability, whereas if $R_0^{\text{SIS}} > 1$, then we'll have instability. Hence, the disease will die out if $R_0^{\text{SIS}} < 1$ and will become endemic if $R_0^{\text{SIS}} > 1$.

Did you spot the mathematical sleight of hand there? (There are actually two!) The answer's in the Appendix D, but try and figure it out for yourself if you can. (Don't be discouraged if you can't; brilliant mathematicians have been missing this for over a century.)

The numerator is the rate of infection, made up of the per-capita transmissibility (a) multiplied by the number of susceptibles (N), while the denominator is the rate of recovery (b). Essentially R_0^{SIS} will be greater than one if the rate of infection exceeds the rate of recovery and less than one if the reverse is true. So this threshold makes some intuitive sense.

More generally though, $\frac{1}{b}$ is actually the average length of time an individual spends infectious, so R_0^{SIS} is the product of the number of interactions (N), the transmissibility (a) and the duration of the infectious period $\frac{1}{b}$. This matches the Anderson & May definition, as we have no background death rate in the SIS model.

6.3.2 The endemic equilibrium for the SIS model

There's another way to do this, however. (In fact there are many, as we'll see shortly.) The other equilibrium is $(\frac{b}{a}, N - \frac{b}{a})$. This is the endemic equilibrium, and, from equation (6.1), we see that the disease-free equilibrium is unstable whenever the endemic equilibrium exists. So, although we haven't proven it, we have good reason to suspect that whenever the endemic equilibrium exists, it's probably stable.

Thus we could use the existence of the endemic equilibrium to predict the long-term outcome of the disease. If the endemic equilibrium doesn't exist, then the disease-free equilibrium will be stable, and hence, in the long run, we won't have any infection. On the other hand, if the endemic equilibrium exists, then the disease will persist (in fact, in the long run, it will persist around this equilibrium).

It follows that the existence of the endemic equilibrium is just as good a method as the Jacobian method for defining an R_0 . Since the condition for the existence of the endemic equilibrium is precisely equation (6.1), we'll thus derive exactly the same R_0^{SIS} as before.

A word of caution, however: for more complicated models, these two methods might derive *different* R_0 values. Both will work as thresholds, but it's not clear which, if either, is necessarily the "true" R_0 .

6.3.3 The SIR model threshold

For this last example, we'll adjust the SIR model slightly to include a constant birth rate π and a background death rate μ , proportional in each class to the number of people in that class. In this case, the removed class are the immune, not the dead (they can't die twice). The model now becomes

$$\begin{aligned}\frac{dS}{dt} &= \pi - aSI - \mu S \\ \frac{dI}{dt} &= aSI - bI - \mu I \\ \frac{dR}{dt} &= bI - \mu R.\end{aligned}$$

The disease-free equilibrium is easily calculated by setting $I = 0$. It follows immediately that $R = 0$ and $S = \frac{\pi}{\mu}$ (why?). The Jacobian matrix is

$$\begin{aligned}J &= \begin{bmatrix} \frac{\partial S'}{\partial S} & \frac{\partial S'}{\partial I} & \frac{\partial S'}{\partial R} \\ \frac{\partial I'}{\partial S} & \frac{\partial I'}{\partial I} & \frac{\partial I'}{\partial R} \\ \frac{\partial R'}{\partial S} & \frac{\partial R'}{\partial I} & \frac{\partial R'}{\partial R} \end{bmatrix} \\ &= \begin{bmatrix} -aI - \mu & -aS & 0 \\ aI & aS - b - \mu & 0 \\ 0 & b & -\mu \end{bmatrix}.\end{aligned}$$

At the disease-free equilibrium, we have

$$J \Big|_{\left(\frac{\pi}{\mu}, 0, 0\right)} = \begin{bmatrix} -\mu & -\frac{a\pi}{\mu} & 0 \\ 0 & \frac{a\pi}{\mu} - b - \mu & 0 \\ 0 & b & -\mu \end{bmatrix}.$$

To calculate the determinant of a 3×3 matrix, we apply the formula

$$\det \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & j \end{bmatrix} = aej + bfg + cdh - ceg - afh - bdj.$$

(See Appendix C.) Fortunately, all but one of these products will be zero in our case. To calculate the eigenvalues we thus have

$$\begin{aligned}
\det(J - \lambda I) &= \det \left(\begin{bmatrix} -\mu & -\frac{a\pi}{\mu} & 0 \\ 0 & \frac{a\pi}{\mu} - b - \mu & 0 \\ 0 & b & -\mu \end{bmatrix} - \lambda \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right) \\
&= \det \begin{bmatrix} -\mu - \lambda & -\frac{a\pi}{\mu} & 0 \\ 0 & \frac{a\pi}{\mu} - b - \mu - \lambda & 0 \\ 0 & b & -\mu - \lambda \end{bmatrix} \\
&= (\mu + \lambda)^2 \left(\frac{a\pi}{\mu} - b - \mu - \lambda \right).
\end{aligned}$$

The eigenvalues are thus

$$\lambda = -\mu, -\mu, \frac{a\pi}{\mu} - b - \mu.$$

The first two eigenvalues are constants and always negative (for a positive death rate μ), so they won't determine our stability. The stability is thus determined by the third eigenvalue. Moving the negatives to the other side and dividing, as before, we thus derive our threshold parameter

$$R_0^{\text{SIR}} = \frac{a\pi}{\mu(b + \mu)}.$$

The total time spent infectious is now $\frac{1}{b+\mu}$, since patients will stop being infectious if they either recover (b) or die (μ). Thus R_0^{SIR} is once again a product of the total number of interactions ($N = \pi/\mu$ in this case), the transmissibility (a) and the time spent infectious $\left(\frac{1}{b+\mu}\right)$. Although this matches the Anderson & May definition, this is just coincidence, since we now include the background death rate.

6.3.4 The endemic equilibrium for the SIR model

After a bit of algebra (try for yourself), we can find the endemic equilibrium:

$$(\bar{S}, \bar{I}, \bar{R}) = \left(\frac{b + \mu}{a}, \frac{\pi}{b + \mu} - \frac{\mu}{a}, \frac{b\pi}{\mu(b + \mu)} - \frac{b}{a} \right).$$

So, once again, the existence of the endemic equilibrium could also be used to determine the threshold parameter, since

$$\begin{aligned}
\frac{\pi}{b + \mu} - \frac{\mu}{a} &> 0 \\
\frac{\pi}{b + \mu} &> \frac{\mu}{a} \\
\frac{a\pi}{\mu(b + \mu)} &> 1,
\end{aligned}$$

which means we would define the same R_0^{SIR} as before. Note that this won't be true in general, as models get more complicated.

6.4 Other methods for calculating R_0

For completeness, other methods for calculating R_0 are listed here, along with an example of their uses in policy. However, if you’re not interested in the wildly diverse methods for calculating thresholds that appear to be the basic reproductive ratio (but often aren’t), feel free to skip ahead to Section 6.5 on page 77, except for the next-generation method (which you’ll need for the exercises). We’ll be encountering one more method for calculating R_0 in the next chapter anyway.

6.4.1 Survival function

The method is, in essence, a first principles definition of R_0 . It always gives the correct R_0 but can be difficult to apply in practice.

Consider a large population, and let $F(a)$ be the probability that a newly infected individual remains infectious for at least time a . This is called the survival probability. Let $b(a)$ denote the average number of newly infected individuals that an infectious individual will produce per unit time when infected for total time a . Then R_0 is given by:

$$R_0 = \int_0^\infty b(a)F(a)da. \quad (6.2)$$

As this expression yields R_0 by definition, this approach will be appropriate for any model in which closed-form expressions can be given for the underlying survival probability $F(a)$ and the infectivity as a function of time, $b(a)$. In particular, it is straightforward to handle situations in which infectivity depends on time since infection or other transmissibilities between states vary with time. This derivation of R_0 is not restricted to systems described by ODEs.

This method can also be naturally extended to describe models in which a series of states are involved in the “reproduction” of an infected individual. As an example of the latter technique, consider epidemic modelling of malaria. An infected human may pass the infection to a mosquito, which may in turn infect more humans. This complete cycle must be taken into account in our derivation of R_0 , which we might expect to yield the total number of infected humans produced by one infected human. In general, if only two distinct infectious states are involved in such an *infection cycle*, $F(a)$ can be defined as the probability that an individual in state 1 at time zero produces an individual who is in state 2 until at least time a . Similarly, $b(a)$ is the average number of new individuals in state 1 produced by an individual who has been in state 2 for time a . In modelling malaria, $F(a)$ could be the probability that a human infected at time zero produces an infected mosquito who remains alive until at least time a . In more concrete terms, $F(a)$ would be

$$F(a) = \int_0^a \text{probability(human infected at time 0 exists at time } t) \\ \times \text{probability(human infected for total time } t \text{ infects mosquito)} \\ \times \text{probability(infected mosquito lives to be age } a - t) dt,$$

while $b(a)$ would simply be the average number of humans newly infected by a mosquito which has been infected for time a . (Note that we could also take the infected mosquito as state 1, deriving an analogous expression which would yield the same value of R_0 .)

Uses in Policy: Luz et al. (2003) used R_0 to evaluate the risk of dengue fever outbreaks in Rio de Janeiro, and to assess possible control measures. R_0 was calculated from the survival function, assuming two spatial compartments with high and low vector density, respectively. The goal of this paper was to assess which of the many unknown parameter values are most important to the model. Luz et al. concluded that field estimates of mosquito mortality and the incubation period of dengue in mosquitos are of critical importance.

6.4.2 Next-generation method

The next-generation method (see Diekmann *et al.* (1990), Diekmann & Heesterbeek (2000), and van den Driessche & Watmough (2002)) is a general method of deriving R_0 in situations in which the population is divided into discrete, disjoint classes. The next-generation matrix can thus be used for models with underlying age structure or spatial structure, among other possibilities. For typical implementations, continuous variables within the population are approximated by a number of discrete classes. This approximation assumes that transmissibilities between states are constant.

In the next-generation method, R_0 is defined as the largest eigenvalue of the next-generation matrix. The formation of this matrix involves determining two compartments, infected and non-infected, from the model.

Let us assume that there are n compartments of which m are infected. We define the vector $\bar{x} = x_i, i = 1, \dots, m$, where x_i denotes the number or proportion of individuals in the i th compartment. Let $F_i(\bar{x})$ be the rate of appearance of new infections in compartment i and let $V_i(\bar{x}) = V_i^-(\bar{x}) - V_i^+(\bar{x})$, where V_i^+ is the rate of transfer of individuals into compartment i by all other means and V_i^- is the rate of transfer of individuals out of the i th compartment. The difference $F_i(\bar{x}) - V_i(\bar{x})$ gives the rate of change of x_i . Note that F_i should include only infections that are newly arising but does not include terms that describe the transfer of infectious individuals from one infected compartment to another.

Assuming that F_i and V_i meet the conditions outlined by Diekmann *et al.* (1990), Diekmann & Heesterbeek (2000) and van den Driessche & Watmough (2002), we can form the next generation matrix FV^{-1} from matrices of partial derivatives of F_i and V_i .

Here, V^{-1} is the inverse of the matrix V . For a 2×2 matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

the inverse is

$$A^{-1} = \frac{1}{\det A} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix},$$

so long as $\det A \neq 0$, obviously. Note that this is only true for 2×2 matrices; also, remember that $\det A = ad - bc$.

We have

$$F = \left[\frac{\partial F_i(\bar{x})}{\partial x_j} \right] \quad \text{and} \quad V = \left[\frac{\partial V_i(\bar{x})}{\partial x_j} \right],$$

where $i, j = 1, \dots, m$, and where \bar{x} is the disease-free equilibrium. The entries of FV^{-1} give the rate at which infected individuals in x_j produce new infections in x_i , times the average length of time an individual spends in a single visit to compartment j . R_0 is given by the largest eigenvalue of the matrix FV^{-1} .

For example, consider a model of malaria. Let us describe the rate of change of the infected human, H_I , and mosquito, M_I , populations by the following equations:

$$\begin{aligned} \dot{H}_I &= \beta_{MH} M_I H_S - (\mu_H + \alpha + \sigma) H_I \\ \dot{M}_I &= \beta_{HM} M_S H_I - \mu_M M_I. \end{aligned}$$

Infected humans are produced by the infection of susceptible humans, H_S , by an infected mosquito with efficacy β_{MH} . We assume that they die with natural death rate μ_H , die due to infection with rate σ and recover from the infection with rate α . Infected mosquitos are produced when susceptible mosquitos, M_S , bite infected humans. We assume that this process has efficacy β_{HM} and assume that infected mosquitos can only leave the infected compartment by dying naturally with rate μ_M . For this system, we find that

$$F = \begin{pmatrix} 0 & \beta_{MH} H_S(0) \\ \beta_{HM} M_S(0) & 0 \end{pmatrix} \quad \text{and} \quad V = \begin{pmatrix} \mu_H + \alpha + \sigma & 0 \\ 0 & \mu_M \end{pmatrix}.$$

Since $\det V \neq 0$, we can determine V^{-1} :

$$\begin{aligned} V^{-1} &= \frac{1}{\mu_M(\mu_H + \alpha + \sigma)} \begin{pmatrix} \mu_M & 0 \\ 0 & \mu_H + \alpha + \sigma \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{\mu_H + \alpha + \sigma} & 0 \\ 0 & \frac{1}{\mu_M} \end{pmatrix}. \end{aligned}$$

We then have

$$\begin{aligned}
FV^{-1} &= \begin{pmatrix} 0 & \beta_{MH}H_S(0) \\ \beta_{HM}M_S(0) & 0 \end{pmatrix} \begin{pmatrix} \frac{1}{\mu_H + \alpha + \sigma} & 0 \\ 0 & \frac{1}{\mu_M} \end{pmatrix} \\
&= \begin{pmatrix} 0 & \frac{\beta_{MH}H_S(0)}{\mu_M} \\ \frac{\beta_{HM}M_S(0)}{\mu_H + \alpha + \sigma} & 0 \end{pmatrix} \\
\det(FV^{-1} - \lambda I) &= \det \begin{pmatrix} -\lambda & \frac{\beta_{MH}H_S(0)}{\mu_M} \\ \frac{\beta_{HM}M_S(0)}{\mu_H + \alpha + \sigma} & -\lambda \end{pmatrix} \\
&= \lambda^2 - \frac{\beta_{MH}\beta_{HM}H_S(0)M_S(0)}{\mu_M(\mu_H + \alpha + \sigma)}.
\end{aligned}$$

Thus

$$R_{0,M} = \sqrt{\frac{\beta_{MH}\beta_{HM}H_S(0)M_S(0)}{(\mu_H + \alpha + \sigma)\mu_M}}.$$

For comparison, we also compute the value of R_0 for this system using the survival function method:

$$\begin{aligned}
R_{0,S} &= \frac{\beta_{MH}\beta_{HM}H_S(0)M_S(0)}{(\mu_H + \alpha + \sigma)\mu_M} \\
&= (R_{0,M})^2.
\end{aligned}$$

The difference here is a matter of definition: the survival function gives the total number of infectives *in the same class* produced by a single infective of that class, while the next-generation matrix gives the mean number of new infectives per infective in any class, *per generation*. Values corresponding to the latter definition thus depend on the number of infective classes in the infection cycle.

For example, suppose you have a mosquito-borne disease where humans infect two mosquitos, while mosquitos infect three humans. For convenience, label these $R_H = 2$ and $R_M = 3$. Then the number of humans infected from a primary human (via mosquitos) is $R_0 = 2 \times 3 = 6$. See Figure 6.1.

However, the next-generation method would calculate $R_{0,N} = \sqrt{6}$, which is a weighted average ($2 < \sqrt{6} < 3$) of the number of infectives each individual produces in the next infection event.

Uses in Policy: Wonham et al. (2004) derived a system of ODEs to describe the behaviour of West Nile virus. Their model consisted of susceptible, infectious, recovered and dead birds, and larval, susceptible, exposed and infectious mosquitos. The next-generation method was used to calculate R_0 from this model, in order to evaluate the ability of the virus to invade the system. The calculated value of R_0 was then interpreted biologically as the square root of the product of (i) the disease R_0 from mosquitos to birds and (ii) the R_0 from birds to mosquitos. Each of these R_0 values was further analysed as a product of disease transmission and infectious lifespan in case (i) and the product of the transmissibility, the number of initially susceptible mosquitos per bird

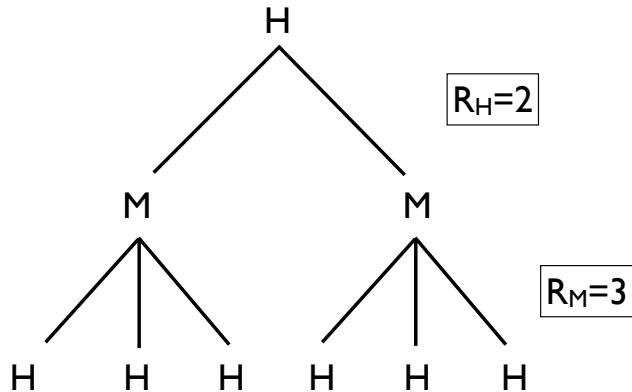


Fig. 6.1. A example of a two-stage basic reproductive ratio.

that survive the exposed period and the infectious lifespan of birds in case (ii). R_0 was then used to establish a threshold mosquito level, above which the virus will invade a constant population of susceptible mosquitos.

The R_0 value derived was then used to evaluate public-health policy markers. Two such policies were evaluated: mosquito control and bird control. It was demonstrated that a small increase in mosquito mortality can lead to a disproportionately large increase in the outbreak threshold. More surprisingly, however, R_0 was used to show that reducing crow densities would have the opposite effect and actually enhance disease transmission (unless extremely low densities limited mosquito biting rates). Thus R_0 was used to show that reducing the initial mosquito population below the calculated threshold would have prevented the West Nile outbreak for New York in 2000. Conversely, bird control would have had the opposite effect.

It should also be noted that, in the dengue example given in the previous section, multiple classes of infectives were defined in the model, but the definition of R_0 used was the number of infected humans per infected human, not the square root of this value as would be obtained by the next-generation operator.

6.4.3 Average age at infection

A related approach, also based on the endemic equilibrium, is that R_0 can be estimated as L/A , where L is the mean lifetime and A is the mean age of acquiring the disease. In brief, we must assume that all individuals are born susceptible, that after acquiring the disease they are no longer susceptible, that the population is at the endemic equilibrium (ie $R_0 > 1$) and that homogenous mixing, particularly among age groups, occurs. While this strong

set of assumptions might never be fully realized in a practical setting, the usefulness of this approach is clear since both L and A are readily measured.

Uses in Policy: This method has been used to calculate pathogens in canines (Laurenson *et al.* 1998).

6.4.4 The final-size equation

The final-size equation is applicable to closed populations only, where the infection leads either to immunity or death. In this situation, the number of susceptibles can only decrease and the final fraction of susceptibles, $s(\infty)$, can be used to estimate R_0 :

$$R_0 = \frac{\ln s(\infty)}{s(\infty) - 1}.$$

This estimate holds when the disease itself does not interfere with the contact process, or when contact intensity is proportional to population density.

Uses in Policy: Using this method, the reproductive number for SARS was estimated in the absence of interventions and in the presence of control efforts in order to determine the effectiveness of public health efforts to reduce transmission (Lipsitch *et al.*, 2003)

6.4.5 Calculation from the intrinsic growth rate

Finally, R_0 may be determined from the intrinsic growth rate of the infected population. This growth rate, often denoted r_0 , is the rate at which the total number of infectives, I , grows in a susceptible population, such that $dI/dt = r_0 I$. Note that this is an *implicit* definition of r_0 ; thus, from a modelling perspective, using r_0 is seldom elegant.

In the simplest possible models, when infectivity is constant throughout the infectious period, R_0 can be estimated as $1 + r_0 L$, where L is the expected duration of the infectious period. (The “one” is necessary in this expression because R_0 reflects the total number of new infections, whereas the overall growth rate r_0 includes the death of the founding individual.) For more complex models, the relation between r_0 and R_0 can be derived by expressing both in terms of the model parameters, exploiting that fact that the largest eigenvalue of the Jacobian, evaluated at the disease-free equilibrium, gives r_0 .

We also note that r_0 itself can be used as a threshold parameter, since $R_0 < 1$ implies $r_0 < 0$. Thus the condition $r_0 < 0$ is actually equivalent to the Jacobian method described in Section 6.3. This method proves useful since r_0 can be readily estimated from incidence data in epidemiology or from viral-load data, for in-host models.

Uses in Policy: Pybus *et al.* (2001) used this method to evaluate differences in epidemic behaviour among Hepatitis C subtypes, based on gene-sequence data.

6.5 Vectorial capacity

For transmission of diseases by blood-sucking vectors, a related concept is the *Vectorial Capacity* (VC), which you may be familiar with. The vectorial capacity is defined as the rate (usually daily) at which a bloodsucking insect population generates new inoculations from a currently infectious case. Vectorial capacity is a measure of potential, rather than actual, rate of transmission, because it includes no parasitological information.

The equation is

$$VC = \frac{ma^2p^n}{-\ln(p)},$$

where m is the ratio of vectors to host (N/H), a is the daily biting rate by the vector on host species, p is the daily survival rate probability for the vector and n is the extrinsic incubation period of parasite in the vector.

When we add information about the competence of the vector and the rate of recovery of the vertebrate host, we can express the basic reproductive ratio in a simple vector transmission model as

$$R_0 = \frac{ma^2p^nbc}{-r \ln(p)},$$

where b is the fraction of mosquitos that infect a vertebrate host when biting (susceptibility of humans), c is the fraction of susceptible mosquitos that become infected when biting (susceptibility of mosquitos) and r is the rate of vertebrate recovery from infectiousness ($1/r$ is the average duration of infection).

This equation has many simplifying assumptions, including constant population densities and age structures, lack of immunity, lack of parasite-induced mortality, exclusion of the latency period and heterogeneity in contact rates between vectors and hosts (non-random biting rates).

The vectorial capacity and other related measures of the potential rate of transmission are particularly advantageous under two circumstances:

1. When the fraction of vectors that are infected is too small to measure reliably; this is typical in epidemic (rather than endemic) areas, where infection rates are often of the order of 0.1% and impossibly large sample sizes (e.g., more mosquitos than can be dissected) are required to detect significant changes in such small rates;
2. When an explanation for the variation in infection rate is needed; in this context, the theory of vectorial capacity clearly explains why the transmission rate is particularly sensitive to changes in the daily survival rate of vectors. The concept of vectorial capacity does not apply to snails as intermediate hosts (e.g., schistosomiasis), though snail abundance and longevity are important epidemiological variables.

Uses in Policy: Hagmann *et al.* (2003) used the low value of R_0 , estimated from the vectorial capacity, to justify the possibility of eliminating malaria from an island in the Gulf of Guinea.

We will discuss vector-borne diseases in more detail in the next chapter.

6.6 Lab work

All these examples are theoretical and do not require Matlab.

6.6.1 Exercises

1. (Next-generation method) Consider an SEIR model

$$\begin{aligned}\dot{S} &= \Omega - \beta SI - \mu S \\ \dot{E} &= \beta SI - (\mu + k)E \\ \dot{I} &= kE - (\gamma + \mu)I \\ \dot{R} &= \gamma I - \mu R,\end{aligned}\tag{6.3}$$

where Ω is the birth rate, μ is the per capita natural death rate, β is the efficacy of infection of susceptible individuals S , k is the rate at which a latent individual becomes infectious and γ is the per capita recovery rate.

- a) Show that

$$F = \begin{pmatrix} 0 & \beta\Omega/\mu \\ 0 & 0 \end{pmatrix}.$$

Hint: only two compartments are infected.

- b) Show that

$$V = \begin{pmatrix} \mu + k & 0 \\ -k & \gamma + \mu \end{pmatrix}.$$

- c) Find V^{-1} .

- d) Finally, show that

$$R_{0,N} = \frac{k\beta\Omega}{(\mu + k)(\mu + \gamma)\mu}.\tag{6.4}$$

(Note that this is also the value of R_0 determined by the survival function method.)

2. (Jacobian) Calculate the Jacobian matrix for model (6.3) and find the eigenvalues for the disease-free equilibrium. Does this match the R_0 value found here? Which method do you like better?
3. (Endemic equilibrium) Consider this model (from Blower *et al.*, 1998) of herpes simplex virus:

$$\begin{aligned}\frac{dX}{dt} &= \pi - Xc\beta \frac{H}{N} - X\mu \\ \frac{dQ}{dt} &= H(\sigma + q) - Q(\mu + r) \\ \frac{dH}{dt} &= Xc\beta \frac{H}{N} - H(\mu + \sigma + q) + rQ,\end{aligned}$$

where X is the susceptible population, Q represents those infected with the virus in the non-infectious latent state, H represents those infected with the virus in infectious state and $N = X + Q + H$. (Other letters are positive parameters.)

- a) Show that, at equilibrium,

$$\begin{aligned}\bar{N} &= \frac{\pi}{\mu} \\ \bar{X} &= \frac{\pi}{\mu} - \frac{\mu + \sigma + q + r}{\mu + r} \bar{H} \\ \bar{Q} &= \frac{\sigma + q}{\mu + r} \bar{H},\end{aligned}$$

where \bar{H} is yet to be determined.

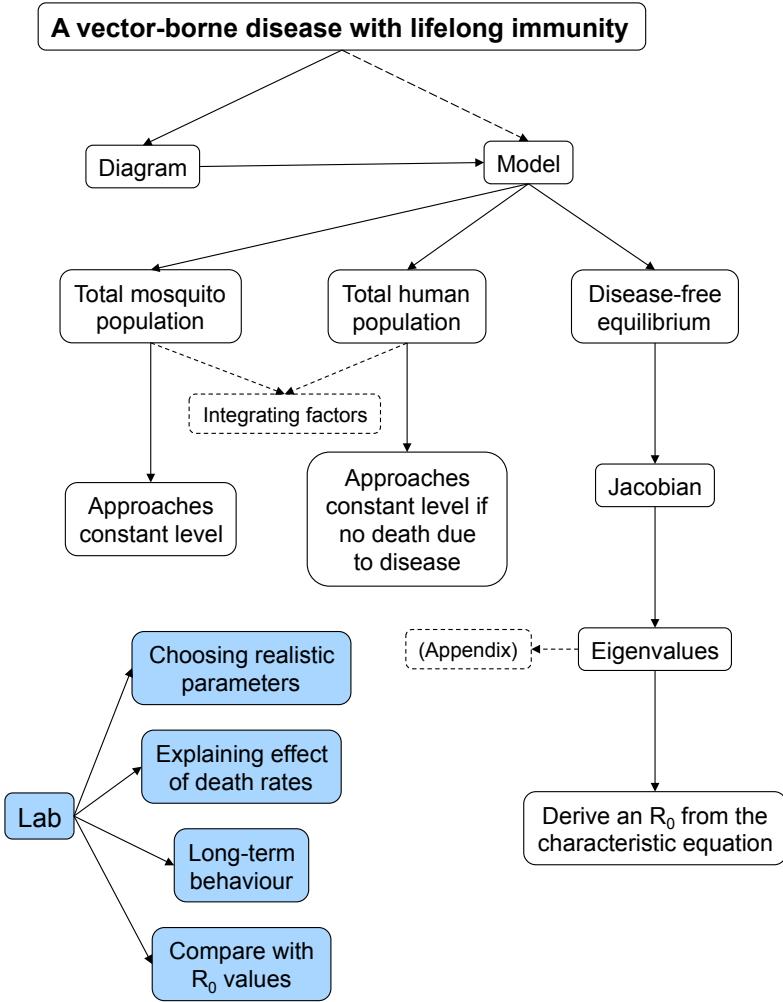
- b) Find the disease-free equilibrium.
- c) Show that, if $\bar{H} \neq 0$, then

$$\bar{H} = \frac{\pi}{\mu} \left[\frac{\mu + r}{\mu + \sigma + q + r} - \frac{\mu}{c\beta} \right].$$

- d) Show that the endemic equilibrium only exists when

$$R_{0,E} \equiv c\beta \left(\frac{r + \mu}{\mu(r + \mu + \sigma + q)} \right) > 1$$

and does not exist if the reverse inequality holds.



A vector-borne disease with lifelong immunity

As we saw previously, ODE models can vary from the simple to the complicated. Thus far, we have been considering the spread of the disease itself as an implicit parameter. In an SIR-type model, people become infected as a result of interacting with already infectious people, but we did not consider the mechanism of that infection. (Not that we needed to.)

For vector-host systems, the disease is spread by the vector. If you stand near someone who has a mosquito-borne disease, you're not at direct risk, unless a mosquito transmits it from them to you. That is, a susceptible mosquito must first take a bloodmeal from an infected human, so that it becomes infected. The now-infected mosquito must then bite a susceptible human, converting them to an infected human. Given the time lags involved, spatial distance between infected humans is not a good measure of the probability of infection.

This chapter is structurally identical to the methods from the previous two chapters but includes more complex dynamics, so the matrices will be bigger and the math more complicated. By the end of this chapter you should be able to see how to build a vector-host model.

7.1 Building a vector-borne disease model

For the moment, we'll consider a vector-borne disease that confers lifelong immunity, such as yellow fever. Thus, susceptible individuals can get infected, whereupon they may die due to the disease, but if they survive they are no longer susceptible. This is, of course, an SIR model, as we saw in previous chapters. However, what differentiates the model in this chapter is that we must also include the dynamics of the mosquito infection into the model.

We'll assume that the parasite has no effect on the mosquito vector. This may or may not be true, but, given their short lifespan, this seems like a reasonable assumption. Since the mosquito lifespan is short, we really need to include the birth rate and the death rate (or else we'd use up all the available

mosquitos quite quickly). We'll also add these in for humans. For humans, there's an additional death rate while infected — you can die from the disease — but for mosquitos we can assume that infection doesn't have a significant impact on their lifespan. Again, this may or may not be true, but isn't an unreasonable assumption.

What does all this mean? It means we have an SIR model for humans and an SI model for mosquitos... but the two aren't independent. We need to consider the dynamics for how a human infects a mosquito and how a mosquito infects a human. A mosquito becomes infected if a susceptible mosquito meets an infected human. So our infection term will actually be a *cross-infection* term, involving both species. The model is shown in Figure 7.1.

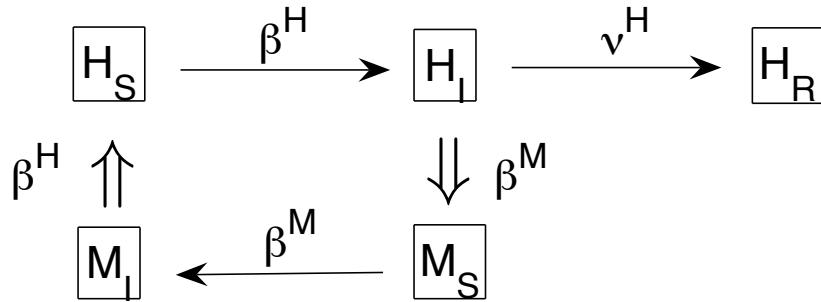


Fig. 7.1. A vector-host model for a disease with permanent immunity.

The differential equations are thus

$$\begin{aligned}
 \frac{dM_S}{dt} &= \lambda^M - \beta^M H_I M_S - \mu^M M_S \\
 \frac{dM_I}{dt} &= \beta^M H_I M_S - \mu^M M_I \\
 \frac{dH_S}{dt} &= \lambda^H - \beta^H M_I H_S - \mu^H H_S \\
 \frac{dH_I}{dt} &= \beta^H M_I H_S - \mu^H H_I - \gamma^H H_I - \nu^H H_I \\
 \frac{dH_R}{dt} &= \nu^H H_I - \mu^H H_R,
 \end{aligned} \tag{7.1}$$

where λ^i is the birth rate ($i = M, H$, for mosquitos and humans respectively), β^i is the rate of infection, μ^i is the background death rate, γ^H is the death rate due to the disease and ν^H is the rate of recovery.

Note: the units of any differential equation must be the same on both sides. The first two equations have units of [mosquitos]/[time]. Thus, each term must ultimately have these units. So λ^M has units [mosquitos]/[time] (since it stands alone) and μ^M has units 1/[time] (since M_S obviously has units [mosquitos]). Thus the term $\beta^M H_I M_S$ must have overall units of [mosquitos]/[time]. Since H_I has units of [humans] and M_S has units of [mosquitos], this means

$$[\beta^M][\text{humans}][\text{mosquitos}] = [\text{mosquitos}]/[\text{time}]$$

$$[\beta^M] = 1/([\text{humans}][\text{time}]).$$

Similarly, β^H has units 1/([mosquitos][time]).

The disease-free equilibrium occurs when there is no disease. Well, obviously. So that means $\bar{M}_I = \bar{H}_I = 0$. Since the derivatives are zero at equilibrium (or else it wouldn't be an equilibrium), this means $\bar{H}_R = 0$, $\bar{M}_S = \frac{\lambda^M}{\mu^M}$ and $\bar{H}_S = \frac{\lambda^H}{\mu^H}$. So we've found the disease-free equilibrium.

7.2 Finding eigenvalues

Calculating eigenvalues involves the same steps as used previously (namely, first find the disease-free equilibrium, calculate the Jacobian matrix and evaluate the determinant of the Jacobian at the disease-free equilibrium). However, this time around, we have a 5×5 matrix to deal with, which isn't so simple.

From Appendix E, the determinant of the Jacobian matrix at the disease-free equilibrium leads to the characteristic equation

$$-(\mu^M + \Lambda)(\mu^H + \Lambda)^2 [(\mu^M + \Lambda)(\mu^H + \gamma^H + \nu^H + \Lambda) - \beta^M \beta^H \bar{M}_S \bar{H}_S] = 0.$$

(We use Λ here, since λ is already taken.)

The first three eigenvalues are always negative ($\Lambda = -\mu^M, -\mu^H$ and $-\mu^H$ respectively), so they do not contribute to questions of stability of the disease-free equilibrium. We can rewrite the part in the square brackets as

$$\Lambda^2 + (\mu^M + \mu^H + \gamma^H + \nu^H)\Lambda + \mu^M(\mu^H + \gamma^H + \nu^H) - \beta^M \beta^H \bar{M}_S \bar{H}_S = 0. \quad (7.2)$$

This is a quadratic, in the form $\Lambda^2 + b\Lambda + c$ with $b > 0$. The sign of c will determine where the parabola crosses the y -axis, but this means the sign of c will also determine the stability. Why is this? Figure 7.2 provides the answer.

The parabola must be pointing upwards since the coefficient of the Λ^2 term is positive. The vertex (the point where the parabola turns around) is to the left of the y -axis, since $b > 0$ ¹.

¹ If $f(\Lambda) = \Lambda^2 + b\Lambda + c$ then $f'(\Lambda) = 2\Lambda + b$. Equating this to zero, gives $\Lambda = -b/2$. When the first derivative is zero, we have a turning point, which is the vertex of the parabola. So if $b > 0$, the vertex is negative; i.e., to the left of the y -axis.

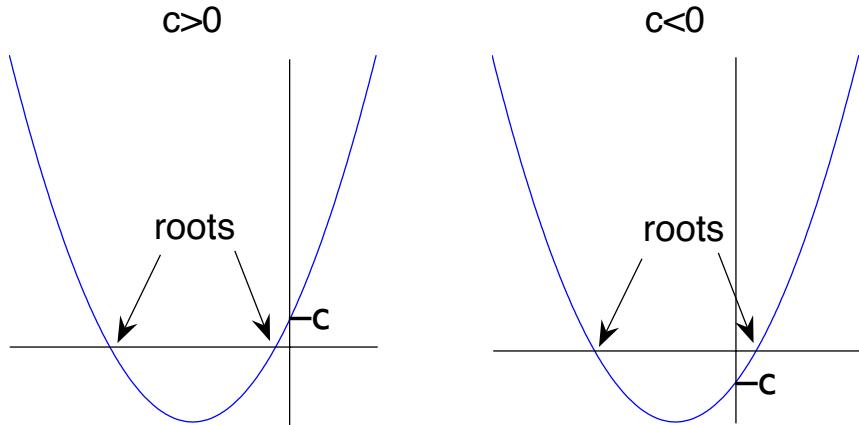


Fig. 7.2. The roots of a quadratic.

The only undecided part is c , the y -intercept. If $c > 0$, then we have the first case of Figure 7.2 and we see that both roots are negative. Hence, since the other three eigenvalues were also negative, the system will be stable (since all roots are negative). If $c < 0$, however, then we have the second case of Figure 7.2 and we see that there is a positive root. It follows that the system would be unstable in this case, since any positive root means instability.

7.3 Deriving a threshold condition

What we'd really like to do now is to derive an R_0 -like threshold from equation (7.2). This is possible, but it involves using the quadratic formula to extract the actual roots, then rearranging the appropriate root in order to derive an R_0 . Which is messy, but just about doable in this case. However, what if we had a cubic equation? Or a quartic?

Another alternative we explored earlier was to solve for the endemic equilibrium and derive an R_0 -like threshold condition from that. Again, this is possible, but a) messy and b) not really tractable in general. Note that, if we did, this R_0 -like condition would almost certainly be different from the one mentioned in the previous paragraph.

There's a third possibility, however. We already have a threshold condition: namely, whether $c < 0$ or $c > 0$. We can rearrange this condition to derive yet another R_0 -like condition. This one is definitely not the same condition we would derive from finding the roots of equation (7.2), or from finding the endemic equilibrium. However, since we don't know if either of those is the true R_0 anyway, we might as well stick to a threshold condition.

Since $c > 0$ implies stability, we have

$$\begin{aligned}
c &= \mu^M(\mu^H + \gamma^H + \nu^H) - \beta^M \beta^H \bar{M}_S \bar{H}_S > 0 \\
\mu^M(\mu^H + \gamma^H + \nu^H) &> \beta^M \beta^H \bar{M}_S \bar{H}_S \\
R_0^c &= \frac{\beta^M \beta^H \bar{M}_S \bar{H}_S}{\mu^M(\mu^H + \gamma^H + \nu^H)} \\
&= \frac{\beta^M \beta^H \lambda^M \lambda^H}{(\mu^M)^2 \mu^H (\mu^H + \gamma^H + \nu^H)},
\end{aligned}$$

since $R_0^c < 1$ when $c > 0$.

This threshold condition tells us important information about how to control the disease. For example, the $(\mu^M)^2$ term in the denominator suggests that enhancement of vector mortality is highly conducive to eradication. That is, the average lifespan of a mosquito is $1/\mu^M$, so if this is small, then $1/(\mu^M)^2$ is even smaller. This helps to lower R_0^c considerably.

We could also see the effects on R_0^c if we had disease-modifying drugs that changed γ^H or ν^H , or sprayed an insecticide that changed μ^M or a larvacide that altered λ^M . It's possible to estimate the effect that such control measures would have before implementing them. Always bearing in mind that our mathematical model is an approximation of the real situation.

7.3.1 The Routh–Hurwitz criterion

What if we have a higher-order characteristic equation? Could we use the same idea? Yes... but with qualification. It's not sufficient just to have positive coefficients. Fortunately, there's a thing called the Routh–Hurwitz criterion to help out here, which is okay for three dimensions but gets complicated beyond that.

Consider the polynomial

$$a_3 \lambda^3 + a_2 \lambda^2 + a_1 \lambda + a_0 = 0. \quad (7.3)$$

Under what circumstances will the roots of (7.3) have negative real part? There are two conditions to ensure this:

1. $a_n > 0$ for all n
2. $a_1 a_2 > a_3 a_0$

For characteristic equations, the leading term is always $a_3 = 1$. And we can use the condition $a_0 = 0$ as our threshold criterion. That means we need to check the signs of a_2 , a_1 and $a_1 a_2 - a_0$. In fact, since we're only concerned with local stability, we can do this with the additional constraint that $a_0 = 0$. That is, at the threshold, one of the roots is zero, so we'd like to check whether the other roots all have negative real part. If so, then the zero root determines stability. We can then rearrange $a_0 = 0$ in the form $R_0 = 1$ and proceed as usual.

For higher-order polynomials, the requirements are more complicated. So in disease modelling, the Routh–Hurwitz criterion is generally only used for third-order polynomials.

7.4 Total population behaviour

Sometimes we can use the symmetry of the equations to discover information about the total population behaviour. For example, if we add the mosquito equations together (the first two equations in model (7.1)), then we have

$$\begin{aligned}\frac{dM_S}{dt} + \frac{dM_I}{dt} &= \lambda^M - \mu^M M_S - \mu^M M_I \\ \frac{d}{dt}(M_S + M_I) &= \lambda^M - \mu^M(M_S + M_I) \\ \frac{dM}{dt} &= \lambda^M - \mu^M M,\end{aligned}$$

where $M = M_S + M_I$, obviously.

What we have here is a single differential equation in one variable (M). But, better than that, this differential equation is linear. So we can use an integrating factor (see Appendix F) to solve it. Once again, don't sweat the details if it looks too heavy for you, but the take-home message is that we can solve a simple (i.e., one-dimensional) equation like this to find the long-term behaviour of the *total* population:

$$\lim_{t \rightarrow \infty} M(t) = \bar{M}_S = \frac{\lambda^M}{\mu^M}.$$

Thus, the total mosquito population approaches a constant level. In particular, the mosquitos never die out on their own (although, depending on the ratio of their birth to death rates, the ultimate mosquito density may be either high or low). Of course, this only tells us about the *total* mosquito population, not the distribution of the disease. For all we know there could be no disease at all, all mosquitos could be infected, or (more likely) it could be some combination of the two.

What this says ecologically is that the overall population is governed by its birth and death rates. The only way into the mosquito population is to be born and the only way out is to die (of natural causes). Everything else is shuffling about between different states (susceptible or infected).

Can we do the same trick with the humans? Answer: not quite. If we add the final three equations together, we get

$$\frac{dH_S}{dt} + \frac{dH_I}{dt} + \frac{dH_R}{dt} = \lambda^H - \mu^H H_S - \mu^H H_I - \gamma^H H_I - \mu^H H_R.$$

The presence of the $\gamma^H H_I$ term means we can't do the same trick (since it's only present for the infected class, not the other two classes). But let's suppose for the moment that $\gamma^H = 0$. That is, there is no death rate due to the disease. This isn't totally unlikely, as some diseases are milder than others. In this case,

$$\frac{d}{dt} (H_S + H_I + H_R) = \lambda^H - \mu^H (H_S + H_I + H_R).$$

Using the same method as before, with integrating factor $e^{\mu^H t}$, we find that

$$H_S + H_I + H_R \rightarrow \bar{H}_S = \frac{\lambda^H}{\mu^H}$$

as $t \rightarrow \infty$. (You can work through the details of this yourself if you're feeling brave enough, as they're almost exactly the same as the mosquito equations we just did.) This tells us about the total population levels for the humans, in the case $\gamma^H = 0$, although again it doesn't tell us how the different classes are represented within this total.

What about the case $\gamma^H \neq 0$, though? We don't have an obvious analytic route to go down (you may be pleased to hear this!), so we'll use Matlab to explore the effects of nonzero γ^H s in the lab.

7.5 Lab work

The problem

Suppose we have a village of 5000 uninfected humans. There are 1000 susceptible mosquitos and ten mosquitos infected with yellow fever. How will these infected mosquitos affect the human population?

- Come up with some reasonable parameters for the birth and death rates and the recovery rate. Consider three cases for the death rate: 1) $\gamma^H = 0$ (no death) 2) $\gamma^H = 0.1$ 3) $\gamma^H = 1$.
- Plot the timecourse of the disease, for both mosquitos and humans, using the different death rates, over 24 hours, a week and 100 days.
- Show a phase-plane analysis of the number of infected mosquitos versus the number of infected humans.

The solution

We want to simulate our model from the previous section, exploring the effects of changing γ^H . We need some parameters to play with. Note that we're not slavishly trying to get accurate parameters here, just a guide to see what sort of behaviour will occur. We can always go back later and change the parameter values if we have better information. But let's think things through, to come up with some semi-reasonable parameter estimates.

The first thing that comes to mind is that mosquitos live for only a matter of weeks, so our timescale should reflect that accordingly. Let's suppose that mosquitos are actively biting for 15 days. We'll pick an average survival time of 60 years for humans. However, these two need to have the same units, so we'll convert 60 years into 21900 days. Remember that μ^M and μ^H aren't

survival times, they're death *rates*, with units of $1/[\text{time}]$. Just as frequency is the reciprocal of the period, so too the death rate is the reciprocal of the survival time. Thus

$$\begin{aligned}\mu^M &= \frac{1}{15} \text{ days}^{-1} \\ \mu^H &= \frac{1}{60} \text{ years}^{-1} = \frac{1}{21900} \text{ days}^{-1}.\end{aligned}$$

What about the birth rate? Remember that λ^M and λ^H must have units $[x]/[\text{time}]$, where x is the number of mosquitos or humans. Let's suppose there are ten mosquitos born every day and one human born every ten days. So

$$\begin{aligned}\lambda^M &= 10 \text{ mosquitos} \cdot \text{day}^{-1} \\ \lambda^H &= 0.1 \text{ humans} \cdot \text{day}^{-1}.\end{aligned}$$

We'll assume an infection period of a week, so

$$\nu^H = \frac{1}{7} \text{ days}^{-1}.$$

Finally, we need to determine the transmissibilities. These are complex interactions of the number of biting events, the chances of the disease being passed from one to another, the immunology of humans and mosquitos, and so on. It's important to stress that we have different transmission rates for host-to-vector infection than for vector-to-host infection.

Let's pick some nice round numbers and suppose that the chance of getting infected if an infected mosquito bites you is ten times the chance of the mosquito getting infected if it bites an infected human. Thus

$$\begin{aligned}\beta^M &= 0.01 \text{ humans}^{-1} \cdot \text{days}^{-1} \\ \beta^H &= 0.1 \text{ mosquitos}^{-1} \cdot \text{days}^{-1}.\end{aligned}$$

(The numbers are chosen at random to reflect different transmission rates but ignore lag times.) We need a function file to account for our ODEs:

```
function pdot=vectorf(t,p)
%This is the ODE for the Yellow fever problem
lambdaM=10;
muM=1/15;
betaM=0.01;
lambdaH=0.1;
muH=1/21900;
betaH=0.1;
gammaH=0;
%gammaH=0.1;
```

```
%gammaH=1;
nuH=1./7;
pdot(1,:)=lambdaM-betaM.*p(4).*p(1)-muM.*p(1);
pdot(2,:)=betaM.*p(4).*p(1)-muM.*p(2);
pdot(3,:)=lambdaH-betaH.*p(2).*p(3)-muH.*p(3);
pdot(4,:)=betaH.*p(2).*p(3)-muH.*p(4)-gammaH.*p(4)-nuH.*p(4);
pdot(5,:)=nuH.*p(4)-muH.*p(5);
```

Remember to save this as “vectorf.m”. Next, we need an M-file to state our initial conditions, call the function file and plot the results.

```
%This is a program to use the vector-host ODE vectorf
clear all
t0=0;
tf=1;
%tf=7;
%tf=100;
p0=[1000,10,5000,0,0];
tspan=[t0 tf];
[t,p]=ode23(@vectorf,tspan,p0);
subplot(1,2,1)
plot(t,p(:,1),t,p(:,2))
title('Mosquitos')
subplot(1,2,2)
plot(t,p(:,3),t,p(:,4),t,p(:,5))
title('Humans')
```

You can run these multiple times by taking the % signs out.

Finally, we need to plot M_I vs H_I . You should be able to figure out how to do that by now (don’t forget to put mosquitos on the x -axis). Use `help plot` if you need a hint (or check page 58). It should look like Figure 7.3.

7.5.1 Exercises

1. Can you determine the direction of the trajectories in Figure 7.3?
2. Do the trajectories cycle around this curve, or do they settle down eventually?
3. When there is no death rate due to infection (i.e., $\gamma^H = 0$), what eventually happens to the population? Has the disease been eradicated?
4. Use your data to calculate R_0^c , for the case $\gamma^H = 0$. (By now you should be able to write a short Matlab code to do this.)

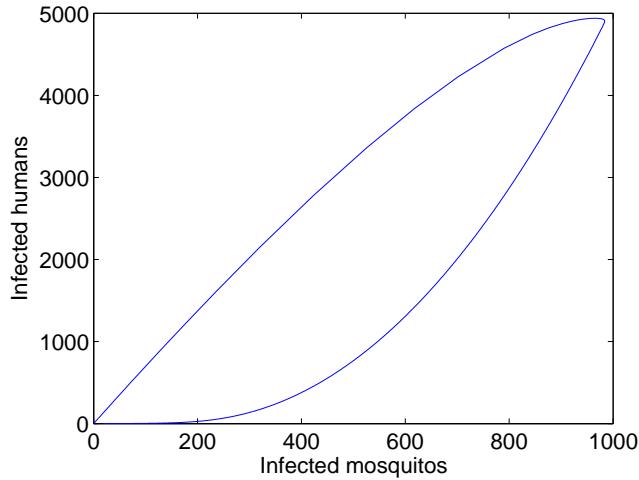
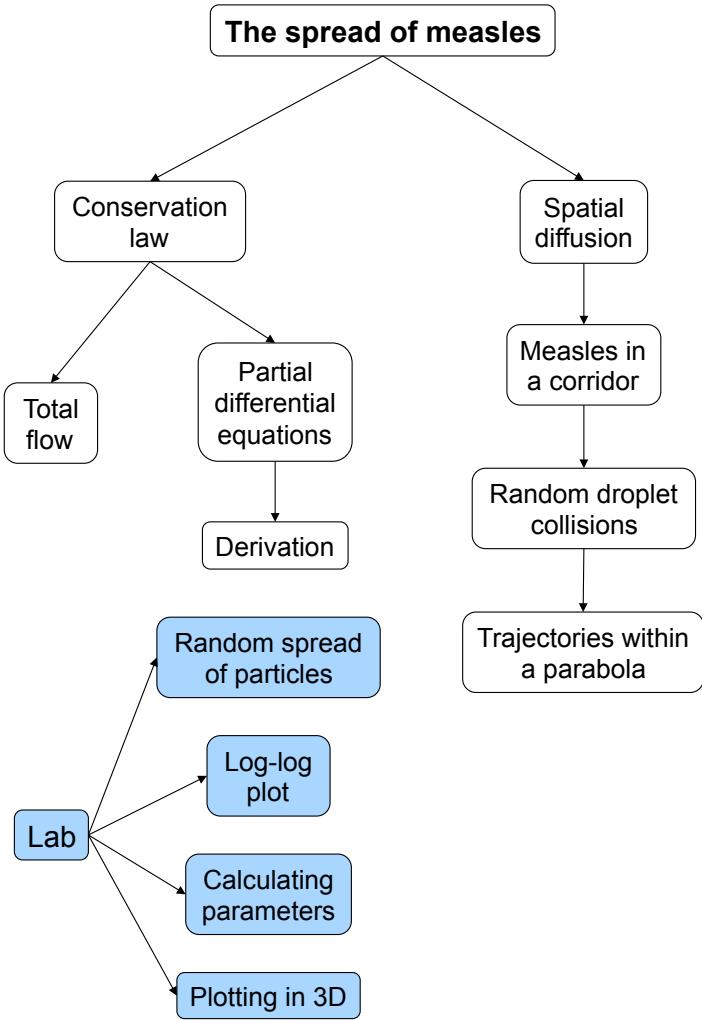


Fig. 7.3. Infected mosquitos vs. infected humans.

5. How does this value compare with your results in this section? Does this knowledge cause you to revise your opinion about eradication?
6. Plot the time series of the infection for $\gamma_H = 0$, $\gamma_H = 0.1$ and $\gamma_H = 1$.
7. Use your time series to find the endemic equilibrium values for $\gamma_H = 0$. Are these what you expected?
8. Name two vector-borne diseases, aside from yellow fever, where immunity is lifelong.
9. What would be needed for the disease to be maintained in humans? What about recurrence of the disease after it has died out?



The spread of measles

When dealing with ODEs previously, we have only been interested in systems that change with respect to one variable, usually time. However, many systems are much more complicated, changing with respect to both space and time. If a disease is introduced into a community, the prevalence at some location will depend both on the location relative to where the initial infection occurred and on the time elapsed since the infection was introduced. The infection at that location changes with space and time.

When dealing with changes in more than one variable, we turn to partial differential equations (PDEs). PDEs allow us to identify the different variables that a system may be dependent on and set up our model accordingly. Naturally, these can get quite complicated. To keep things simple, we'll concentrate on one-dimensional spatial problems — although these are really two dimensional problems, because they also involve time.

By the end of this chapter, you should be able to see how PDEs can tell us about movements through space and time.

8.1 The Conservation Law

The law of Conservation of Mass says that mass can neither be created nor destroyed. One of the implications of this law is that in modelling the spread of molecules (like infected measles droplets, as we'll see shortly), we must create models which do not allow for the spontaneous creation or disappearance of any of our droplets. The mathematical form of this law can be used as the first step in creating models to account for the movement of such droplets.

Consider a stretch of space Δx long (labelled from x to $x + \Delta x$) along which infected measles droplets are flowing (we restrict ourselves to one spatial dimension here, because the math gets complicated; remember that we already have two dimensions, since we have both a spatial dimension and time).

We will represent the density of infected measles droplets at a particular time and space as $U(x, t)$, with units of droplets/length. The flow of the

droplets F (in units of droplets/time) will be related in some manner to the density of the droplets. That is, $F = F[U(x, t)]$. Figure 8.1 illustrates this system, with the inward and outward flow.

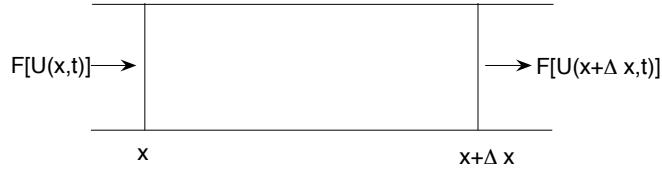


Fig. 8.1. The total flow of droplets in a segment of space of length Δx .

From Figure 8.1, we can see that the total flow on this interval is the inward flow minus the outward flow, or $F[U(x, t)] - F[U(x + \Delta x, t)]$. The total flow represents the change in total number of droplets with respect to time. There is another way to calculate the total flow, however. Reconsider Figure 8.1, looking at a small length of space dx wide, as in Figure 8.2.

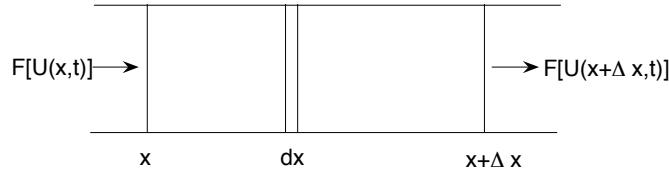


Fig. 8.2. The flow of droplets within a small length of space.

Just as the mass of an object is calculated by summing the density of small sections of it, the total number of droplets in this stretch of space can be calculated by summing the density of these small portions over the total interval. A near-infinite sum of tiny partitions can be approximated by an integral over the entire interval from x to $x + \Delta x$. (Everyone always forgets that an integral is just an infinite sum of tiny pieces.)

Thus, the number of droplets is

$$\int_x^{x+\Delta x} U(x, t) dx.$$

If this is the total number of droplets, then the total flow is the change in this number with respect to time. Thus, the total flow is

$$\frac{d}{dt} \int_x^{x+\Delta x} U(x, t) dx.$$

Equating our two expressions for total flow gives us

$$\frac{d}{dt} \int_x^{x+\Delta x} U(x, t) dx = F[U(x, t)] - F[U(x + \Delta x, t)].$$

We can bring the $\frac{d}{dt}$ term into the integral on the left, but if we do so it will become a *partial derivative*, namely $\frac{\partial}{\partial t}$, since U changes with respect to x as well as t . Partial derivatives are used when the outcome depends on several variables but only one of these is differentiated.

We'll also write the right-hand side in a slightly different manner:

$$\int_x^{x+\Delta x} \frac{\partial U(x, t)}{\partial t} dx = -\{F[U(x + \Delta x, t)] - F[U(x, t)]\}. \quad (8.1)$$

(All we did to the right-hand side was swap the negatives around.)

How can we simplify the integral? For simplicity, let's let $f(x) = \frac{\partial U(x, t)}{\partial t}$. When taking the integral of $f(x)$ from a to $a + \Delta a$, we are calculating the area under the integral from a to $a + \Delta a$. If Δa is extremely small, this area can be approximated by a rectangle with width Δa and height $f(a)$. See Figure 8.3. Thus if we let $\Delta a \rightarrow 0$, we can make the following approximation:

$$\int_a^{a+\Delta a} f(u) du \approx f(a) \Delta a. \quad (8.2)$$

If we apply (8.2) to Equation (8.1), with $f(x) = \frac{\partial U(x, t)}{\partial t}$, we have

$$\frac{\partial U(x, t)}{\partial t} \Delta x \approx -\{F[U(x + \Delta x, t)] - F[U(x, t)]\}.$$

Divide both sides by Δx :

$$\frac{\partial U(x, t)}{\partial t} \approx -\frac{\{F[U(x + \Delta x, t)] - F[U(x, t)]\}}{\Delta x}. \quad (8.3)$$

Since we are letting $\Delta x \rightarrow 0$, we can treat Δx as we usually would dx (or dx in this case). We thus have the (partial) derivative of F with respect to x on the right-hand side of this equation. Remember the definition of a partial derivative:

$$\frac{\partial z(u, v)}{\partial u} = \lim_{\Delta u \rightarrow 0} \frac{z(u + \Delta u, v) - z(u, v)}{\Delta u}.$$

So if we take the limit as $\Delta x \rightarrow 0$ in equation (8.3), we'll have the partial derivative $\frac{\partial F[U(x, t)]}{\partial x}$.

We have thus derived the conservation law

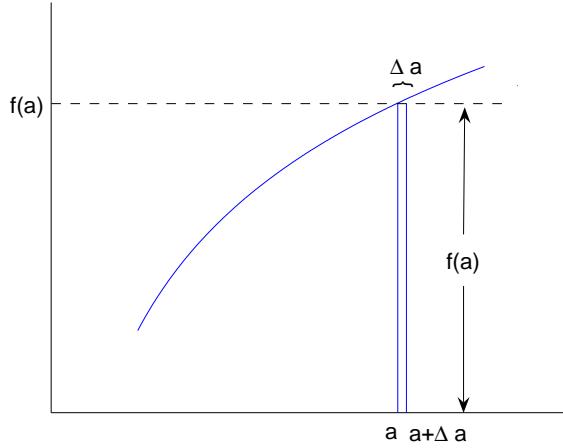


Fig. 8.3. Approximating a thin area by a rectangle. Area $\approx f(a) \times \Delta a$.

$$\frac{\partial U(x, t)}{\partial t} = -\frac{\partial F[U(x, t)]}{\partial x},$$

which is often rewritten as

$$\frac{\partial U}{\partial t} + \frac{\partial F(U)}{\partial x} = 0.$$

This is the conservation equation. Like the conservation law of energy, which says that nothing can be created or destroyed, it says that measles droplets cannot spontaneously appear or disappear; they are all accounted for in the overall flow of droplets.

8.1.1 Diffusion

Diffusion is the term given to “random spreading”, such as that seen when we drop dye into water, as you’ve probably seen in chemistry classes. It involves the movement from high concentration of particles to low concentration. It is a result of Brownian Motion, which is the vibration of the molecules of the medium and the molecules in question. Brownian Motion can be affected by several factors, including temperature.

For diffusion, $F(U) = -D \frac{\partial U}{\partial x}$, where D is the diffusion constant, which depends on the viscosity of the medium. The flow is proportional to the change in density; if the density is decreasing over distance, then $\frac{\partial U}{\partial x}$ is negative. Differentiating, we have

$$\begin{aligned}\frac{\partial F}{\partial x} &= \frac{\partial}{\partial x} \left(-D \frac{\partial U}{\partial x} \right) \\ &= -D \frac{\partial^2 U}{\partial x^2}.\end{aligned}$$

The diffusion equation is then

$$\frac{\partial U}{\partial t} = D \frac{\partial^2 U}{\partial x^2}. \quad (8.4)$$

8.2 Measles in a corridor

Harry is infected with measles. After he sneezes in a corridor at school, the infectious molecules (airborne droplets) spread out from the epicentre (him) and random molecular collisions may knock them left or right. (For simplicity, we're only assuming one spatial dimension, hence the corridor.) His sneeze has some initial velocity, but, after a while, the movement of infected droplets will be largely determined by the way in which they interact with each other, not the initial velocity.

If there are s seconds between collisions and each collision displaces a droplet by $\pm r$, what is the mean distance we can expect the measles to travel in either direction after time t ?

With each random collision, a droplet moves a distance r to the left or right. We can plot the movements of an infectious droplet in Figure 8.4 if we take the horizontal axis to be time and the vertical axis to be space (so $+y$ represents the left direction, and $-y$ represents the right direction).

Assume each collision (left or right) is independent of the previous one and let r_n denote the displacement from 0 at the n th step (positive if to the left and negative if to the right). For the path in Figure 8.4, $r_1 = r$, $r_2 = r$, $r_3 = -r$, $r_4 = r$, $r_5 = -r$, $r_6 = -r$, $r_7 = -r$, $r_8 = r$, $r_9 = -r$, $r_{10} = -r$, $r_{11} = r$. The total displacement (in the y direction) after n collisions is

$$y_n = r_1 + r_2 + r_3 + \cdots + r_n.$$

So in Figure 8.4, $y_1 = r$, $y_2 = 2r$, $y_3 = r$, $y_4 = 2r$, $y_5 = r$, $y_6 = 0$, $y_7 = -r$, $y_8 = 0$, $y_9 = -r$, $y_{10} = -2r$, $y_{11} = -r$. On average, y_n will be zero, since the law of probability states that the collisions will likely send the droplets left as often as right.

We want to look at the mean distance, without the individual components cancelling each other out, so we'll look at the sum of squares. If we square y_n , we get

$$\begin{aligned}y_n^2 &= (r_1 + r_2 + r_3 + \cdots + r_n)^2 \\ &= r_1^2 + r_2^2 + r_3^2 + \cdots + r_n^2 + 2(r_1 r_2 + r_1 r_3 + \cdots).\end{aligned}$$

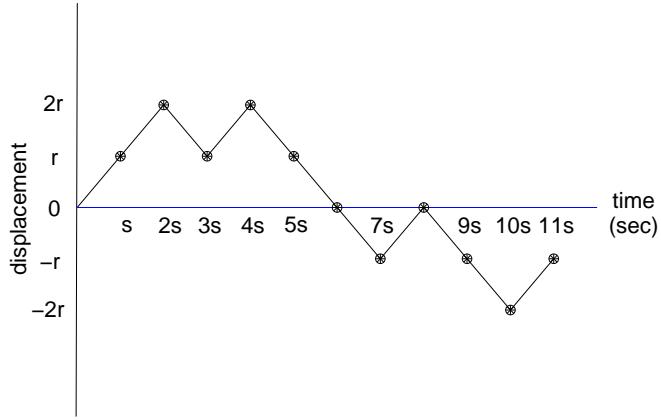


Fig. 8.4. The trajectory of an airborne measles droplet due to random collisions.

Since r_k is either $+r$ or $-r$, each term in the parentheses has an equal chance of being positive (the product of two positive terms or two negative terms) or negative (the product of one positive term and one negative term). Since each term will be of the same magnitude, r^2 , they should all cancel out, on average (i.e., for large n). This leaves us with the first n terms, but each has the same value (r^2). So we thus have the approximation

$$y_n^2 = nr^2.$$

Looking back at Figure 8.4, we see that since there are s seconds between each collision, then, after n steps, sn seconds will have elapsed.

$$\begin{aligned} t &= sn \\ n &= \frac{t}{s} \\ y^2 &= nr^2 \\ &= \left(\frac{t}{s}\right)r^2. \end{aligned}$$

The fraction $\frac{r^2}{s}$ is a positive constant (since s is positive), so we can rename it by some (positive) constant $C^2 = \frac{r^2}{s}$. (We use C^2 since we know for a fact that C^2 can't be negative.)

Thus

$$\begin{aligned} y_n^2 &= C^2 t \\ y_n &= \pm C \sqrt{t}. \end{aligned}$$

Graphing this, we have Figure 8.5.

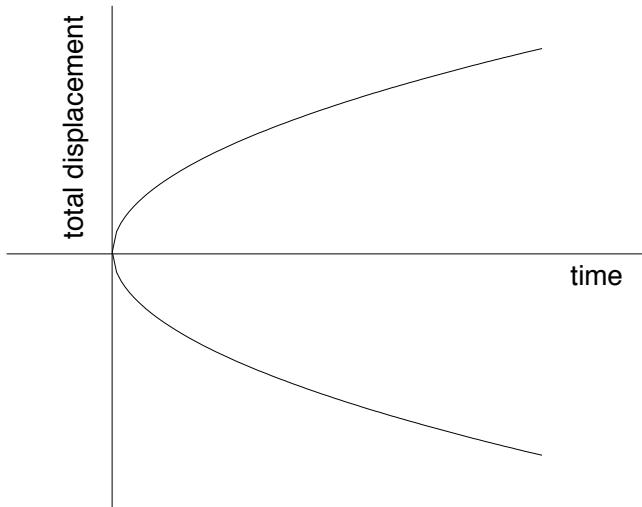


Fig. 8.5. Average trajectory of droplets over time.

What does this tell us? If we kept track of a whole lot of droplets, the tendency as a whole would be to follow paths within this parabolic trajectory. It's reasonable to assume that the spatial spread of droplets is approximately normally distributed at any given time. However, the distribution is wider for later times than it is for earlier times. Thus, the parabola traces one standard deviation from the average vertical displacement of 0. So 70% of the droplets should be found within this parabola at any time.

This tells us important information about the spread of measles over time, given one spatial dimension. This is very useful when determining how stringent any quarantine measures should be, for example. It also gives us a basic idea of how to determine the spread over two or three dimensions, although the mathematics is a lot harder.

8.3 Lab work

The problem

We want to track the spread of measles in a corridor. The time interval between collisions is $s = 1$ second, and each collision displaces an infected droplet by $\pm r = \pm 0.01$ mm. Assume there are N droplets. Find the average

displacement and the average squared displacement after n steps by calculating both y_n and y_n^2 for each droplet (i.e., find y_{av} and y_{av}^2) for the following values of n and N :

n (number of steps)	100	100	100
N (number of droplets)	15	200	1000

- What does y_{av} approach for large N ?
- What does y_{av}^2 approach for large N ?
- Include as part of your answer a plot of both of them on the same graph for $N = 15$ and $N = 200$.
- Suppose we want to fit the following model to the data: $y_{\text{av}}^2 = Dt^\lambda$. Rewrite this equation in a manner which may be compared to the computational data to determine the values of D and λ using the `polyfit` command. Use the data from the $n = 100$, $N = 200$ run.
- Is λ what you think it should be?
- Plot $\pm\sqrt{y_{\text{av}}^2}$ vs. t . On the same set of axes, plot the path taken by 15 droplets for $n = 1000$. What can you say about the way the disease spreads?

The solution

First we need to create a matrix that consists of the displacements of a number of droplets over time. In each time step, the droplet undergoes a displacement of $\pm r$. If we keep track of whether it was $+r$ or $-r$ for each droplet by placing these in a row of a matrix, then we can find the total displacement by simply adding all the entries in that row. So a row with three $+r$'s and two $-r$'s would add up to a total displacement of $+r$ for one particular droplet.

If we did this for every droplet (in subsequent rows), then we could find total displacements for each. Finding the average displacement is simply done by adding up all N total displacements and dividing by N . The average squared displacement can be found by squaring the total displacements and adding them together before dividing by N .

The displacement matrix will be a matrix with a bunch of 1's and -1's, multiplied by the scalar r . More importantly, it is a *random* matrix of 1's and -1's. Since there are N droplets, we will have N rows. We are also keeping track of n steps. Thus, we create an $N \times n$ random matrix.

To do this, we can use the `rand` function to generate a uniform distribution of random numbers between 0 and 1. This isn't quite what we want, however, so we have to convert these to either +1 or -1. We can use an `if ... then ... else` to separate the random numbers: those that are above 0.5 will be changed to 1 while those below 0.5 will be changed to -1.

```

clear all
n=100;
N=15;
r=0.01;
a=rand(N,n); % (a is the random matrix)
for i=1:N
for j=1:n
if a(i,j)>0.5
a(i,j)=1;
else
a(i,j)=-1;
end
end
end
a=r.*a;

```

Note that there's no output yet. Although if you want to see the random matrix, you can drop the last semicolon (;).

We'll put this together with some more pieces in a moment. These first few program fragments are to show you how the individual pieces in the whole program work.

In order to sum rows, columns or parts of rows or columns of a matrix, we use the `sum` command. However, we won't just sum the total of each row; instead we'll create an $N \times n$ matrix of running totals (i.e., the third column consists of the total displacement at the third step). Our "running total" matrix y should look like the following, where every row represents an individual droplet and r_1 is its first displacement, r_2 its second and so on:

$$y = \begin{bmatrix} r_1 & r_1 + r_2 & r_1 + r_2 + r_3 & \cdots & r_1 + r_2 + r_3 + \cdots + r_n \\ \vdots & \vdots & \vdots & & \vdots \\ r_1 & r_1 + r_2 & r_1 + r_2 + r_3 & \cdots & r_1 + r_2 + r_3 + \cdots + r_n \end{bmatrix} (N \times n).$$

We want rows to be added up to the j th column, for all j from 1 to n (so we'll need a `for` loop). We wish to do this for all N rows (so we'll need two `for` loops). Mathematically, we wish each entry of our "running total" matrix y to be equal to the sum of row i , columns 1 to j . This requires the following code:

```

for i=1:N
for j=1:n
y(i,j)=sum(a(i,1:j));
end
end

```

(Again, try dropping the semicolon (`;`) to see what happens.)

Here we are telling the `sum` command to look at the i th row, columns 1 through j , of the matrix a . We want Matlab to sum all the columns of our “running total” matrix: this adds the displacements for all N droplets at each step. Dividing the answer by N , we have a row matrix which represents the average displacement at all n steps. A similar approach yields the average squared displacement.

```
yav=sum(y)./N;
yav2=sum(y.^2)/N;
```

We also want to plot the average displacement and the two average displacements as calculated by the squaring method. That is, we have calculated y_{av}^2 , we now find the two new values of yav that represent the positive and negative square roots of y_{av}^2 . We will place these values in the first two rows of some new matrix Y .

```
Y(1,:)=yav2.^0.5;
Y(2,:)=-(yav2.^0.5);
```

We thus want to plot the variables `yav` and `Y` vs. time. How do we create our time matrix?

If there is a time s between steps then the first step will be at time $(1)(s)$, the second step will be at time $(2)(s)$ and the last (n th) step at time $(n)(s)$. The time matrix will go from s to $n.*s$ with increments of 1 (since n increases by 1 each time). Thus we can easily create our time matrix t and plot our variables. We add this code to the others, along with titles and `gtext` labels to label particular curves and using `hold on` to keep the $N = 15$ run when we rerun our program for $N = 200$.

```
clear all
n=100;
N=15;
%N=200;
r=0.01;
s=1;
a=rand(N,n); %(a is the random matrix)
for i=1:N
  for j=1:n
    if a(i,j)>0.5
      a(i,j)=1;
    else
```

```

a(i,j)=-1;
end
end
end
a=r.*a;
for i=1:N
for j=1:n
y(i,j)=sum(a(i,1:j));
end
end
yav=sum(y)./N;
yav2=sum(y.^2)/N;
Y(1,:)=yav2.^0.5;
Y(2,:)=-(yav2.^0.5);
t=s : n.*s;
plot(t,yav,t,Y)
title('The spread of measles for N=15 droplets')
xlabel('time')
ylabel('Average displacement and average displacement^2')
gtext('N=15')
%gtext('N=200')
hold on

```

You'll have to rerun this, making the obvious changes, to get the results for $N = 200$.

The next part of the problem involves simplifying $y_{\text{av}}^2 = Dt^\lambda$ so that we can calculate values of D and λ from our experimental data. To do this, we take the natural logarithm of both sides:

$$\begin{aligned}
y_{\text{av}}^2 &= Dt^\lambda \\
\ln(y_{\text{av}}^2) &= \ln(Dt^\lambda) \\
\ln(y_{\text{av}}^2) &= \ln(D) + \ln(t^\lambda) \quad (\text{since } \ln(ab) = \ln a + \ln b) \\
\ln(y_{\text{av}}^2) &= \ln(D) + \lambda \ln(t) \quad (\text{since } \ln(a^c) = c \ln a).
\end{aligned}$$

Thus, a log-log plot of y_{av}^2 vs. t would have a slope of λ and an intercept of $\ln(D)$. Since the slope and intercept are the coefficients of a linear polynomial, we can use the `polyfit` function to determine these. The command `polyfit` finds the coefficients of a polynomial $p(x)$ of degree n that fits the data in a least squares sense. The command `polyval` returns the value of a polynomial of degree n evaluated at x . We can use these two commands to fit polynomials of any degree to our data. Of course, if the degree is higher than the number of data points, Matlab will give you a warning.

We can define the first two rows of a new matrix L to be $\ln(t)$ and $\ln(y_{\text{av}}^2)$ respectively and then apply a linear fit. The program is again similar, but don't forget to make sure $N = 200$ this time.

```

clear all
n=100;
N=200;
r=0.01;
s=1;
a=rand(N,n); %(a is the random matrix)
for i=1:N
for j=1:n
if a(i,j)>0.5
a(i,j)=1;
else
a(i,j)=-1;
end
end
end
a=r.*a;
for i=1:N
for j=1:n
y(i,j)=sum(a(i,1:j));
end
end
yav2=sum(y.^2)/N;
t=1.*s : n.*s;
plot(log(t),log(yav2),'*')

```

Now go to Tools → Basic Fitting and apply a linear fit to the data. Click on the \rightarrow button and apply a linear fit. Write down the values p_1 and p_2 from the linear polynomial.

Run your program again and apply the linear fit. Notice how p_1 and p_2 have changed slightly? Try doing it all a third time.

Due to the random portion of this problem, the exact numbers you arrive at will differ slightly each time, but p_1 should be close to 1 (λ should be 1, since we proved in the previous section that $y^2 = C^2 t$) and p_2 should be in the range $(-9.4$ to $-9.0)$, making D equal to approximately $e^{-9.2}$.

Determining D helps us find the best model of the form we'd suspected, but determining λ validates our guesswork in the first part of this chapter. Remember that we'd made some approximations when deriving the equations, so the fact that $\lambda \approx 1$ tells us that our approximations were reasonable.

We now use these values of D and λ to create our mean distance plot and include the actual trajectories of 15 droplets to see if they behave in the manner that we are assuming they do. We are interested in the running total displacements of the 15 droplets (our y matrix), and we create new y values to fit our $y_{\text{av}}^2 = e^{-9.2} t$ formula. Let's assign the positive and negative y_{av} values

$(\pm\sqrt{e^{-9.2}t})$ to the first and second rows of a new matrix Y (new because we don't have such a variable in our present program, although we did have such a variable in our first program). Make sure $n = 1000$ and $N = 15$.

```

clear all
n=1000;
N=15;
r=0.01;
s=1;
D=exp(-9.2);
a=rand(N,n); % (a is the random matrix)
for i=1:N
  for j=1:n
    if a(i,j)>0.5
      a(i,j)=1;
    else
      a(i,j)=-1;
    end
  end
end
a=r.*a;
for i=1:N
  for j=1:n
    y(i,j)=sum(a(i,1:j));
  end
end
t=1.*s : n.*s;
Y(1,:)=(D.*t).^(0.5);
Y(2,:)=-(D.*t).^(0.5);
plot(t,y,t,Y)
title('Mean distance plot')
xlabel('time')
ylabel('Displacement')

```

These graphs will be slightly different each run, due to the random nature of the displacements, of course.

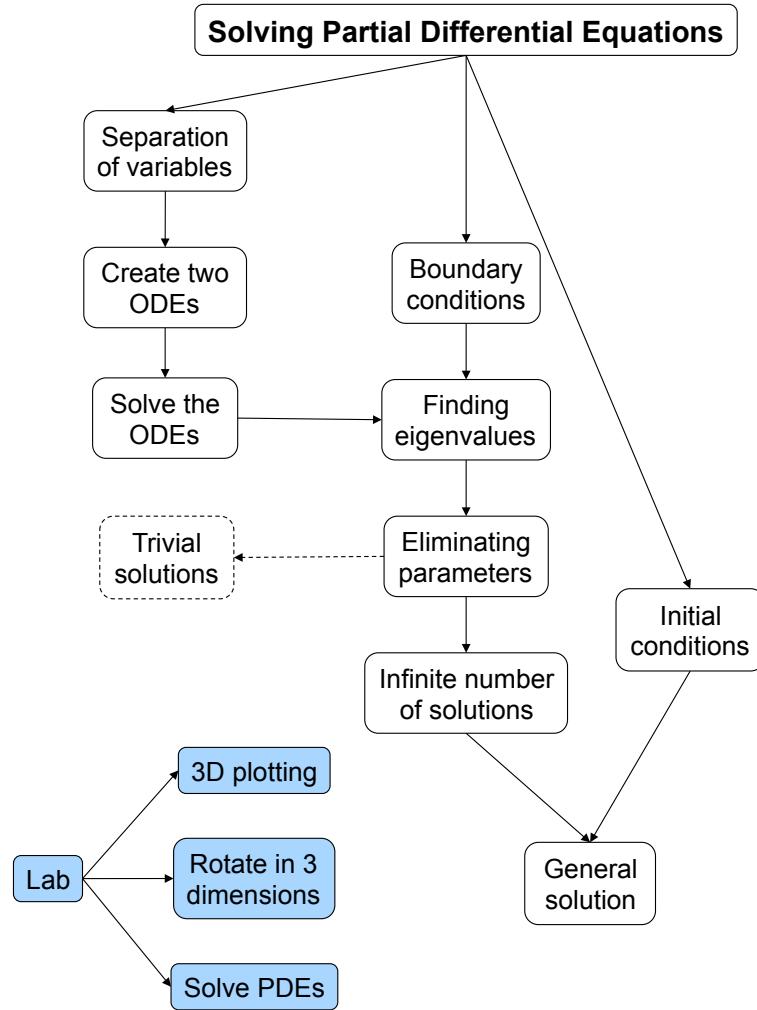
Analysis of the results

The standard average stays pretty much at zero, as we expect; there should be approximately as many droplets moving to the left as to the right at any given moment. However, our squared-derived average gives us much more interesting information: the mean distance travelled in the positive and negative directions by the droplets. This is why we calculate the squared average, to

acquire much more useful information about how the system will behave. Both the first and last sections of this lab demonstrate that the mean positive and negative y values correspond to the values associated with a parabola, should they be plotted vs. time. Keep in mind that this does not represent a physical parabola in two space dimensions; we are sticking to a one-dimensional problem here, and the droplets are spreading to the left or right. The values of the mean distances correspond to parabolic values with respect to time.

8.3.1 Exercises

1. Try running the first program several times, without changing N (thus overlaying different solutions on top of one another). How convincing is this?
2. Based on the mean distance output, does it seem reasonable that 70% of the trajectories are within the parabola?
3. Use the command `axis([0 1000 -2 2])` to change scale the y-axis (you can do this in the command window; you don't have to rerun the program). Now how does it look?
4. Try increasing N in the mean distance program (say $N = 50$). How do the results look now?
5. Based on these simulations, what would be a reasonable quarantine distance, if you wanted to be sure to contain more than 70% of the droplets?



Solving Partial Differential Equations

This chapter provides a method to analytically solve the PDEs from the previous chapter. Solving PDEs in general is extremely difficult and often outright impossible. Possible solutions could be functions of x , t or more likely both, in some manner.

By the end of this chapter, you should be able to turn PDEs into ODEs, apply boundary conditions, solve the ODEs and reconstruct the original solution. Or at least appreciate how complicated mathematical analysis of spatial modelling can get.

Note: The labs for this chapter don't require knowledge of the theory here, so if you get lost or want to skip the mathematics entirely, you can still do the lab (except for the last question).

9.1 Separation of variables

One approach is to assume that solutions are of the form $X(x)T(t)$, where X and T are yet-to-be-determined functions. That is, the solution can be a function of both x and t (which we'd expect), but it must be made up of a function purely of x and a function purely of t multiplied together. This is called *separation of variables*. For example, $xt + x = x(t + 1) = X(x)T(t)$. Obviously, there might be other solutions not of this form (e.g., e^{xt}), but it does give us a place to start.

We'll solve equation (8.4), the diffusion equation. For simplicity of notation, we represent partial derivatives with respect to a particular variable by subscripts of those variables. So equation (8.4) becomes

$$U_t = D U_{xx}.$$

Using separation of variables, we thus have

$$\begin{aligned} U &= XT \\ U_t &= X\dot{T} \\ U_x &= X'T \\ U_{xx} &= X''T, \end{aligned}$$

where we use $(\)'$ to denote spatial derivatives and $(\)^{\dot{}}$ to denote time derivatives. Since both X and T are solely functions of one variable (x and t respectively), their derivatives are ordinary derivatives, not partial.

Putting this into the diffusion equation, we have

$$\begin{aligned} X\dot{T} &= DX''T \\ \frac{\dot{T}}{DT} &= \frac{X''}{X}. \\ \uparrow & \uparrow \\ \text{independent} & \text{independent} \\ \text{of } x & \text{of } t \end{aligned}$$

We've moved all the variables involving time to one side and all the variables involving space to the other side. (The D could go on either side, but it's usually easier to stick constants with the time equation.)

However, something miraculous has occurred. The left-hand side of this equation is dependent only on t , while the right-hand side is dependent only on x . If we change t from, say, 2 to 54, we might assume the value of the left-hand side would change. If so, the right-hand side would have to change as well, but it is independent of t ! It doesn't have to change just because t changes. The only way for the left-hand side to equal the right-hand side at all times is if each side is always equal to the very same constant. That is, for $t = 4, 78$ etc or for $x = 89, -0.0002$ etc, both $\frac{\dot{T}}{DT}$ and $\frac{X''}{X}$ must always equal the same constant.

We thus have two independent ODEs, which are much easier to solve than PDEs. It turns out that if this constant is positive or zero, then there will be no (nontrivial) solutions. See Appendix G if you're interested in this. So we'll set both sides to equal the same negative constant, which we'll call $-\lambda^2$. Thus

$$\frac{\dot{T}}{DT} = \frac{X''}{X} = -\lambda^2 \quad (9.1)$$

$$\dot{T} = -\lambda^2 DT \quad (9.1)$$

$$X'' + \lambda^2 X = 0. \quad (9.2)$$

Solving equation (9.1) is relatively straightforward:

$$\begin{aligned}
\frac{1}{T} \frac{dT}{dt} &= -\lambda^2 D \\
\int \frac{1}{T} dT &= -\lambda^2 D \int dt \\
\ln T &= -\lambda^2 D t + c \\
T &= k e^{-\lambda^2 D t}, \tag{9.3}
\end{aligned}$$

where $k = e^c$.

To solve equation (9.2), we write

$$\begin{aligned}
\left(\frac{d}{dx} \right)^2 X + \lambda^2 X &= 0 \\
\left[\left(\frac{d}{dx} \right)^2 + \lambda^2 \right] X &= 0 \\
\left[\frac{d}{dx} + i\lambda \right] \left[\frac{d}{dx} - i\lambda \right] X &= 0,
\end{aligned}$$

using the difference of two squares $a^2 + b^2 = (a+ib)(a-ib)$ and where $i = \sqrt{-1}$. (We can only use this method because the equation has constant coefficients.) It's not worth dwelling on complex numbers if you haven't seen them before, but they're useful to solve polynomials. Then again, if you haven't seen complex numbers before, you're probably not reading this chapter in terribly much detail either!

Now we don't want $X = 0$, so either $\left[\frac{d}{dx} + i\lambda \right] X = 0$ or $\left[\frac{d}{dx} - i\lambda \right] X = 0$. These are both first-order, linear ODEs (despite the complex numbers), which are easy to solve. The first one is

$$\begin{aligned}
\frac{dX}{dx} + i\lambda X &= 0 \\
\frac{1}{X} dX &= -i\lambda dx \\
\int \frac{1}{X} dX &= -i\lambda \int dx \\
\ln X &= -i\lambda x + k_1 \\
X &= A_1 e^{-i\lambda x} \quad (\text{where } A_1 = e^{k_1}) \\
&= A_1 \cos \lambda x - A_1 i \sin \lambda x,
\end{aligned}$$

since $e^{i\theta} = \cos \theta + i \sin \theta$ and $e^{-i\theta} = \cos \theta - i \sin \theta$. By similar reasoning, the other solution is

$$\begin{aligned}
X &= A_2 e^{i\lambda x} \\
&= A_2 \cos \lambda x + A_2 i \sin \lambda x.
\end{aligned}$$

The general solution will be a linear combination of these two solutions... but they're both of the same form anyway. So the general solution is

$$X = B \cos(\lambda x) + C \sin(\lambda x).$$

(Since we don't know the coefficients, we can absorb the i into them.) This allows us to get the general solution.

9.2 Boundary conditions

However, we could also impose boundary conditions on our measles problem. If we assume the corridor is of length L with Harry standing at one end, then infectious droplets can't pass through the walls at either end, so there will be no particles at $x = 0$ or $x = L$ (since there's a wall at these points). This corresponds to the boundary conditions $U(0, t) = 0$ and $U(L, t) = 0$. (If we were to specify a certain distribution of the original sneeze at $t = 0$ then we would also have an initial condition $U(x, t) = f(x)$.)

Bear in mind that we can't have $T(t) = 0$ (or else the entire solution is zero), so the first boundary condition is

$$\begin{aligned} U(0, t) &= X(0)T(t) = 0 \\ \Rightarrow X(0) &= 0 \\ B \cos(0) + C \sin(0) &= 0 \\ B(1) + C(0) &= 0 \\ B &= 0 \\ \text{Therefore } X &= C \sin(\lambda x). \end{aligned}$$

We don't want $C = 0$, or else the solution would then be $x \equiv 0$, which we've already ruled out.

The second boundary condition gives us

$$\begin{aligned} U(L, t) &= X(L)T(t) = 0 \\ \Rightarrow X(L) &= 0 \\ C \sin(\lambda L) &= 0 \\ \Rightarrow \lambda L &= n\pi \quad (\text{since we don't want } C = 0) \\ \lambda &= \frac{n\pi}{L}. \end{aligned}$$

Because the sine function is zero at regular intervals ($n\pi$), our boundary conditions have allowed us to find an infinite number of λ values.

We thus have, for each n ,

$$\begin{aligned} X_n(x) &= C_n \sin\left(\frac{n\pi}{L}x\right) \\ T_n(t) &= k_n e^{-(\frac{n\pi}{L})^2 Dt} \quad (\text{from equation (9.3)}). \end{aligned}$$

We've put subscripts to indicate that each coefficient will be different for different n . The general solution is the product $U_n(x, t) = X_n(x)T_n(t)$, or

$$U_n(x, t) = A_n e^{-(\frac{n\pi}{L})^2 D t} \sin\left(\frac{n\pi}{L}x\right),$$

where A_n represents the value of $k_n C_n$ for each n . Not only are there an infinite number of λ values, but each has its own particular corresponding function (a different function depending on the value of n).

Thus there is more than one answer, depending on what n is. It does not represent a physical variable in the sense that time or distance does, but nevertheless there are various *different* solutions, as well, since the solutions corresponding to different n values are linearly independent. Addition or scalar multiplication (or both) of independent solutions always yields another viable solution. In the end, we determine which solution(s) best represent a particular case (i.e., which one of the mathematically correct solutions actually mimics “real life”) only through experimentation and observation.

We can write the general solution as

$$U(x, t) = \sum_{n=0}^{\infty} A_n e^{-(\frac{n\pi}{L})^2 D t} \sin\left(\frac{n\pi}{L}x\right). \quad (9.4)$$

9.3 Initial conditions

The final condition we could impose would be an initial condition. As an example, let's suppose we imposed the initial condition (corresponding to Harry's initial sneeze) as

$$U(x, 0) = 2 \sin \frac{\pi x}{L} - 3 \sin \frac{5\pi x}{L} + 14 \sin \frac{6\pi x}{L}. \quad (9.5)$$

So if we set $t = 0$ in our general solution (9.4), we get

$$U(x, 0) = \sum_{n=0}^{\infty} A_n \sin\left(\frac{n\pi}{L}x\right). \quad (9.6)$$

The exponential term has disappeared, because it's equal to 1 when $t = 0$.

All we need to do now is match (9.5) to (9.6). But this is extremely straightforward, since we just match the coefficients. Thus $A_1 = 2$, $A_5 = -3$ and $A_6 = 14$. Better yet, all other A_n 's match to zero! So the final solution only involves $n = 1, 5, 6$. Thus

$$U(x, t) = 2e^{-(\frac{\pi}{L})^2 D t} \sin \frac{\pi x}{L} - 3e^{-(\frac{5\pi}{L})^2 D t} \sin \frac{5\pi x}{L} + 14e^{-(\frac{6\pi}{L})^2 D t} \sin \frac{6\pi x}{L}. \quad (9.7)$$

9.4 Lab work

This lab explores the solution found in (9.7). You don't need to understand any of this chapter in order to do these exercises, except for the last one.

The problem

Use Matlab to plot the surface $z(x, y) = \sin(0.2xy)$ over the range $-7 \leq x \leq 7$, $-7 \leq y \leq 7$.

The solution

We need to transform our x and y variables into a “mesh”. That is, a grid of equally-spaced intervals, so that Matlab can plot the surface above each of the grid points. The command `meshgrid` will do this for us. We then create the function using the new, meshed, variables.

To plot the surface, we use the function `mesh`.

```
x=-7:0.1:7;
y=-7:0.1:7;
[X Y]=meshgrid(x,y);
z=sin(0.2.*X.*Y);
mesh(X,Y,z)
xlabel('x')
ylabel('y')
zlabel('z')
```

Looks great, huh? Except it's not quite obvious what this complicated function actually looks like. Fortunately, Matlab has an amazing solution to this problem.

Go to Tools → Rotate3D. This allows you to rotate the figure in any direction, using the mouse. Simply click on the figure and move the mouse around. You can swivel it around, move it up and down, even upside down.

Rotate your figure until you can clearly see the point in the middle. This is a very special kind of point, called a “saddle point”. It's a maximum in one direction (from underneath) and a minimum in another direction (from above). Just like a horse's saddle, where the point you sit on is also both a maximum (in the direction across your legs) and a minimum (in the direction of the horse's spine) simultaneously.

9.4.1 Exercises

1. Plot the following 3-dimensional surfaces over the range $-7 \leq x \leq 7$, $-7 \leq y \leq 7$:

- (a) $z(x, y) = \sin(x + y)$
- (b) $z(x, y) = \log(xy)$
- (c) $z(x, y) = \exp(-0.01x^2y^2)$

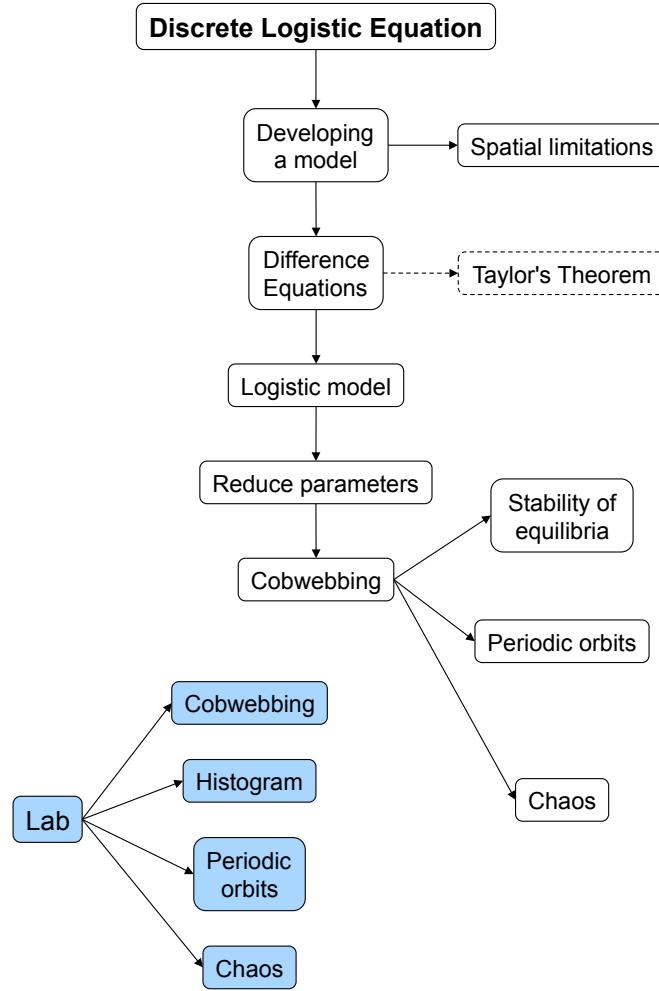
(d) $z(x, y) = \frac{\sin \sqrt{x^2+y^2}}{\sqrt{x^2+y^2}}$ (Note that the Matlab command for the square root is `sqrt`.) Is there a problem when $(x, y) = (0, 0)$? What's happening at this point?

Use the `Rotate3D` tool to examine each figure from different viewpoints.

2. Plot the 3-dimensional solution (9.7) with $D = 0.01$ and $L = 3$ over the range and the following ranges of time:
 - (a) $0 \leq t \leq 1$
 - (b) $0 \leq t \leq 5$
 - (c) $0 \leq t \leq 10$.
3. Describe in words what is happening in this plot.
4. Is this solution biologically reasonable? Explain.
5. (Theoretical) Try solving the boundary value PDE

$$\begin{aligned} \frac{\partial^2 U}{\partial t^2} &= c^2 \frac{\partial^2 U}{\partial x^2} \\ U(0, t) &= 0 \\ U(L, t) &= 0 \\ U(x, 0) &= 7 \sin \frac{3\pi x}{L} - 2 \sin \frac{5\pi x}{L} \\ \frac{\partial U(x, 0)}{\partial t} &= 0 \end{aligned}$$

using the methods outlined in this chapter. (Hint: All but two of the A_n 's will be zero.)



10

The discrete logistic equation

This chapter deals with finite difference equations, which are another way of examining changes to populations. However, instead of using a continuous variable (time), we use a discrete variable (which may be in “clumps” of time, or discrete events, such as annual changes) to model distinct changes from one period to another.

By the end of this chapter, you should be able to construct a simple difference equation, use the cobwebbing method to explore equilibrium points and their stability and see an introduction to the exciting worlds of periodic orbits and chaos.

10.1 Developing a nonlinear model

10.1.1 Spatial limitations

The growth of populations eventually becomes constrained by spatial considerations. A bacterial culture in a petri dish might exhibit a geometric growth rate for a time, but eventually the culture starts to fill the ecological niche (space and nutrients are not limitless) and the growth rate becomes sub-geometric.

Spatial considerations are hugely important in real-world applications, as we'll see in the coming chapters. Your chances of catching a disease will vary according to where you live in the world, places where you work, travel and visit and even the type of environment your neighbourhood is in. You're more at risk of West Nile virus in suburban areas with lots of green spaces than you are in an inner-city, urban environment.

However, spatially explicit models are also notoriously complicated, mathematically. The actual models might not be so hard to formulate, but the analysis gets very difficult, very quickly. In this chapter, we're not going to develop spatially explicit models, but we are going to consider spatial limitations as an implicit factor in our models.

10.1.2 Difference equations

My savings account gets \$10 added into it each Friday. So each week, my current savings are equal to the previous week's savings plus \$10. If “ a ” represents the amount of money in the account then, mathematically, the amount of money I have in the n^{th} week is

$$a_n = a_{n-1} + 10.$$

This is known as a “difference equation”. The value of a variable changes by a prescribed amount (not always a simple constant) from the $(n - 1)^{\text{st}}$ time period to the n^{th} period. More generally, the new variable is calculated according to some function f of the old value; it is “updated” according to a kind of “update rule”:

$$a_n = f(a_{n-1}), \quad (10.1)$$

where f is some function. The subscript n doesn't exactly count time; it is an integer which counts the *number of time periods* since the initial “start time” at $n = 0$. This sort of equation is called a *difference equation*.

It is important to distinguish difference equations from differential equations (DEs). Difference equations deal with finitely spaced, instantaneous “steps”, whereas DEs are based on continuous changes in variables. In the above example, changes (increases in the amount of money in my savings account) only occur once a week; they are specific, instantaneous events, and in between them nothing happens. As we've seen, DEs (both ordinary and partial) are used to model systems that change continuously, such as the spread of influenza through a population. The epidemiological spread of disease is well modelled by continuous dynamics. In reality, of course, influenza is not spread continuously, but rather from person to person. However, as the population gets larger and larger, then the contribution from each individual gets smaller and smaller... so, in the limit, the spread would become continuous, and we'd find ourselves dealing with DEs.

Let's consider equation (10.1), but we'll drop the subscripts for the moment. For the annual spread of (say) smallpox through a community, it is reasonable to assume

- $f(a = 0) = 0$ (no one infected, no subsequent infections)
- $f(a) > 0$ whenever $a > 0$ (how do you have a negative population?)
- f is differentiable (a common assumption in applied mathematics; it means the function is “smooth” without unrealistic jumps or corners).

A linear growth rate, for example, would be $f(a) = ra$, with $r > 0$. However, spatial considerations require us to look for something that is less than linear, or else the disease would go on spreading forever. This means we want the growth rate (the derivative) to be slowing down as a increases. If the derivative is slowing down, then *its* derivative will be negative. Thus

$$f''(a) < 0 \quad \text{for all } a > 0.$$

By Taylor's Theorem (see Appendix H), we can rewrite f as

$$f(a) = f(0) + af'(0) + \frac{1}{2}a^2 f''(0) + O(a^3), \quad (10.2)$$

where $O(a^3)$ indicates that the lowest term in the remainder is of the order of a^3 , which we can approximate by zero (since a is small, a^3 is tiny, so $O(a^3)$ is negligible compared to the other terms). Now, if $f(a) > 0$ when $a > 0$, then the first derivative at $a = 0$ must be positive (or possibly zero, but that's not very interesting). We'll sharpen this up and assume

$$f'(0) = r > 0.$$

Since $f'' < 0$, we can write

$$f''(0) = -2b < 0$$

(the 2 is inserted merely to balance out the $\frac{1}{2}$ in equation (10.2)). Substituting into our simplified Taylor's expansion gives

$$\begin{aligned} f(a) &= af'(0) + \frac{1}{2}a^2 f''(0) \quad (\text{Remember } f(0) = 0) \\ &= ar + \frac{1}{2}a^2(-2b) \\ &= ra - ba^2. \end{aligned}$$

Putting it all together, our model, known as the “discrete logistic equation”, is

$$a_n = ra_{n-1} - ba_{n-1}^2. \quad (10.3)$$

This is definitely not linear (the a_{n-1}^2 term guarantees this). Thus, we cannot expect clear, predictable behaviour out of this equation all the time.

Intuitively, this equation can be seen to include a natural (linear) growth rate ($a_n = ra_{n-1}$), minus a “competition” term ($-ba_{n-1}^2$) which would increase dramatically as the disease spreads and infectious individuals “compete” for the same limited susceptible individuals.

Let's just note at this point that we've managed to derive a model using the barest of assumptions (linear growth and a nonlinear spatially limiting term). As we analyse this model in this chapter and the next, try to remember just how simple this model was to formulate. We'll see all manner of crazy dynamics from this very basic model, which tells us a lot about how much is going on when we step into the world of nonlinearity.

We can plot equation (10.3) with the updated value a_n on the y -axis and the previous value a_{n-1} on the x -axis. See Figure 10.1.

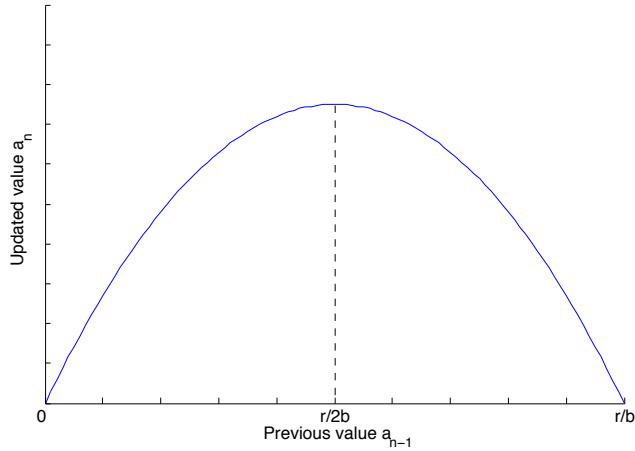


Fig. 10.1. Difference equations are visualised by plotting the previous value versus the updated value.

Since we had a lot of information around $a = 0$, we expect our model to be fairly accurate in this region. The presence of a maximum makes sense: there would be a maximum number of infections that the community could ever support, no matter what the growth rate of the infection was. Further away from $a = 0$ however, the results are potentially less accurate; we had no information about this area, and one of our key assumptions (“throwing away” the higher-order terms in the Taylor expansion) relied on the assumption that a was small.

So the applicability of the model has limitations, like all models. But let’s put aside these suspicions and see what it tells us anyway. Note that the model only makes mathematical sense on the interval $[0, r/b]$ (why?). Let’s rescale to make life easier:

$$x = \frac{b}{r}a \quad \Rightarrow a = \frac{r}{b}x.$$

Substituting this into equation (10.3), we have

$$\begin{aligned} a_n &= ra_{n-1} - b(a_{n-1})^2 \\ \left(\frac{r}{b}x_n\right) &= r\left(\frac{r}{b}x_{n-1}\right) - b\left(\frac{r}{b}x_{n-1}\right)^2 \\ \frac{r}{b}x_n &= \frac{r^2}{b}x_{n-1} - \frac{r^2}{b}(x_{n-1})^2 \\ x_n &= rx_{n-1} - rx_{n-1}^2. \end{aligned}$$

We have thus reduced our equation to

$$x_n = rx_{n-1}(1 - x_{n-1}). \quad (10.4)$$

Now we only have one parameter to deal with (r), and the parabola is contained within $0 \leq x \leq 1$. Note that x doesn't directly represent the infected population (it is off by the ratio used to define it). Nevertheless, reducing the number of parameters is always helpful, especially in more complicated systems.

We wish to follow the dynamics (the time progression) of the system. To start, we draw a point whose x -axis coordinate ("old value") must be equal to the initial condition x_0 and whose y -axis coordinate ("updated value") must be equal to the first "output" value x_1 ; this is the point (x_0, x_1) . In the next iteration, x_1 is now our old value (x_{n-1}), and x_2 is the new x_n value; thus the next point will be (x_1, x_2) . And so on. See Figure 10.2.

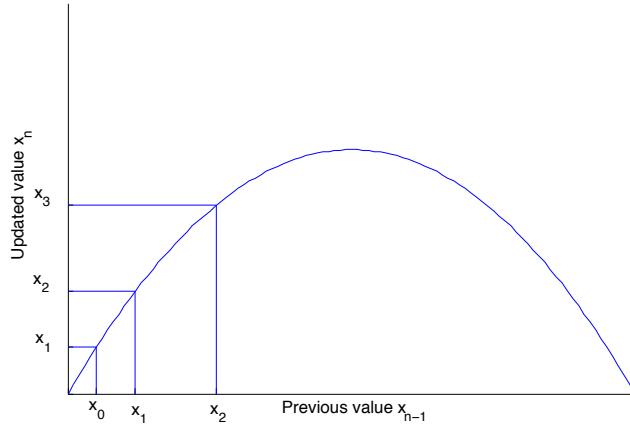


Fig. 10.2. Iterating the dynamics.

10.2 Cobwebbing

There is an easier way to do this, called *cobwebbing*. Note that the old y -axis value always becomes the next x -axis value. Thus, if we drew a horizontal line from one point (x_{k-1}, x_k) to the line $x_n = x_{n-1}$ ($y = x$), we would be at the point (x_k, x_k) . Drawing a vertical line up or down to meet the curve, we would encounter the next point (x_k, x_{k+1}) . To illustrate this, in Figure 10.3, we start at the point $(0, x_0)$ to find the points illustrated above in Figure 10.2.

We will take advantage of cobwebbing to aid us in determining the stability of any equilibria that may exist. As in DEs, equilibrium points of difference equations are those where the system is, and remains, in steady state. Their

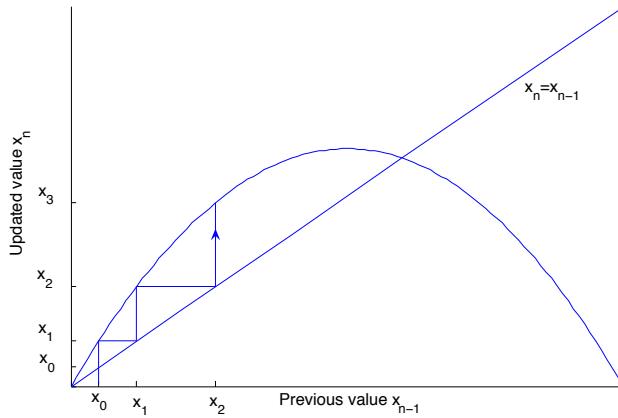


Fig. 10.3. Cobwebbing is the same as iterating in Figure 10.2, but it's easier to visualise the trajectories.

stability tells us important information about the long-term behaviour of the system. We can guess that 0 might be an equilibrium — smallpox can't just appear out of nowhere — but it's not clear what other points might be, unless we do a little mathematics.

Note that whenever the curve crosses the line $x_n = x_{n-1}$, there must be an equilibrium. (Convince yourself of this if you're not sure.) Mathematically, this means equilibria x^* must satisfy

$$\begin{aligned} x^* &= rx^*(1 - x^*) \\ (1 - r)x^* + r(x^*)^2 &= 0 \\ x^* [(1 - r) + rx^*] &= 0. \end{aligned}$$

This is a quadratic in x^* , so the roots are

$$x^* = 0, \quad x^* = \frac{r - 1}{r}.$$

If $r = 1$, then 0 is the only equilibrium. If $r > 1$, then an equilibrium point will exist such that $0 < x^* < 1$. If $r < 1$, then the nonzero equilibrium would be negative, and we aren't interested in that part of the graph, since negative populations are a physical impossibility. Graphically, $r > 1$ corresponds to a curve "high enough" to cross the $x_n = x_{n-1}$ line, not just meet it at 0. The case $r \leq 1$ corresponds to a curve that lies totally below this line on the range $0 < x \leq 1$, meeting it only at 0.

10.2.1 Stability in the discrete logistic equation

We've seen stability before, in the context of DEs. Here, we're fundamentally talking about the same concept, but there's a more intuitive way of "seeing"

stability when we use cobwebbing. An equilibrium point is stable if nearby points move towards it and unstable if nearby points move away from it.

To explore stability, we start first with the $r < 1$ case, investigating the stability of the (only) equilibrium 0 by starting out near it and following the dynamics by cobwebbing. Figure 10.4 shows the case for $r = 0.5$. It's clear that, in this case, 0 is a stable equilibrium, since the cobwebbing heads towards it.

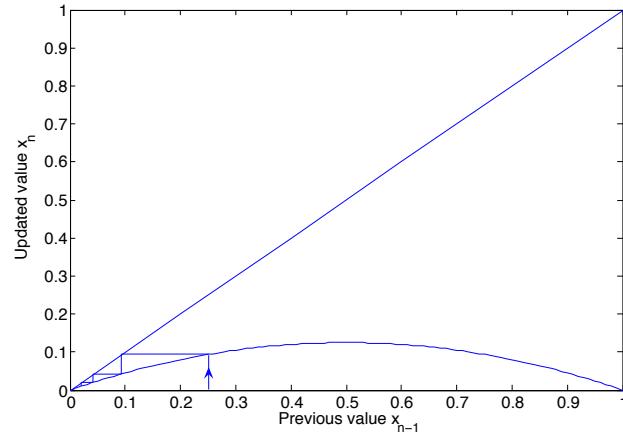


Fig. 10.4. Logistic map for $r = 0.5$.

Next, we look at the $r \geq 1$ case. Figure 10.5 shows the case for $r = 2$. Here we see that not only is 0 an unstable equilibrium, but that the other equilibrium is stable. The cobwebbing heads away from 0 and towards the other point. We can calculate the exact value of this other equilibrium point:

$$\begin{aligned} x^* &= \frac{r-1}{r} && \text{from above} \\ &= \frac{2-1}{2} \\ &= 0.5. \end{aligned}$$

To examine stability mathematically, we can use Taylor's theorem again (Appendix H). If we start with an initial point x_0 very close to x^* , we could write

$$x_0 = x^* + \epsilon,$$

where ϵ is very small indeed (although it may be positive or negative). Then by Taylor's theorem (Appendix H),

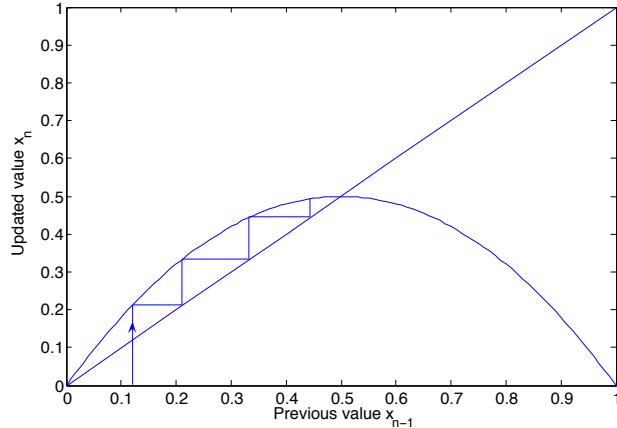


Fig. 10.5. Logistic map for $r = 2$.

$$\begin{aligned}
 x_1 &= f(x_0) = f(x^* + \epsilon) \\
 &= f(x^*) + \epsilon f'(x^*) + O(\epsilon^2) \\
 &\approx x^* + \epsilon f'(x^*),
 \end{aligned}$$

since $f(x^*) = x^*$, and we can approximate $O(\epsilon^2)$ by zero since ϵ is small. If $f'(x^*) > 0$, then x_1 and x_0 lie on the same side of x^* , whereas if $f'(x^*) < 0$, then x_1 and x_0 lie on opposite sides. More importantly though, if $|f'(x^*)| > 1$, then x_1 is *further* from x^* than x_0 is. That is, successive trajectories are travelling away from the equilibrium, so it must be unstable. Conversely, if $|f'(x^*)| < 1$, then x_1 is closer to x^* than x_0 is and hence we have a stable equilibrium.

10.2.2 More complex behaviour

The cobwebbing results for $r = 3.2$ and $r = 4$ are shown in Figures 10.6 and 10.7, respectively. Each was iterated 100 times. What's going on here?

In the first case, the population ends up oscillating between two different values forever. This is called a periodic orbit, since the orbit cycles periodically between two (or possibly more) different states. Periodic behaviour is very common in nature, as seen in the cycling of the seasons, the beating of your heart or the winter outbreaks of influenza.

In the latter case, the system is undergoing “chaos”. Chaotic systems are highly unpredictable, as you can guess. Tiny fluctuations at one stage will lead to huge differences later on. This is disastrous news for our attempts to use computers to model such systems, because tiny fluctuations in the computation (e.g., roundoff errors) will have huge impacts later on. Chaotic systems occur in nature: the weather (this is why the weatherman can't predict

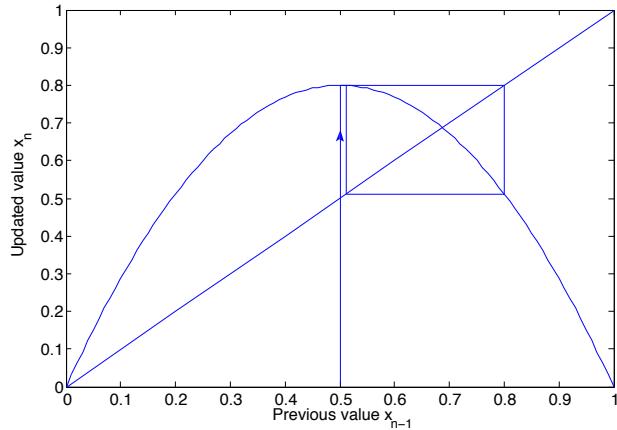


Fig. 10.6. Logistic map for $r = 3.2$.

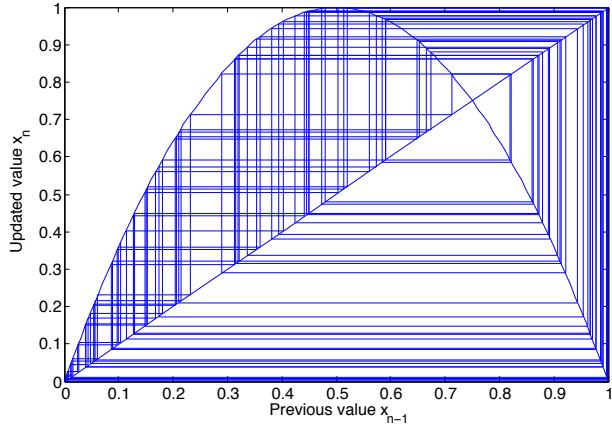


Fig. 10.7. Logistic map for $r = 4$.

more than a few days ahead and also why he's frequently wrong), a fibrillating heart (hopefully not yours) and the unpredictable population of cannibalistic flour beetles (who sometimes eat eggs or pupae, but not always).

Figure 10.7 illustrates the difference between linear and nonlinear systems; a linear system would never give weird unpredictable results like these. We'll explore some more of these issues in the next chapter.

10.3 Lab work

The problem

We wish to explore the discrete logistic equation for various values of the parameters, beginning with cobwebbing. This lab emphasises the exploration of different parameters over the creation of code (since the necessary programs are somewhat difficult for beginning Matlabbers). Below you will find the programs necessary for the exercises, where you can explore the discrete logistic equation.

The solution

The following program plots the discrete logistic equation, the $x_n = x_{n-1}$ ($y = x$) line and cobwebs for as many iterations as the user wants. It also allows the user to choose different r and initial values. Note that it runs continuously, always erasing the old figure (the only way to escape the program is to use the **Control-C** break). Pauses are inserted so that it will be easier to follow the cobwebbing over time (if you choose a very large number of iterations, you will want to place a % in front of the pause lines to speed things up).

Be sure to type the codes in precisely as-is, because they're designed to give outputs on screen.

```
%Cobwebbing for the discrete logistic equation
while 1<2 %(will run continuously)
    clear all
    r=input('Choose r, the bifurcation parameter, in (0,4) ');
    n=input('Choose n, the number of iterations ');
    x=zeros(1,n+1);
    x(1)=input('Choose x(1), the initial point, in [0,1] ');
    clf
    s=0:0.2:1;
    plot(s,s,'b')
    t=0:0.01:1;
    hold on
    v=r.*t.*(1-t);
    plot(t,v,'g')
    pause(1)
    w=zeros(1,3);
    y=w;
    for i=2:n+1
        x(i)=r.*x(i-1).*(1-x(i-1));
    end
```

```

plot([x(1) x(1)], [x(1) x(2)])
for i=2:n
    plot([x(i-1) x(i)], [x(i) x(i)])
    pause(1)
    plot([x(i) x(i)], [x(i) x(i+1)])
end
xlabel('Previous value x_{n-1}')
ylabel('Updated value x_n')
title('Logistic Map x(n)=r*x(n-1)*(1-x(n-1))')
end

```

Beyond $r = 3$, it may be difficult to tell exactly what's going on. Another way to look at the same problem is to create a histogram, counting the number of times the population reaches a certain value over time. Note that, in both this and the cobwebbing program, small errors are introduced since the computer you are using can only handle so many decimal points; eventually these may add up to slightly skew your results.

The following program determines an individual trajectory (with a user-defined value of r) for 10,000 iterations and counts the amount of time it spends in different subintervals of $[0, 1]$, each with width 0.001. The starting point is always 0.5, and a histogram is plotted at the end. Note that it is also designed to run "forever" (but you may have to hit Control-C TWICE to exit this program).

```

%Histogram for the discrete logistic equation
clear all
x0=0.5;
n=10000;
bins=0:0.001:1;
while 1<2
    r=input('What is r (between 0 and 4)?');
    clf
    y=zeros(1,n);
    y(1)=r.*x0.*(1-x0);
    for k=2:n
        y(k)=r.*y(k-1).*(1-y(k-1));
    end
    hist(y,bins);
    xlabel('Population values scaled to lie between 0 and 1')
    ylabel('The number of times pop^n lies within value x and
x+\Delta x')
    title('Frequency histogram for the Logistic map')

```

```
end
```

10.3.1 Exercises

Throughout these exercises, keep track of your results for specific r values. See if you can build up a “big picture” of what is going on. These programs will come in handy for the next chapter’s lab, so be sure to save your work.

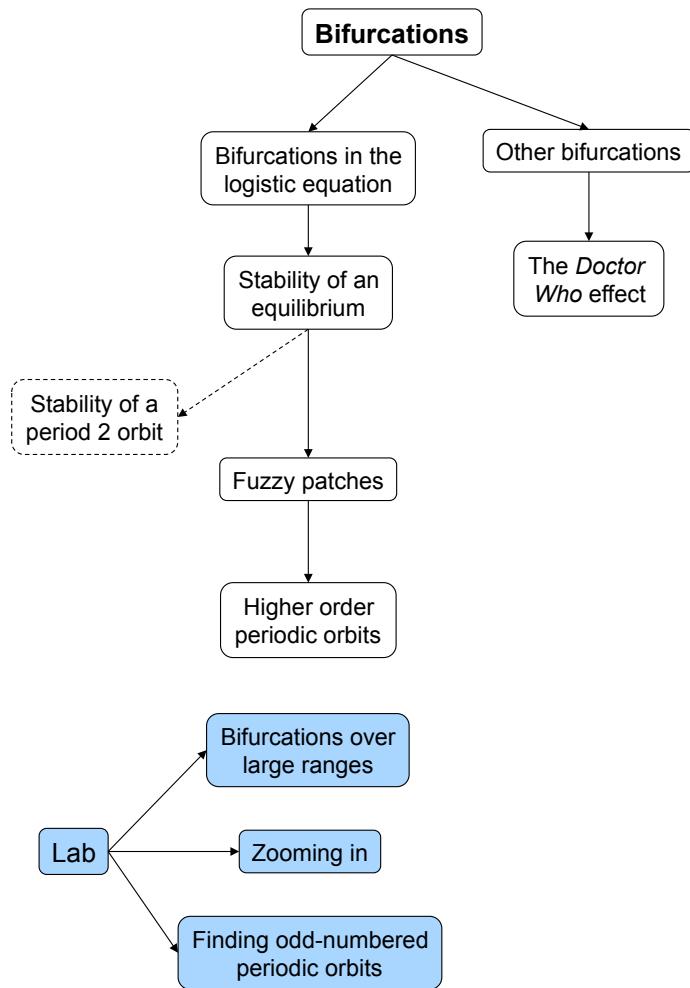
1. For a value of $r \in (0, 2)$, sketch the graph of the parabola by hand and cobweb by hand starting from the initial value of 0.5.
2. Now, using the cobwebbing program, repeat for the same initial value and r (for 20 iterations), comparing this to your sketch.
3. For the same value of r , run the program with several initial values. For each run, describe the long-term behaviour of the web and the corresponding time evolution. Can you draw any conclusions about the dependence on the starting value?
4. Run the program starting from 0.5, for a range of r values from 1 to 3. Describe the long-term behaviour of the web, and the corresponding time evolution.
5. Run the program starting from 0.5 for a range of r values from 3 to 4. Try to describe the long-term behaviour of the web, and the corresponding time evolution (in the chaotic regime, you might want to use very large n values and put a % in front of the pause commands).
6. Use the histogram program to explore the same r values as you investigated above. Can you clear up some of the confusion you may have had for some r values?
7. Can you find a value of r which produces exactly 2 spikes? 4 spikes? 8 spikes? Draw the histograms in each case.
8. Can you find a value of r that produces exactly 3 spikes (this is tricky)?
9. Can you find a value of r that produces exactly 5 spikes (this is even harder)?
10. Run your cobwebbing program again and see what the trajectories look like.
11. Consider the difference equation $x_{n+1} = \lambda \sin(\pi x_n)$ on the range $0 \leq x_n \leq 1$.
 - a) Use Taylor’s theorem to show that 0 is an equilibrium point.

- b) Can you show that there's a second equilibrium point, assuming λ is large enough? (Note: you won't be able to find it explicitly, but it's still possible to show that there has to be one.)
- c) Use your cobwebbing program to investigate the system when

$$\lambda = 0.2, 0.5, 0.8, 0.9.$$

- d) Use your histogram program to confirm your findings for these values.

12. Bonus: Consider the difference equation $x_{n+1} = \alpha \sin^2 x_n$. Use your cobwebbing program to explore what happens for different values of α . How does the system change depending on where you start your iterations?



11

Bifurcations

This chapter introduces one of the most important concepts that mathematics can bring to biology that biologists are rarely taught. Bifurcations involve the system changing in some fundamental way when a certain parameter reaches a critical value. In fact, we've already seen a bifurcation, back in Chapter 6. When the parameter R_0 changes from being above 1 to below 1, there's a fundamental change in the system; that is, the disease stops being endemic and will die out on its own.

Bifurcations are intricately tied to stability. A system may undergo a bifurcation when an equilibrium changes from being unstable to stable (as in the R_0 example above; the disease-free equilibrium changes from unstable to stable) or when a new equilibrium (or periodic orbit) is created.

By the end of this chapter, you should be able to determine the stability of an equilibrium of a difference equation, understand how one determines the stability of periodic orbits of a difference equation, understand how to read a bifurcation diagram and have further reasons to develop a healthy scepticism about relying too much on computers.

11.1 Stability in difference equations

In the previous lab, we saw the stability of equilibria of the logistic equation change as a certain parameter, r , changed. First there was only one equilibrium, at 0 (Figure 10.4), then its stability changed from stable to unstable when a new (stable) equilibrium was created (Figure 10.5). As r increased, this new equilibrium became unstable, and a periodic orbit was created (Figure 10.6). As r increased further, there were more periodic orbits and then chaos (Figure 10.7).

The chaotic regime can be thought of as a regime where the detailed long-term behaviour of the deterministic system is effectively unpredictable because of sensitivity to the initial conditions. This system is an example of what is called “the period-doubling route to chaos”.

However, as r increased even further, we were back to periodic orbits again (assuming you could find the Period 3 and 5 orbits in the lab).

How can we describe what's happening here in a way that's easy to understand? As always, the best way to "see" what's going on will be to draw a picture. We'll have to walk through some of the mathematics to understand the issues here, but we'll use that to build up our *bifurcation diagram*. This is a diagram that illustrates the changing nature of equilibria, periodic orbits etc, with the parameter of interest plotted on the x -axis. It's a neat conceptual way of visualising the various aspects of the system at a glance, as we'll see.

11.1.1 The range $0 < r < 3$

From Section 10.2 of the last chapter, we had two equilibria, 0 and $\frac{r-1}{r}$. Let's call these \bar{x}_1 and \bar{x}_2 , respectively. (Notice that our nontrivial equilibrium is different for different values of r .)

In Section 10.2.1, we saw that $|f'(x^*)| < 1$ was our condition for stability of an equilibrium x^* . So let's differentiate the function:

$$\begin{aligned} f(x) &= rx(1-x) \\ f'(x) &= r(1-2x). \end{aligned}$$

We thus have

$$\begin{aligned} f'(\bar{x}_1) &= r \\ f'(\bar{x}_2) &= r \left(1 - 2 \cdot \frac{r-1}{r}\right) = 2 - r. \end{aligned}$$

It follows that \bar{x}_1 is stable if and only if $|r| < 1$, while \bar{x}_2 is stable if and only if $|2 - r| < 1$. This condition is $-1 < 2 - r < 1$ or $1 < r < 3$. This is the same result we obtained from cobwebbing: up to 1, the zero equilibrium was the only stable one, while from 1 to 3, only the nonzero equilibrium was stable.

We can plot this in a bifurcation diagram, measuring the stability of the two points against the parameter r . In such a diagram, stable phenomena are represented by solid curves, while unstable phenomena are represented by dashed curves. In the lab, the program can't "see" unstable equilibria, since it can never reach them, so the bifurcation programs only show us stable orbits. Unless you happen to hit the equilibrium exactly, which is unlikely, you'll always be moving away from an unstable equilibrium, so the computer only finds stable phenomena.

In Figure 11.1, we see that, for $0 < r < 1$, there is only one equilibrium, and it is stable. For $1 < r < 3$, there are two equilibria (not counting negative values), one stable and one unstable. Naturally, the nontrivial equilibrium will be different for different values of r , which accounts for the curve.

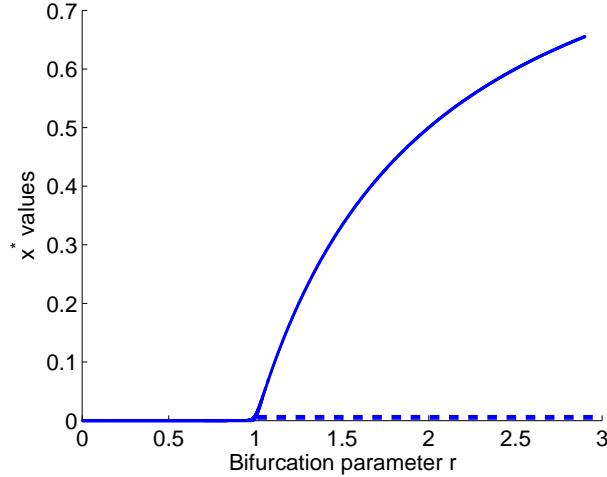


Fig. 11.1. Bifurcation diagram for $0 < r < 3$.

11.1.2 The range $3 < r < 3.45$

What happens beyond $r > 3$? In Section 10.2.2, we saw that the population around $r = 3.2$ oscillated between two points, w_1 and w_2 . The system has undergone another bifurcation as r changed from $r < 3$ to $r > 3$, splitting from period one (an equilibrium) to period two (a periodic orbit). “Period two” means

$$w_2 = f(w_1) \\ \text{and} \quad w_1 = f(w_2).$$

Therefore

$$w_1 = f[f(w_1)] \equiv g(w_1).$$

Let’s examine the stability of $g(x)$. We can use the definition of g to derive an explicit expression for $g(x)$:

$$\begin{aligned} g(x) &= f[f(x)] \\ &= rf(x)\{1 - f(x)\} \\ &= r[rx(1 - x)]\{1 - [rx(1 - x)]\} \\ &= r^2x(1 - x)[1 - rx(1 - x)]. \end{aligned}$$

This might look like Figure 11.2.

Why “might”? Because, depending on the value of r , the curve might cross the line, once (if the two “hills” were small enough, 0 would be the only

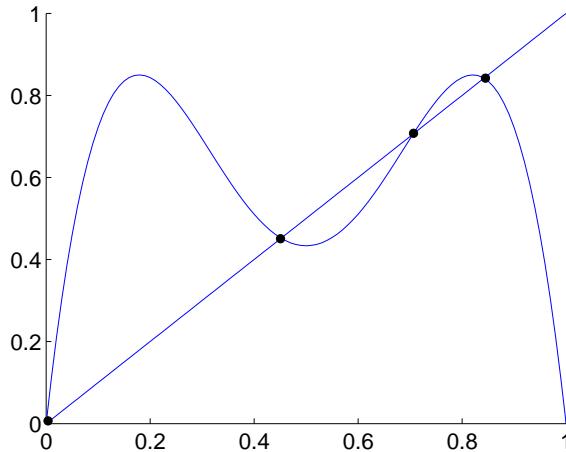


Fig. 11.2. Equilibrium points of $g(x)$.

equilibrium), twice (if the “valley” were sufficiently raised or the second “hill” small enough), three times (if the “valley” just touched the line), or four times (as we see here). What’s happening here is another bifurcation, but this time of periodic orbits.

Two of these equilibria we’ve seen before: the first and third equilibria are \bar{x}_1 and \bar{x}_2 from before. How do we know this for sure? Well, since \bar{x} is an equilibrium of f , then it must be an equilibrium of g . Thus

$$\begin{aligned} g(\bar{x}) &= f[f(\bar{x})] \\ &= f(\bar{x}) \quad (\text{since } f(\bar{x}) = \bar{x}) \\ &= \bar{x} \quad (\text{for the same reason}). \end{aligned}$$

What about the stability of these points? We know that both \bar{x}_1 and \bar{x}_2 are unstable for $r > 3$, and we saw in the last chapter that we have a periodic orbit when $r > 3$. So our other two points are stable, at least for a while.

What’s happening thus far is that our nonzero equilibrium “continues” on from the region of only one stable equilibrium, but it is no longer stable. So our bifurcation diagram looks like Figure 11.3.

We also saw in the previous lab that the Period 2 zone only existed for a certain range beyond 3. After that, they split into Period 4 points and so on. We can determine the extent of this range using the same arguments as we made for the singular solution zone. Thus we look for conditions on r such that $|g'(w_1)|$ and $|g'(w_2)| < 1$ (for stability of the two points).

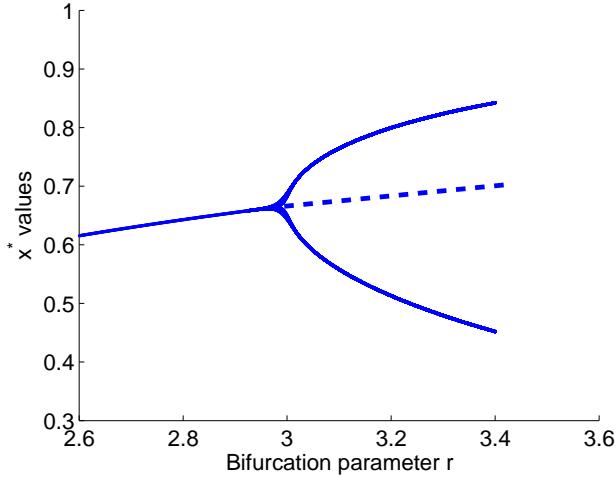


Fig. 11.3. Bifurcation diagram for the appearance of Period 2 points.

$$\begin{aligned} g'(x) &= \frac{d\{f[f(x)]\}}{dx} \\ &= f'[f(x)] \cdot f'(x) \end{aligned}$$

using the chain rule and the definition of g . Next, since $f(w_1) = w_2$ and $f(w_2) = w_1$,

$$\begin{aligned} g'(w_1) &= f'(w_2)f'(w_1) \\ &= f'(w_1)f'(w_2) \quad \text{swapping these around} \\ &= g'(w_2). \end{aligned}$$

We now look to see when $|g'(w_1)| = |f'(w_2)f'(w_1)| < 1$, with the understanding that the results will apply to $|g'(w_2)|$ as well. It turns out (see Appendix I) that

$$|g'(w_1)| = |4 + 2r - r^2|.$$

We are interested in the boundaries themselves; that is, the values of r for which $|g'(w_1)| = 1$, or $g'(w_1) = \pm 1$. If $g'(w_1) = +1$, then

$$\begin{aligned} 4 + 2r - r^2 &= 1 \\ r^2 - 2r - 3 &= 0 \\ (r - 3)(r + 1) &= 0. \end{aligned}$$

In this case, the boundaries are $r = 3$ or -1 . We'll ignore the negative root (we already said we weren't interested in this region), and the 3 corresponds

to the left-most boundary for the Period 2 zone. That is, as we suspected, the Period 2 points only exist for $r > 3$.

For the other case, $g'(w_1) = -1$, we have

$$\begin{aligned} 4 + 2r - r^2 &= -1 \\ r^2 - 2r - 5 &= 0 \\ r &= \frac{2 \pm \sqrt{4 + 20}}{2} \\ &= \frac{2 \pm 2\sqrt{6}}{2} \\ &= 1 \pm \sqrt{6} \\ &\approx 3.45 \end{aligned}$$

(after we ignore the negative root). This means that the Period 2 zone should extend from $r = 3$ to $r \approx 3.45$.

However, when you run your bifurcation program in the lab for (say) $2.8 < r < 3.5$, you'll notice a sort of fuzzy patch precedes each bifurcation, starting slightly before the points we've just calculated. What's causing this? Answer: computer round-off error. The area near the bifurcations is so sensitive (nonlinear systems generally are, anyway) that we aren't getting the proper results. The range and "amplitude" of the fuzzy area will vary from computer to computer and, in some cases, depend on the program used to create the graph (Matlab, C++, Fortran etc).

11.1.3 The range $r > 3.45$

For higher periods, we could try to repeat the same analysis as for Period 2. From this, we'd have a seventh order polynomial, out of which we could extract three factors (one for the original nontrivial equilibrium, two for the now-unstable Period 2 points), leaving us with a fourth order polynomial to solve (corresponding to Period 4). This is much too hard to do in general, so we take the lessons we've learnt from the theory surrounding Period 2 and rely on the computer simulations to guide us through higher-order periods.

We'll see in the lab that not only do we get Period 4, Period 8 and so on, these start to happen quicker and quicker as r increases. We not only get chaos, which is pretty interesting in itself, but there are also Period 3, 5, 7 and so on orbits *after* the appearance of chaos. Who knew that our little logistic equation, that was created from the barest of assumptions with only a very slight nonlinearity to it, could be capable of such diverse behaviour?

The one key point to remember: nonlinear systems are complicated and sometimes subject to computer error, so be careful.

11.2 The *Doctor Who* effect

If you ever have access to video equipment, here's a fun game you can play. Point the camera at a TV showing its own output. As you'd expect, you'll see a TV that's watching a TV, which is itself watching a TV and so on. It looks a little curved, because there's a (very) slight delay in the signal reaching the camera. Also, the image is less sharp as you go down the chain, because the camera digitises the information, meaning it isn't quite an exact copy.

Now do something exciting to perturb the system: wave your hand in front, or strike a match. What you get is a continual feedback loop, where the camera is watching a picture, digitising the information and sending it back to the screen... but by the time it does, the delay means that at the next instant it's not *quite* watching the same picture. Because of the tiny errors, what you'll see will be a magnificent looping effect as the picture continually updates itself. (An easier way to see this is to watch the opening credits of old *Doctor Who* episodes, which were created in exactly this way.)

What you're seeing is a bifurcation from a stable equilibrium (the TVs all watching themselves) to a chaotic system (the *Doctor Who* effect). Bifurcations are all around us, and it's a pity they're not better known outside mathematics.

In last chapter's lab, we saw some other difference equations and the way their equilibria and periodic orbits changed stability as their parameters changed. In this chapter's lab, we'll create a program that will draw the bifurcation diagrams for us, at least for the stable orbits.

11.3 Lab work

The following program is designed to plot the infected population after a large number of iterations (equal to the final population value(s) for a certain range of r values) versus the parameter r over a user-specified range of r .

```
%Bifurcation in the discrete logistic equation
clear all
rmin=input('Input the minimum value of r: ');
rmax=input('Input the maximum value of r: ');
n=850;
N=100;
for m=1:n
    r=rmin+(rmax-rmin)*m/n;
    x=0.5;
    for k=1:N
        x=r.*x.*(1-x);
```

```

end
X(1)=x;
R(1)=r;
for k=1:N
    X(k+1)=r.*X(k).*(1-X(k));
    R(k+1)=r;
end
plot(R,X,'.')
hold on
end
title('Bifurcation plot')
xlabel('Bifurcation parameter r')
ylabel('x values')

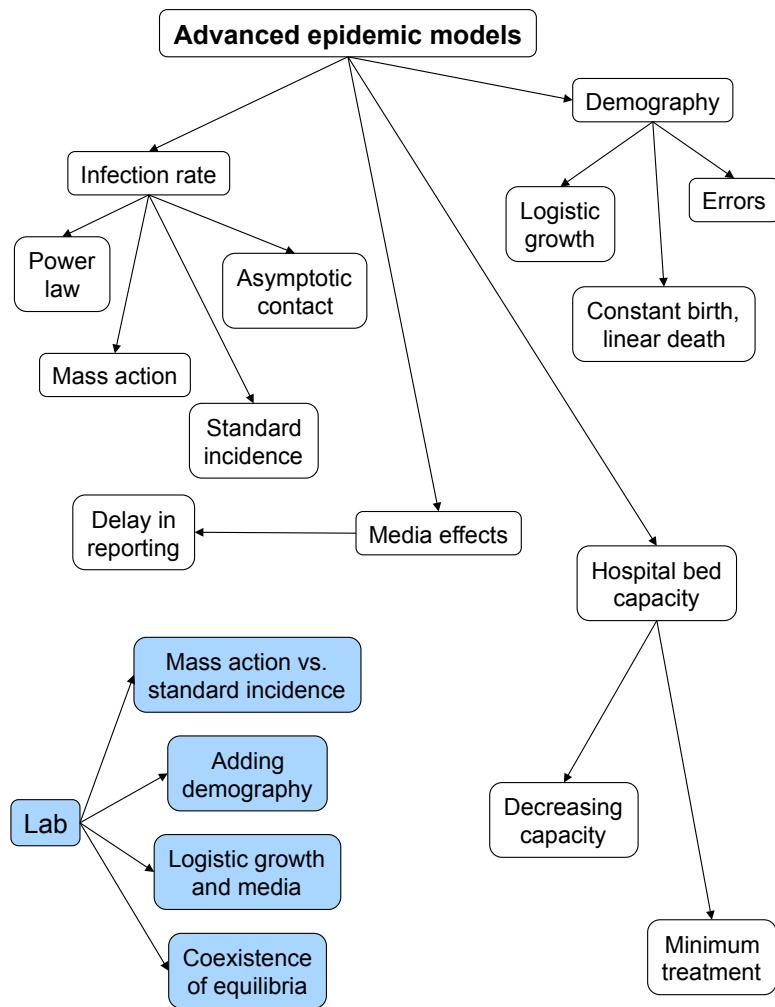
```

11.3.1 Exercises

1. Run the bifurcation program for the range $0 < r < 4$.
2. Now that you have run the program once, you may use the `axis` command to “zoom in” on a region. Keep in mind that the detail will be poorer as you zoom in further.
3. Rerun the program for the smaller r range. This will take a lot longer, but will give you better detail within that range.
4. Can you use this to find the r values to produce 3 and 5 spikes in the histogram program (if you didn’t already)? How about 7 spikes?
5. Consider the difference equation $x_{n+1} = \lambda \sin(\pi x_n)$. Use your program to draw the bifurcation diagram on the range $0 < \lambda < 1$.
6. Consider the difference equation $x_{n+1} = \alpha \sin^2 x_n$. Use your program to draw the bifurcation diagram for $\alpha > 0$. What’s happening here?

Part II

Advanced material



12

More advanced epidemic models

This chapter deals with a range of different options for creating more advanced models, depending on the circumstances. It examines choices that can be made for the infection term, for demography and inclusion of factors like the effects of media or available hospital beds. Throughout this chapter, we're going to focus on the differential equation for the susceptibles, with the understanding that there's a corresponding equation for infected individuals.

By the end of this chapter, you should understand some of the choices available for tweaking the basic models and adding in more advanced options.

12.1 The infection rate

In Chapter 5, we used the infection term

$$S' = -\beta SI,$$

on the grounds that there was no transmission if $S = 0$ or $I = 0$. Which is a condition we'd like to have, but there are other formulations we could use that would still have this property. In mass-action transmission, the parameter β is the average ratio of infective contact per infected individual per unit time. It has units of [people⁻¹·time⁻¹]. This is also known as *density-dependent transmission*. For a disease like the flu, the more people around you who are infected, the more likely you are to catch it. So the contact rate depends on the density of people around you.

Another form of infection is

$$S' = -\frac{cSI}{S + I}.$$

This is called *standard incidence*. In this case, c represents the number of contacts and the probability of becoming infected. It has units of [1/time]. Note that if the population were constant, so that $S + I = N$, then we could write

$$S' = -\frac{cSI}{N} = -\beta SI$$

with $\beta = c/N$. That is, the standard incidence is a general case of mass-action transmission. Of course, populations aren't usually constant when you have births and deaths, but it's an approximation that can be justified in some circumstances.

This is also known as *frequency-dependent transmission*. For a disease like HIV, it doesn't matter how many people around you have HIV, it only matters how many of them you have sex with. So the contact rate depends on the frequency of contacts.

One difference between the two is that the former has no bound if the population gets large, while the latter flattens out to a maximum as the population increases. Mass action is simpler, however, so when taking derivatives (to calculate the Jacobian, for instance), it's a lot easier to work with. Which one you choose depends on the nature of the disease and the size of the population you're dealing with.

They aren't the only choices, of course. Another option that still has no transmission when $S = 0$ or $I = 0$ is

$$S' = -\beta S^p I^q$$

where p and q are two further parameters that we might have some control over. For example, if it takes two zombies to infect a single human, then we might choose $p = 1$ and $q = 2$. This is called a *power relationship* and is an obvious generalisation of mass-action transmission.

Another possibility is

$$S' = -\frac{\beta S^p I^q}{1 - \epsilon + \epsilon(S + I)},$$

where ϵ ranges between 0 and 1. This is called *asymptotic contact* and generalises all of the above. If $\epsilon = 0$ then we have the power relationship. If $\epsilon = 1$ then we have the generalised version of standard incidence. Note that this assumes that S and I have the same units as 1 and ϵ (i.e., no units). So this works when populations are proportions, not individuals.

As you can see, the possibilities are potentially limitless. Of course, your choice of model depends on the biology, and there are advanced techniques to determine which type of model might be best in a given situation. However, there's also the issue of mathematical tractability: unless you're running a purely numerical analysis, you also need to be able to analyse these models. (And a purely numerical investigation can miss things.) Even simple models can give rise to very complicated behaviour. We don't necessarily need the perfect model to understand possible outcomes; that is, sometimes we are happy to trade accuracy for insight.

12.2 Demography

One issue that often arises in disease modelling is to determine what's happening to the demographics of the population, independent of the disease. This issue comes up a lot in ecological modelling, as you can imagine. However, even if our main focus is the infection, it's still important to include the background dynamics of birth and death.

You may be unsurprised to learn that there are several ways to do this. In Chapter 6, we saw perhaps the most straightforward way to do this: constant birth and proportional death. Let's generalise this slightly by writing

$$S' = \pi(S) - dS,$$

where $\pi(S)$ is a growth function and d is the background death rate. The latter is quite a useful formulation, because it turns out that this means individuals will be alive for $1/d$ time units. It's not the only way to include the death rate, but it is the easiest.

Except... why not a constant death rate? Surely something like

$$S' = \pi(S) - \tilde{d}$$

would be a simpler way to describe background death: individuals simply leave at a constant rate with units of [people/time]. The problem is, having a negative in a differential equation runs the risk of making the population negative. (We saw this issue in the exercises in Chapter 5 in fact!) Specifically, if the population is initially zero (so that $S(0) = 0$) and $\pi(0) = 0$ (so if there's nobody in the population, there'll be no growth, which is reasonable), then the equation is

$$S'(0) = -\tilde{d} < 0.$$

That is, initially there's nobody in the population (which is fine) and the population is decreasing, meaning that shortly afterwards there will be negative people (which isn't fine at all). Obviously, any model that allows for negative populations isn't a good one. To put it another way, we'd like to avoid negative people \odot . So we don't want to have negative terms that don't allow those terms to be zero when the population is zero. It follows that the simplest form of the background death rate is the term $-dS$.

Okay, so what about the growth rate? In this case, we don't have any concerns about populations being negative, so we can let $\pi(S)$ be a constant if we like. This is what we did in Chapter 6, so we would have something like

$$S' = r - dS, \tag{12.1}$$

where r is the recruitment rate of new people per unit time. It represents constant birth or immigration (or both) and has units of [people/time]. This

is a pretty easy equation. It's linear, so we could solve it without too much trouble. Or we could look at the equilibrium values:

$$S' = 0 \Rightarrow r - d\bar{S} = 0$$

$$\bar{S} = \frac{r}{d}.$$

That is, populations will equilibrate at the value $\frac{r}{d}$. Note that if $S(0) < \frac{r}{d}$, then $S' > 0$, whereas if $S(0) > \frac{r}{d}$, then $S' < 0$. So solutions that start above $\frac{r}{d}$ will be decreasing, while solutions that start below $\frac{r}{d}$ will be increasing. For this simple model, all solutions approach the equilibrium. So the outcome looks like Figure 12.1.

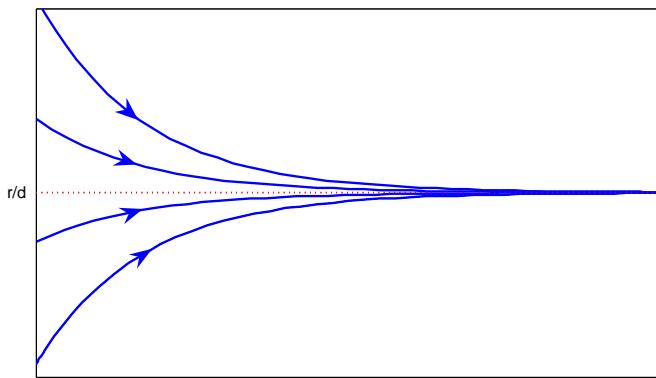


Fig. 12.1. Constant growth, linear death.

Another option is *logistic growth*, with

$$\pi(S) = rS \left(1 - \frac{S}{K}\right).$$

In fact, since there's already a linear factor rS , we can absorb the death rate into this and simply have

$$S' = rS \left(1 - \frac{S}{K}\right).$$

This has the property that there is no change if $S = 0$ (so no growth if there are no individuals) or $S = K$ (where the population is at equilibrium). We call K the *carrying capacity*. Since $S' < 0$ for $S > K$ and $S' > 0$ for $0 < S < K$, it follows that solutions below the carrying capacity will increase, while those above will decrease. It turns out that all nonzero solutions will converge to the carrying capacity, so solutions look like Figure 12.2.

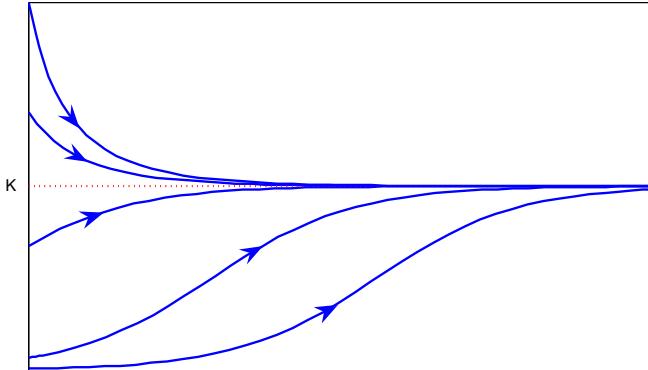


Fig. 12.2. Logistic growth.

The difference between Figures 12.1 and 12.2 is that solutions that start with sufficiently small initial conditions are much “flatter” at first. That is, they persist at low levels for some time before heading towards the equilibrium. They get there in the end, but it’s a slower process. More formally, there’s a point of inflection that changes some of the solution curves from concave up to concave down. The value of $\frac{r}{d}$ in Figure 12.1 is like the carrying capacity in logistic growth.

In the logistic growth example, we have nonlinear growth and linear death (whether explicitly written out or absorbed into the linear term). In the earlier case, we had constant birth and linear death. We saw that we couldn’t have constant death. But surely the next simplest version would be linear growth and linear death?

Unfortunately, problems arise here. Suppose we had something like

$$S' = rS - dS,$$

where rS represents linear growth and dS linear death. However, there’s an issue when both terms have the same order. We can rewrite this model as

$$S' = (r - d)S.$$

and in fact we can explicitly solve this one:

$$S(t) = S(0)e^{(r-d)t}.$$

If $r > d$, then solutions increase to infinity, which is impossible. If $r < d$, then the entire population dies out in the absence of disease, which isn’t very likely. But why not set $r = d$ and assume that the population is constant? It’s a nice idea, but unfortunately it means that we’re on a knife-edge case, where tiny fluctuations in the birth or death rate could result in catastrophic consequences.

Alas, a number of disease models make this assumption, with the population sometimes being held in check by the disease itself. This isn't very realistic though, so this should be avoided when creating models. In general, the birth and death terms can be whatever you want, so long as they're not the same order and so long as populations can't become negative.

12.3 The effects of media

In any pandemic, the media can kind of go crazy. Sometimes they panic, over-reporting on every suspected case. Sometimes they ignore the disease for much longer than they should. Nevertheless, the media is a powerful tool for affecting people's behaviour. It's not always a straightforward response though.

How can we build this into our models? We'll focus on the transmission rate, although it's not the only possibility. The transmission rate might change due to people mixing less as a result of the media reporting on the disease. There are several ways we might include this, but the basic idea is that the transmission rate should decrease as the number of infected individuals increases. In a pandemic, you still have to go to work, but you might not go to a hockey game. With less mixing of susceptibles and infecteds, the contact rate decreases, which is part of the transmission rate.

Having determined which form of demography we might want, let's return to the infection term. Using (say) logistic growth and standard incidence, our model might look like

$$S' = rS \left(1 - \frac{S}{K}\right) - \beta(I) \frac{SI}{S+I}. \quad (12.2)$$

This combines two things we'd previously seen, except that β is no longer a constant.

The most obvious function to consider is

$$\beta(I) = \beta_0 e^{-mI},$$

so that the transmission function decreases as the number of infected individuals increases, eventually approaching zero as the whole population becomes infected. See Figure 12.3.

However, the media response is not instantaneous. Information takes time to be released and reported. Yes, even in our 24-hour news cycle! So if it takes τ days for the health system to release numbers of infected people, then that will introduce a delay into the system. Our media-transmission function then becomes

$$\beta(I) = \beta_0 e^{-mI(t-\tau)}.$$

(Remember that I is a function of time, so β is as well.) See Figure 12.4.

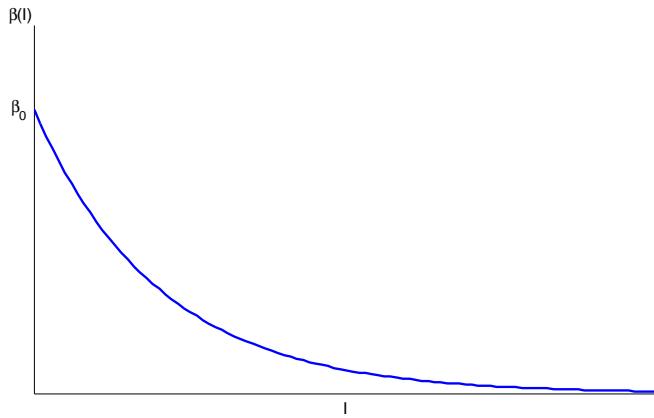


Fig. 12.3. The media-affected transmission function decreases as the number of infected individuals increases.

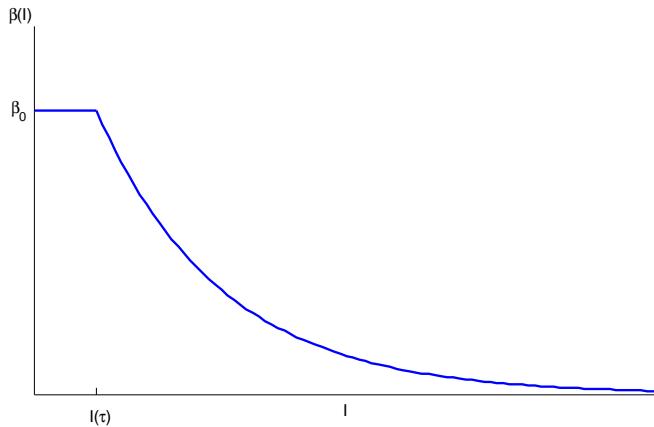


Fig. 12.4. After an initial delay, the media-affected transmission function decreases as the number of infected individuals increases.

12.4 Hospital bed capacity

Thus far, we've dealt with the infection and demographics, but there's another place where we might have some control over our equations: the recovery rate. Medical resources will determine the treatment and recovery rate, but resources may be limited. We'll consider the hospital bed capacity, which will have a maximum and a minimum. In a pandemic, the hospital may be inundated, but some minimum treatment will still occur. For diseases like swine flu (H1N1), a large number of people survived due to respirators in hospitals that essentially did the breathing for them while their lungs healed.

In the 1918 Spanish flu pandemic (also H1N1), such patients would have died. So there's a relationship between recovery and hospital bed capacity.

For this, we'll need to consider a full SIR model like

$$\begin{aligned} S' &= r - dS - \frac{\beta SI}{S + I + R} \\ I' &= \frac{\beta SI}{S + I + R} - dI - \nu I - \mu(b, I)I \\ R' &= \mu(b, I)I - dR. \end{aligned} \tag{12.3}$$

Here we're using standard incidence with a denominator involving all populations (since we've added in the R class). There's a background death rate d in all classes and a disease-specific death rate ν in the infected class.

The recovery rate is μ , which is often considered constant. However, here we will formulate the recovery rate in such a way that it incorporates the effect of the capacity and limited resources of the healthcare system, as defined by the parameter $b > 0$ representing the hospital bed capacity. Conditions that we want are:

- μ is positive for $b > 0$.
- When $I = 0$, μ is constant, so that $\mu(b, 0) = \mu_1 > 0$, where μ_1 is the maximum per capita recovery rate due to sufficient healthcare resources.
- As I increases, μ should decrease. That is, the more infected people overwhelm the hospital, the slower treatment and hence recovery will be.
- As the number of infections increase, the available resources cannot meet the demand, but some individuals will still get treated and recover. Thus μ has a lower bound μ_0 such that $\lim_{I \rightarrow \infty} \mu(b, I) = \mu_0 > 0$.
- As b increases, μ should increase. That is, as more beds become available, recovery should be higher.

The simplest function that satisfies these requirements is

$$\mu(b, I) = \mu_0 + (\mu_1 - \mu_0) \frac{b}{I + b}.$$

See Figures 12.5 and 12.6.

12.5 Lab work

The problem

Is there a difference between mass-action infection and standard incidence without demography?

The solution

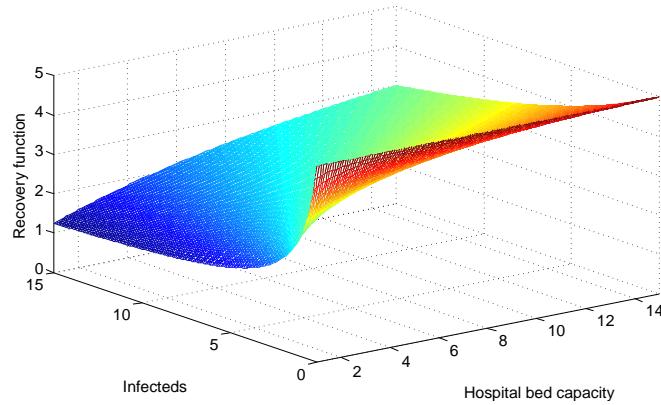


Fig. 12.5. The recovery rate $\mu(b, I)$ with $\mu_0 = 1$ and $\mu_1 = 5$.

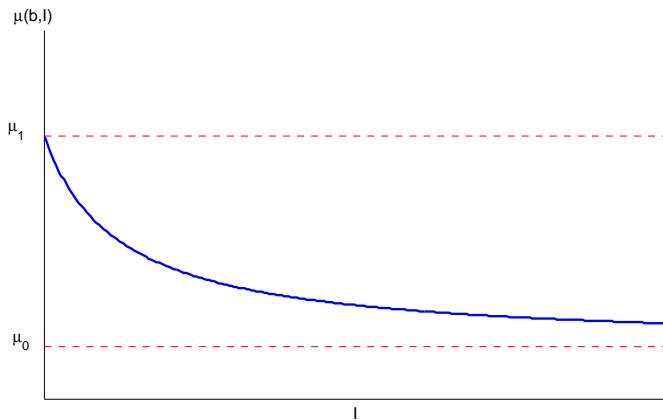


Fig. 12.6. The recovery rate as a function of I for a given number of hospital beds b .

First let's set up our m-file. We need initial conditions, so let's suppose there are five susceptible individuals and 1 infected individual. We'll run the output for 3 time units.

```

clear all
t0=0;
tau=3;
x0=[5 1];
tspan=[t0 t0+tau];
[t,x]=ode45(@standardincf,tspan,x0);

```

```
plot(t,x(:,1),t,x(:,2), '--')
xlabel('time')
ylabel('population')
```

Now for our function file. We'll do standard incidence first. The parameter values don't really matter, so let's randomly choose $c = 3$. The code for this is thus:

```
function xprime=standardincf(t,x)
c=3;
xprime(1,:)=-c.*x(1).*x(2)./(x(1)+x(2));
xprime(2,:)=c.*x(1).*x(2)./(x(1)+x(2));
```

Next, we'll change the equations for mass-action transmission. We could write a new function file with a new name and consequently adjust the m-file... or we could just comment out the equations we used and put in the revised equations. This saves us having to rename everything. We'll keep the name c but we need an N for $\beta = c/N$ to have a comparison. Since we had six individuals, let's choose $N = 6$. So the revised function file is:

```
function xprime=standardincf(t,x)
c=3;
N=6;
xprime(1,:)=-c.*x(1).*x(2)./(N);
xprime(2,:)=c.*x(1).*x(2)./(N);
%xprime(1,:)=-c.*x(1).*x(2)./(x(1)+x(2));
%xprime(2,:)=c.*x(1).*x(2)./(x(1)+x(2));
```

The results are shown in Figure 12.7. (Note that the labels are drawn in by hand, using the “Plot Tools” interface and then “Textarrow” from the “Insert” menu.)

Do you notice any difference? They look pretty similar to the naked eye, but how would we know for sure? One way is to look at the output at the end of the simulation. If we write $x(\text{length}(x), :)$, this will give us the final row of the x values.

For mass action, we find $\text{ans} = 0.0037\ 5.9963$. Note that the sum is exactly 6. In fact, if we write $x(\text{length}(x), 1)+x(\text{length}(x), 2)$, then the answer is 6.

For standard incidence, we find $\text{ans} = 0.0037\ 5.9963$. However, if we write $x(\text{length}(x), 1)+x(\text{length}(x), 2)$, then the answer is 6.000. What's going on here? One is exact and the other is a decimal, but with zeros. This suggests that the answer isn't precisely 6 but rather something very

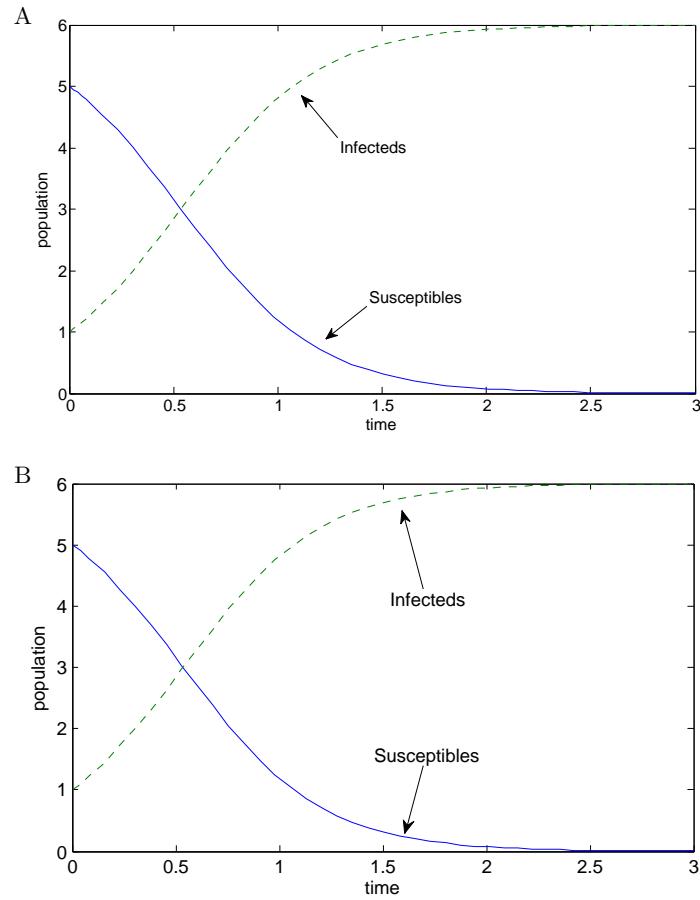


Fig. 12.7. A. The output for standard incidence. B. The output for mass action.

close to it. So if we write `x(length(x),1)+x(length(x),2)-6`, then we find `ans=-8.8818e-16`.

So do they produce the same output or not? Well, technically no... but if you have two real-world processes with output that's within 10^{-15} of each other, then they might as well be identical.

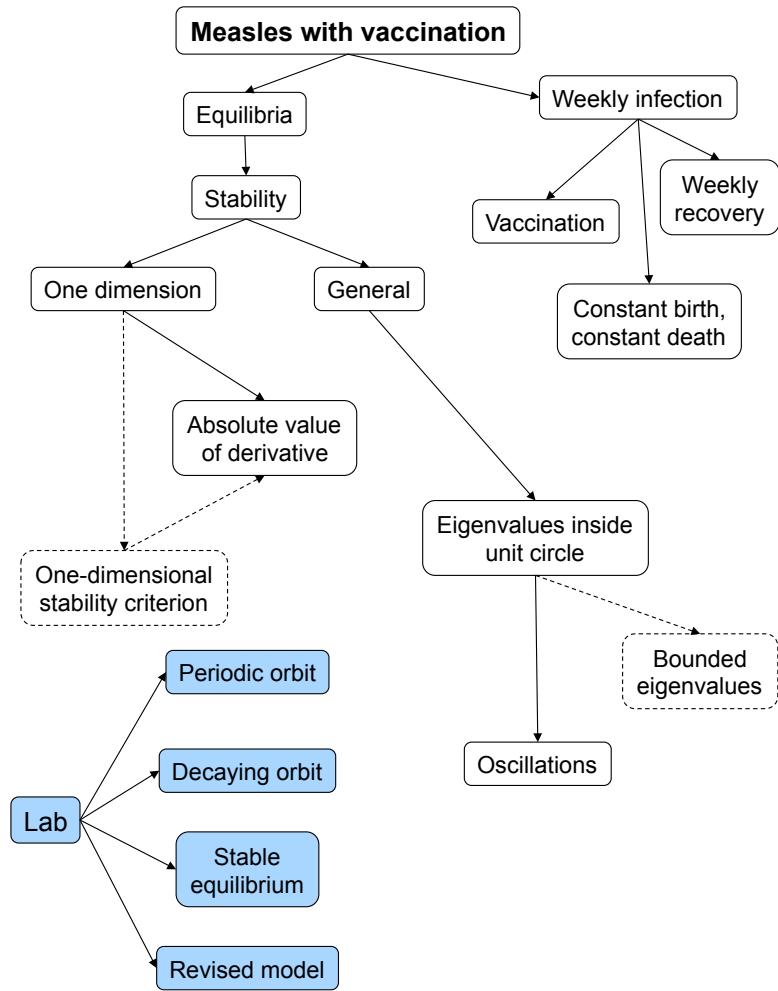
There's a simpler way to approach this, however. Without demography, $S' + I' = 0$ in both cases, so the population $N = S + I$ must be a constant. Hence

$$\frac{c}{S + I} = \frac{c}{N} = \beta,$$

so the two sets of differential equations really are equal. It's interesting that Matlab almost but didn't quite tell us that. This is the kind of error we should expect to live with.

12.5.1 Exercises

1. Suppose the demography follows the form of equation (12.1) with $r = 5$ and $d = 0.1$ in the S' equation, with the same death rate (but no births) in the I' equation.
 - a) Adjust your programs to find $x(\text{length}(x), :)$ for both mass action and standard incidence. What do you conclude now?
 - b) Rerun your programs for 30 time units. What is this telling you?
 - c) What if $r = 12$ and $d = 2$? Why does this produce what it does?
 - d) Now repeat the process with logistic growth instead, with the same parameter values and $K = 50$. What's happening now? Try running for longer (e.g., 300 time units) to see the full effect.
2. Adjust your program to model an infection with the form of equation (12.2).
 - a) Run the program for $m = 0.1$ and $\beta_0 = 1$.
 - b) Plot the two functions $\beta(I)$ on the same graph.
 - c) What's the effect of including the $-dS$ term in the S' equation versus leaving it out?
3. Consider system (12.3).
 - a) Calculate the disease-free equilibrium.
 - b) Calculate R_0 . (Hint: you may have to bound $\mu(b, I)$.)
 - c) Show that there are potentially two endemic equilibria. (You don't have to find them explicitly to do this.)
 - d) If two endemic equilibria coexist, what can you conclude about their stability?



13

Measles with vaccination

As seen in Chapter 10, when dealing with discrete time, there can be unexpected results. There we dealt with only one-dimensional models, so here we'll investigate a higher-order model. The key to these models is having an underlying timestep that drive the dynamics.

By the end of this chapter, you should be comfortable with discrete-time modelling, both in construction and analysis. You should be able to appreciate some of the similarities and some of the differences between discrete-time and continuous models.

13.1 The model

Suppose we have a constant number of births and deaths B , so that the number of births equals the number of deaths each week. This isn't necessarily true in the long term, but on a weekly basis it's an okay assumption.

Individuals recover from measles within a week, so there are only new measles cases each week. Newborns are born susceptible, while everyone catches the measles (or gets vaccinated), so only recovered individuals die.

Because time is discrete, we can use difference equations instead of ODEs. This means that, instead of having a derivative, we will have update each class at each timestep. Of course, many updates don't change the state of that variable, so if there were susceptibles last week, then there will be susceptibles this week. The basic model is an SIR model, as shown in Figure 13.1.



Fig. 13.1. The progression of measles.

The model is given by

$$S_{t+1} = S_t - \alpha I_t S_t + B$$

$$I_{t+1} = \alpha I_t S_t$$

$$R_{t+1} = R_t + I_t - B.$$

↑

Anyone infected last week is now recovered.

What changes each week are that measles can infect you and you then recover the following week, which is why the solo I_t term doesn't appear in the second equation but instead moves individuals to the recovered class at the next timestep. You might also be born (as a susceptible) or die (as a recovered individual, which is too bad for you, but at least you survived the measles). We're ignoring death due to disease, which is a reasonable assumption in the developed world, although it might not be in the developing world.

Key differences between this and ODE models are:

- birth and death are both constant, not proportional to population size
- the infectious period is the same length as the time step.

Let's add vaccination, which takes some susceptibles directly to the recovered class. If a proportion p of susceptibles are vaccinated each week, then the model becomes

$$S_{t+1} = (1 - p)S_t - \alpha I_t S_t + B$$

$$I_{t+1} = \alpha I_t S_t$$

$$R_{t+1} = R_t + I_t - B + pS_t.$$

See Figure 13.2.

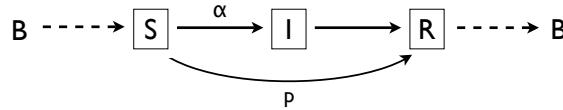


Fig. 13.2. The addition of vaccination.

13.2 Finding equilibria

First, let's look at the whole population. If we set $N_t = S_t + I_t + R_t$, then, adding the equations together, we have

$$N_{t+1} = S_t + I_t + R_t = N_t.$$

That is, we've *deduced* that the population size remains constant over time. Hence we could write $N_t = N$. Furthermore, S_t and I_t do not depend on R_t . That is, the third equation decouples from the model. We can thus look at S_{t+1} and I_{t+1} only.

Equilibria must satisfy $S_{t+1} = S_t$ and $I_{t+1} = I_t$. That is, there should be no change in time. We could call these fixed values (if they exist) \bar{S} and \bar{I} . Plugging them into the model, we have

$$\begin{aligned}\bar{S} &= (1-p)\bar{S} - \alpha\bar{I}\bar{S} + B \\ \bar{I} &= \alpha\bar{I}\bar{S}.\end{aligned}$$

From the second equation, either $\bar{I} = 0$ or $\bar{S} = \frac{1}{\alpha}$. So let's look at each case.

$$\begin{aligned}\bar{I} &= 0 & \bar{S} &= \frac{1}{\alpha} \\ \bar{S} &= \bar{S} - p\bar{S} + B & \frac{1}{\alpha} &= (1-p)\frac{1}{\alpha} - \bar{I} + B \\ p\bar{S} &= B & \bar{I} &= B - \frac{p}{\alpha}, \quad \text{if } B - \frac{p}{\alpha} > 0. \\ \bar{S} &= \frac{B}{p}.\end{aligned}$$

Hence the equilibria are

$$(\bar{S}, \bar{I}) = \left(\frac{B}{p}, 0 \right), \left(\frac{1}{\alpha}, B - \frac{p}{\alpha} \right).$$

Note that if there's no vaccination ($p = 0$), then there's no disease-free equilibrium. (If there's no disease, then the system blows up, because the death assumption assumes that everyone passes through to the recovered stage.) The endemic equilibrium only exists if $B - \frac{p}{\alpha} > 0$.

13.3 Stability in a discrete model

We can find the Jacobian quite easily:

$$\begin{aligned}J_p &= \begin{bmatrix} 1-p-\alpha I_t & -\alpha S_t \\ \alpha I_t & \alpha S_t \end{bmatrix} & p > 0 \\ J_0 &= \begin{bmatrix} 1-\alpha I_t & -\alpha S_t \\ \alpha I_t & \alpha S_t \end{bmatrix} & p = 0.\end{aligned}$$

The linearisation is thus

$$\begin{pmatrix} S_{t+1} \\ I_{t+1} \end{pmatrix} = J_i \begin{pmatrix} S_t \\ I_t \end{pmatrix} \text{ for } i = 0, p.$$

Although there's a Jacobian matrix, the conclusions aren't quite the same as in the continuous case. This is a two-dimensional system, but let's first figure out how stability works for a one-dimensional discrete-time system.

13.3.1 Stability in one dimension

An equilibrium \bar{x} of $x_{t+1} = f(x_t)$, where f and f' are continuous, is (locally) stable if $|f'(\bar{x})| < 1$ and unstable if $|f'(\bar{x})| > 1$.

See Appendix J for details.

Let's unpack this a little. First, note that, although we are dealing with a discrete system, the function f still needs to be continuous, because we need its derivative. That is, $\frac{df}{dx}$ has to exist and be defined (at least near \bar{x}), which means we need to be able to differentiate the function. Continuity is a necessary requirement for differentiability.

Second, in a one-dimensional system, $f'(\bar{x})$ is just the (single) eigenvalue of the (1×1) Jacobian matrix $f'(\bar{x})$. That is, the eigenvalue and the Jacobian matrix are the same, because we only have one dimension, and a 1×1 matrix is just a number. We'll generalise this shortly.

Third, the condition tells us what happens if the absolute value is larger or smaller than 1, but we don't know what happens if it's equal to 1. It turns out that the results there aren't predictable... but they weren't predictable when the real part of λ was zero in the continuous case either. (See Page 66.)

Finally, there's nothing to specify that $f'(\bar{x})$ has to be real, because we have a perfectly well defined absolute value for complex numbers as well. (We can't have a complex eigenvalue without its conjugate, so we don't get complex solutions in a one-dimensional system, but we will in higher-order systems.) What this condition really means is we get stability when the eigenvalues are within the unit circle. Compare this to the continuous case, where stability occurred when the eigenvalues were in the left half of the plane.

13.3.2 Stability in two dimensions

By analogy, the equilibrium $\bar{\vec{x}}$ of $\vec{x}_{t+1} = \vec{f}(\vec{x}_t)$ is stable if all eigenvalues of the Jacobian matrix have absolute value less than 1 and unstable if there is an eigenvalue with absolute value greater than 1.

First let's deal with $p = 0$. The Jacobian at the endemic equilibrium is

$$\begin{aligned}
J_0 \left(\frac{1}{\alpha}, B \right) &= \begin{bmatrix} 1 - \alpha B & -1 \\ \alpha B & 1 \end{bmatrix} \\
\det \left(J_0 \left(\frac{1}{\alpha}, B \right) - \lambda I \right) &= \det \begin{bmatrix} 1 - \alpha B - \lambda & -1 \\ \alpha B & 1 - \lambda \end{bmatrix} \\
&= (1 - \alpha B - \lambda)(1 - \lambda) + \alpha B \\
&= \lambda^2 - (2 - \alpha B)\lambda + 1 \\
\lambda_{1,2} &= \frac{2 - \alpha B \pm \sqrt{(2 - \alpha B)^2 - 4}}{2} \\
&= \frac{2 - \alpha B \pm \sqrt{\alpha B(\alpha B - 4)}}{2}.
\end{aligned}$$

There are two cases here.

(i) If $\alpha B > 4$, then

$$\begin{aligned}
\lambda_2 &= \frac{2 - \alpha B - \sqrt{\alpha B(\alpha B - 4)}}{2} \\
&< \frac{2 - 4 - \sqrt{\alpha B(\alpha B - 4)}}{2} \\
&= -1 - \frac{\sqrt{\alpha B(\alpha B - 4)}}{2}
\end{aligned}$$

and the part under the square root is positive. It follows that $|\lambda_2| > 1$, and hence the equilibrium is unstable. (Note that we don't need to check λ_1 , because instability only requires a single equilibrium with absolute value larger than 1. Nothing that λ_1 can do could change the stability.)

(ii) If $\alpha B < 4$, then we will be taking the square root of a negative. However, this is no problem, because the absolute value of a complex number $z = u + iv$ is $|z| = \sqrt{u^2 + v^2}$. Hence

$$\begin{aligned}
|\lambda_{1,2}| &= \sqrt{\left(\frac{2 - \alpha B}{2} \right)^2 + \frac{\alpha B(\alpha B - 4)}{4}} \\
&= \sqrt{\frac{4 - 4\alpha B + (\alpha B)^2 + 4\alpha B - (\alpha B)^2}{4}} \\
&= 1.
\end{aligned}$$

Therefore this equilibrium is not asymptotically stable. If $\alpha B < 4$, then, because of the complex eigenvalues, solutions oscillate around the endemic equilibrium. They may oscillate towards it, away from it or stay nearby. We can't tell without more sophisticated methods.

Next, let's deal with $p \neq 0$. The Jacobian at the disease-free equilibrium is

$$J_p \left(\frac{B}{p}, 0 \right) = \begin{bmatrix} 1 - p - \alpha B/p & -\alpha B/p \\ 0 & \alpha B/p \end{bmatrix}.$$

The eigenvalues here are

$$\lambda_{3,4} = 1 - p, \frac{\alpha B}{p}.$$

How did we get these so quickly? Answer: because the matrix is upper triangular (i.e., only zeroes below the diagonal), which means the eigenvalues are the diagonal elements. Note: this isn't true in general, of course. Or you can use the formula from Appendix C.

Since $0 < p < 1$, $|\lambda_3| < 1$. So the disease-free equilibrium is stable if

$$|\lambda_4| = \left| \frac{\alpha B}{p} \right| < 1$$

i.e., if $\alpha B < p$,

$$\text{and we can define } R_0 = \frac{\alpha B}{p}.$$

The endemic equilibrium $\left(\frac{1}{\alpha}, B - \frac{p}{\alpha}\right)$ only exists if $R_0 > 1$; i.e., if $\alpha B > p$. The Jacobian at this equilibrium is

$$\begin{aligned} J_p \left(\frac{1}{\alpha}, B - \frac{p}{\alpha} \right) &= \begin{bmatrix} 1 - p - \alpha \left(B - \frac{p}{\alpha} \right) - \alpha \left(\frac{1}{\alpha} \right) \\ \alpha \left(B - \frac{p}{\alpha} \right) \alpha \left(\frac{1}{\alpha} \right) \end{bmatrix} \\ &= \begin{bmatrix} 1 - \alpha B - 1 \\ \alpha B - p \ 1 \end{bmatrix} \\ \det(J_p - \lambda I) &= \det \begin{bmatrix} 1 - \alpha B - \lambda & -1 \\ \alpha B - p & 1 - \lambda \end{bmatrix} \\ &= (1 - \alpha B - \lambda)(1 - \lambda) + \alpha B - p \\ &= \lambda^2 - (2 - \alpha B) + 1 - p \\ \lambda_{5,6} &= \frac{2 - \alpha B \pm \sqrt{\Delta}}{2} \end{aligned}$$

where

$$\begin{aligned} \Delta &= (2 - \alpha B)^2 - 4(1 - p) \\ &= \alpha^2 B^2 - 4\alpha B + 4p. \end{aligned}$$

Once again, we have two cases.

(i) If $\Delta \geq 0$, then the roots are real. The only restriction on αB is that it is greater than p . First we have

$$\begin{aligned} \lambda_5 &= \frac{2 - \alpha B + \sqrt{(\alpha B)^2 - 4\alpha B + 4p}}{2} \\ &< \frac{2 - \alpha B + \sqrt{(\alpha B)^2}}{2} \\ &= 1. \end{aligned}$$

We'll show that $\lambda_5 > -1$ in Appendix K. Remember that, for an absolute value, there are two things to check.

Understanding what happens with λ_6 is more challenging. Let's try a few values. If $\alpha B = 1$, then

$$\begin{aligned}\lambda_6 &= \frac{2 - 1 - \sqrt{1 - 4(1) + 4p}}{2} \\ &= \frac{1 - \sqrt{-3 + 4p}}{2}.\end{aligned}$$

For the range of p that gives real roots (e.g., p slightly smaller than 1), we have $|\lambda_6| < 1$ and hence stability.

Conversely, if $\alpha B = 4$, then

$$\begin{aligned}\lambda_6 &= \frac{2 - 4 - \sqrt{16 - 4(4) + 4p}}{2} \\ &= \frac{-2 - \sqrt{4p}}{2} \\ &= -1 - \sqrt{p} < -1.\end{aligned}$$

Hence $|\lambda_6| > 1$.

Since it's possible for λ_6 to be both smaller than or greater than 1, we can conclude that the equilibrium is sometimes stable and sometimes unstable. (Note: We weren't able to show this in general, but testing two values was enough to conclude that instability is at least possible. We'll explore this case in more detail in the lab.)

(ii) If $\Delta < 0$, then we have complex roots, so

$$\begin{aligned}|\lambda_{5,6}| &= \sqrt{\left(\frac{2 - \alpha B}{2}\right)^2 + \frac{-\Delta}{4}} \\ &= \sqrt{\frac{4 - 4\alpha B + \alpha^2 B^2 - \alpha^2 B^2 + 4\alpha B - 4p}{4}} \\ &= \sqrt{\frac{4 - 4p}{4}} \\ &= \sqrt{1 - p} \\ &< 1.\end{aligned}$$

Hence the endemic equilibrium is stable whenever $R_0 > 1$ and complex roots arise.

13.4 Lab work

The problem

Write a program to model the discrete-time system for measles vaccination for the cases $p = 0$, $\alpha B > p$ (with $\Delta < 0$) and $\alpha B < p$.

The solution

First we'll need to pick some parameters. Let's choose $B = 115$ and $\alpha = 0.3 \times 10^{-4}$, so $\Delta = (\alpha B)^2 - 4\alpha\beta + 4p = -0.0138 + 4p$. We'll need small values of p (which we can vary) to keep this negative. We also need some initial conditions, so we'll choose $S(0) = 3 \times 10^4$, $I(0) = 200$ and $R(0) = 0$. Unlike ODEs, discrete-time systems are really easy to program in Matlab, because they're just stepping discretely through time, which is basically exactly what Matlab does. So a simple `for` loop will suffice.

```

clear all
S=3e4;
I=200;
R=0;
%p=0;
p=0.002;
%p=0.005;
B=115;
alpha=0.3e-4;
r=1000;
for n=1:r-1
    S(n+1)=(1-p).*S(n)-alpha.*I(n).*S(n)+B;
    I(n+1)=alpha.*I(n).*S(n);
    R(n+1)=R(n)+I(n)-B+p.*S(n);
end
subplot(1,3,1)
plot([1:r],S,[1:r],[B./p B./p], '--r')
xlabel('Time in Weeks')
ylabel('Susceptibles')
subplot(1,3,2)
plot(I)
xlabel('Time in Weeks')
ylabel('Measles Cases')
subplot(1,3,3)
plot(I,S)
xlabel('Measles Cases')
ylabel('Susceptibles')

```

Here we've set up various values for p that can be adjusted by commenting in or out. The three plots illustrate a) the time series for susceptibles, b) the time series for infected individuals and c) the phase portrait.

For $p = 0$, we see a periodic orbit. Solutions oscillate around the endemic equilibrium, which is neither asymptotically stable nor unstable, as predicted by λ_2 (since $\alpha B < 4$ in this case). There are large, recurring outbreaks of measles. See Figure 13.3.

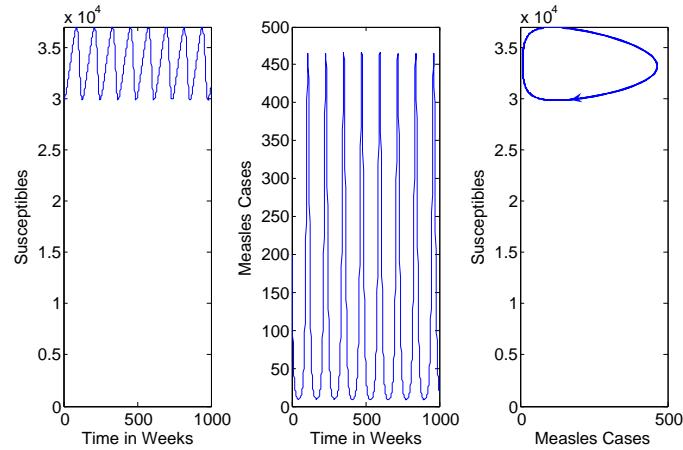


Fig. 13.3. Solutions oscillate around the endemic equilibrium.

For $p = 0.002$ (so $\alpha B > p$), we see a decaying orbit. That is, solutions oscillate towards the endemic equilibrium. There are still recurring peaks but at lower numbers than without vaccination. See Figure 13.4.

For $p = 0.005$ (so $\alpha B < p$), solutions do not oscillate but instead move directly to the disease-free equilibrium, since $|\lambda_4| < 1$. There is no endemic equilibrium. See Figure 13.5.

13.4.1 Exercises

1. What kinds of behaviour can occur in the case when $\Delta \geq 0$ and $\alpha B > p$? If both the disease-free equilibrium and the endemic equilibrium are unstable, what happens? Adjust your parameters from your program to explore this.
2. A more realistic epidemic model for measles vaccination might assume that the birth and death rates are proportional to the size of the class and that some people are born infected (although no one is born with immunity). We can also add other factors, such as a recovery rate that isn't in lock with the timestep and standard incidence. Such a model has the form

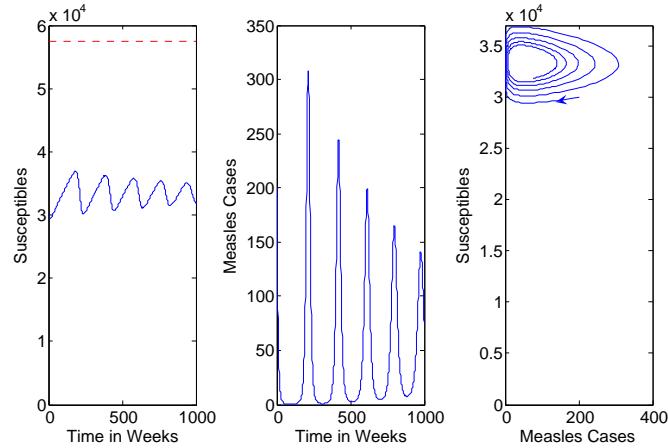


Fig. 13.4. Solutions oscillate towards the endemic equilibrium. The horizontal line indicates $S = B/p$, the disease-free equilibrium.

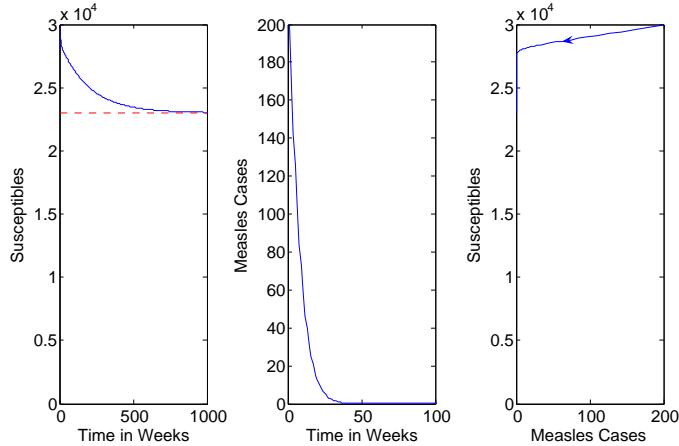


Fig. 13.5. Solutions move directly towards the disease-free equilibrium, indicated by the horizontal line. Note the shortened timescale for the measles time series.

$$\begin{aligned}
 S_{t+1} &= (1-p)S_t - \frac{\beta}{N} I_t S_t + b(R_t + I_t) \\
 I_{t+1} &= \frac{\beta}{N} I_t S_t + (1-b-\gamma)I_t \\
 R_{t+1} &= (1-b)R_t + \gamma I_t + pS_t,
 \end{aligned}$$

where $p, b, \gamma, \beta > 0$ and $\beta = \frac{\alpha}{N}$ (N is the total population). The parameter b is the per-capita number of births, γ is the rate of recovery, p is the

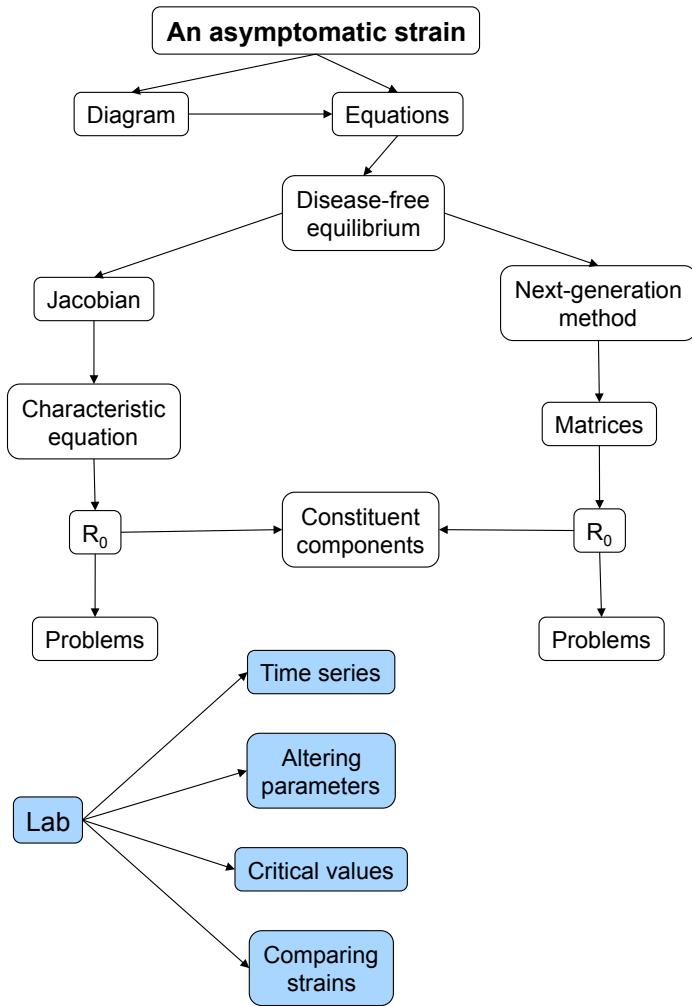
proportion of people vaccinated and β is the number of successful contacts in time t to $t + 1$. Suppose that $0 < p + \beta < 1$ and $0 < b + \gamma < 1$.

- a) Show that if $S_0 + I_0 + R_0 = N$, then $S_t + I_t + R_t = N$.
- b) Show that if $S_0, I_0, R_0 > 0$ and $S_0 + I_0 + R_0 = N$, then solutions are positive for all time; i.e., $S_t, I_t, R_t > 0$ for all t .
- c) Find the disease-free equilibrium $(\bar{S}, \bar{I}, \bar{R})$.
- d) Show that the basic reproductive ratio is

$$R_0 = \frac{\beta b}{(b + \gamma)(b + p)}.$$

(Hint: find the Jacobian, and remember that we have stability if $|\lambda| < 1$.)

- e) Use the substitution $R_t = N - S_t - I_t$ to reduce the three dimensional SIR model to a two-dimensional SI model.
- f) Find the equilibria of this SI model.
- g) Show that the endemic equilibrium exists if and only if $R_0 > 1$.
- h) Show that the disease-free equilibrium is locally asymptotically stable if $R_0 < 1$.
- i) If $\beta = 0.5$ and $b = \gamma = 0.05$, find the minimum proportion of people \bar{p} who should be vaccinated in order to eradicate the disease.
- j) Write a program to model the spread of measles with the above parameters for both $p < \bar{p}$ and $p > \bar{p}$.



14

A disease with an asymptomatic class

In this chapter, we're going to look at a disease where the infection may split into a symptomatic class and an asymptomatic class, like influenza. Asymptomatic individuals may not know they're infected (which means they're unlikely to seek treatment or stay home), but they may also transmit the infection at a lower rate, due to a lower viral load. They may also recover faster — or not, depending.

Unlike Chapter 7, immunity may not be permanent. Recovered individuals can lose their immunity due to mutation of the virus and return to the susceptible class. This makes things a lot more complicated.

By the end of this chapter, you should at least have some appreciation of the complexities involved in realistic disease modelling. This chapter consolidates a lot of what we've seen previously.

14.1 Asymptomatic infection

Many diseases do not confer permanent immunity, but rather provide only temporary immunity. Influenza is one such disease, because the virus mutates each year. Previously, we used delay differential equations to describe this effect, but they're not essential. So here we'll move from susceptible to infected to recovered and then have the ability to return to the susceptible class.

A tweak we can add is to suppose that some proportion of infected individuals are asymptomatic. That is, they are still infected (and infectious), but they may have different transmission and recovery rates. We're not explicitly modelling symptoms per se, so the effect is to split the infectious class in two. One thing to be careful of: the word “proportion”. Upon infection, a fraction of individuals q will move to the regular infected class, while the remainder $(1 - q)$ will move to the asymptomatic class. See Figure 14.1.

Note that susceptibles can be infected by either symptomatic or asymptomatic individuals. So there are *four* infection terms floating around, characterised by $q\beta$, $(1 - q)\beta$, $q\beta_A$ and $(1 - q)\beta_A$. So the differential equations are

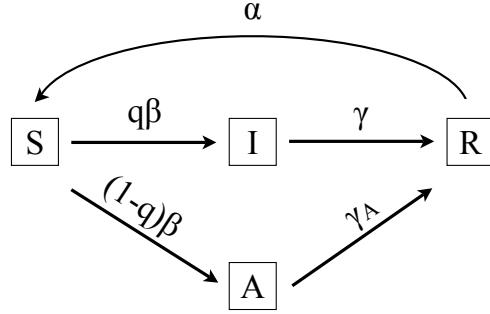


Fig. 14.1. A model for influenza with temporary immunity and an asymptomatic class.

given by

$$\begin{aligned}
 S' &= \lambda - \beta SI - \beta_A SA + \alpha R - \mu S \\
 I' &= q\beta SI + q\beta_A SA - \mu I - \gamma I \\
 A' &= (1-q)\beta SI + (1-q)\beta_A SA - \mu A - \gamma_A A \\
 R' &= \gamma I + \gamma_A A - \alpha R - \mu R.
 \end{aligned}$$

The disease-free equilibrium is given by

$$(\bar{S}, \bar{I}, \bar{A}, \bar{R}) = \left(\frac{\lambda}{\mu}, 0, 0, 0 \right).$$

(You should probably be able to find this for yourself by now.)

Next, we'd like to find the reproduction number. We're going to do this in two ways, to illustrate some of the issues that may arise. So we'll outline the key steps for each.

14.1.1 Using the characteristic polynomial

Here are the steps we need to take:

1. Find the Jacobian matrix.
2. Evaluate the Jacobian at the disease-free equilibrium
3. Determine the characteristic equation.
4. Use information about the coefficients to develop an R_0 .
5. Identify any problems that arise.

Differentiating, the Jacobian is

$$J = \begin{bmatrix} -\beta I - \beta_A A - \mu & -\beta S & -\beta_A S & \alpha \\ q\beta I + q\beta_A A & q\beta S - \mu - \gamma & q\beta_A S & 0 \\ (1-q)\beta I + (1-q)\beta_A A & (1-q)\beta S & (1-q)\beta_A S - \mu - \gamma_A & 0 \\ 0 & \gamma & \gamma_A & -\alpha - \mu \end{bmatrix}$$

$$J|_{\text{DFE}} = \begin{bmatrix} -\mu & -\beta\bar{S} & -\beta_A\bar{S} & \alpha \\ 0 & q\beta\bar{S} - \mu - \gamma & q\beta_A\bar{S} & 0 \\ 0 & (1-q)\beta\bar{S} & (1-q)\beta_A\bar{S} - \mu - \gamma_A & 0 \\ 0 & \gamma & \gamma_A & -\alpha - \mu \end{bmatrix}.$$

The characteristic equation is then

$$\det(J|_{\text{DFE}} - \Lambda I) = (-\mu - \Lambda)(-\alpha - \mu - \Lambda) \det M,$$

where

$$M = \begin{bmatrix} q\beta\bar{S} - \mu - \gamma - \Lambda & q\beta_A\bar{S} \\ (1-q)\beta\bar{S} & (1-q)\beta_A\bar{S} - \mu - \gamma_A - \Lambda \end{bmatrix}.$$

The first two eigenvalues are always negative, while the characteristic equation for M is

$$\begin{aligned} 0 &= (q\beta\bar{S} - \mu - \gamma - \Lambda)((1-q)\beta_A\bar{S} - \mu - \gamma_A - \Lambda) - q(1-q)\beta\beta_A\bar{S}^2 \\ &= \Lambda^2 + \Lambda [\mu + \gamma - q\beta\bar{S} + \mu + \gamma_A - (1-q)\beta_A\bar{S}] \\ &\quad + (\mu + \gamma)(\mu + \gamma_A) - q\beta\bar{S}(\mu + \gamma_A) - (\mu + \gamma)(1-q)\beta_A\bar{S}. \end{aligned}$$

Rearranging the constant term of the characteristic equation, we have

$$\begin{aligned} (\mu + \gamma)(\mu + \gamma_A) &= q\beta\bar{S}(\mu + \gamma_A) + (\mu + \gamma)(1-q)\beta_A\bar{S} \\ R_0 &= \frac{q\beta\bar{S}}{\mu + \gamma} + \frac{(1-q)\beta_A\bar{S}}{\mu + \gamma_A}. \end{aligned}$$

What problems arise here? Answer: we didn't consider the coefficient of the linear term in the characteristic equation. Consider

$$b = \mu + \gamma - q\beta\bar{S} + \mu + \gamma_A - (1-q)\beta_A\bar{S}.$$

If $b > 0$, then the largest eigenvalue will be negative if $R_0 < 1$ and will be positive if $R_0 > 1$. What happens if $b < 0$? It's certainly possible that it might be, as there are negative terms in it. If $b < 0$, then the largest eigenvalue will be positive, regardless of the constant term. So even if $R_0 < 1$, the disease-free equilibrium will be unstable.

Evaluating b at the threshold

There's one more thing we can do, however, which is evaluate b at the threshold $R_0 = 1$. At the threshold, the constant term is zero, so if the other eigenvalue is negative, then small changes in the constant term won't change the negativity of the other eigenvalue. That is, if $b < 0$ when $R_0 = 1$, we will have local stability.

Using the condition $R_0 = 1$ gives us a further constraint that can allow some of the negative terms to be cancelled. Substituting

$$\mu + \gamma = q\beta\bar{S} + \frac{\mu + \gamma}{\mu + \gamma_A}(1 - q)\beta_A\bar{S}$$

gives

$$\begin{aligned} b|_{R_0=1} &= q\beta\bar{S} + \frac{\mu + \gamma}{\mu + \gamma_A}(1 - q)\beta_A\bar{S} - q\beta\bar{S} + \mu + \gamma_A - (1 - q)\beta_A\bar{S} \\ &= \frac{\mu + \gamma - (\mu + \gamma_A)}{\mu + \gamma_A}(1 - q)\beta_A\bar{S} + \mu + \gamma_A \\ &= \frac{\gamma - \gamma_A}{\mu + \gamma_A}(1 - q)\beta_A\bar{S} + \mu + \gamma_A \end{aligned}$$

This will be positive if $\gamma \geq \gamma_A$; i.e., if recovery from symptomatic infection is faster than asymptomatic infection. So if asymptomatic infection lasts longer, then R_0 is a threshold. However, this is unlikely to be true for most diseases, because the viral load for asymptomatic infection is lower, so recovery should be quicker.

This condition isn't sharp though. So we could have $\gamma < \gamma_A$ and still have $b|_{R_0=1} > 0$. We'll explore this case in the lab.

Splitting R_0 into its constituent components

Suppose $q = 1$, so that nobody is asymptomatic. Then, from the model, the asymptomatic equation becomes

$$A' = -\mu A - \gamma_A A,$$

so $A \rightarrow 0$ as $t \rightarrow \infty$. So we might as well set $A = 0$. (More formally, if $A(0) = 0$, then $A \equiv 0$.) Then the infected equation becomes

$$\begin{aligned} I' &= \beta S I - \mu I - \gamma I \\ &= (\beta S - \mu - \gamma)I, \end{aligned}$$

so we can set

$$R_0^I = \frac{\beta\bar{S}}{\mu + \gamma}.$$

This is the reproduction number for a single infected individual in a wholly susceptible population. In particular, if $R_0^I < 1$, then $I' < 0$, whereas if $R_0^I > 1$, then $I' > 0$.

Next suppose that $q = 0$, so that everybody is asymptomatic. The infected equation becomes

$$I' = -\mu I - \gamma I,$$

so we can set $I = 0$. The asymptomatic equation becomes

$$\begin{aligned} A' &= \beta_A S A - \mu A - \gamma_A A \\ &= (\beta_A S - \mu - \gamma_A) A. \end{aligned}$$

We can then define

$$R_0^A = \frac{\beta_A \bar{S}}{\mu + \gamma_A}.$$

This is the reproduction number for a single asymptomatic individual in a population of susceptibles.

We can thus write

$$\begin{aligned} R_0 &= \frac{q\beta \bar{S}}{\mu + \gamma} + \frac{(1-q)\beta_A \bar{S}}{\mu + \gamma_A} \\ &= qR_0^I + (1-q)R_0^A. \end{aligned}$$

That is, R_0 can be split into a linear combination of the individual reproduction numbers for each of the two substrains of the virus.

14.1.2 Using the next-generation method

Here are the steps we need to take using this method:

1. Identify the infected classes.
2. Determine the matrix of new infections and the matrix of transfer terms.
3. Find the largest eigenvalue to calculate R_0 .
4. Identify any problems that arise.

First, let's note that the only infected classes are the I and A classes. So this will reduce our problem to a two-dimensional one. The vector of new infections is

$$\mathcal{F} = \begin{bmatrix} q\beta SI + q\beta_A SA \\ (1-q)\beta SI + (1-q)\beta_A SA \end{bmatrix}.$$

We are expressing the middle two equations in the form $\mathcal{F} - \mathcal{V}$, so the vector of transfer terms is

$$\mathcal{V} = \begin{bmatrix} \mu I + \gamma I \\ \mu A + \gamma_A A \end{bmatrix}.$$

Differentiating, we have the matrices

$$F = \begin{bmatrix} q\beta S & q\beta_A S \\ (1-q)\beta S & (1-q)\beta_A S \end{bmatrix} \quad V = \begin{bmatrix} \mu + \gamma & 0 \\ 0 & \mu + \gamma_A \end{bmatrix}$$

The reproduction number is the largest eigenvalue of FV^{-1} , so (using Appendix C) we have

$$\begin{aligned} FV^{-1} \Big|_{\text{DFE}} &= \begin{bmatrix} q\beta\bar{S} & q\beta_A\bar{S} \\ (1-q)\beta\bar{S} & (1-q)\beta_A\bar{S} \end{bmatrix} \begin{bmatrix} \frac{1}{\mu+\gamma} & 0 \\ 0 & \frac{1}{\mu+\gamma_A} \end{bmatrix} \\ &= \begin{bmatrix} \frac{q\beta\bar{S}}{\mu+\gamma} & \frac{q\beta_A\bar{S}}{\mu+\gamma_A} \\ \frac{(1-q)\beta\bar{S}}{\mu+\gamma} & \frac{(1-q)\beta_A\bar{S}}{\mu+\gamma_A} \end{bmatrix}. \end{aligned}$$

Next we need to find the largest eigenvalue, so (using Appendix C again) we have

$$\begin{aligned} \det(FV^{-1} - \Lambda I) &= \det \begin{bmatrix} \frac{q\beta\bar{S}}{\mu+\gamma} - \Lambda & \frac{q\beta_A\bar{S}}{\mu+\gamma_A} \\ \frac{(1-q)\beta\bar{S}}{\mu+\gamma} & \frac{(1-q)\beta_A\bar{S}}{\mu+\gamma_A} - \Lambda \end{bmatrix} \\ &= \Lambda^2 - \Lambda \left(\frac{q\beta\bar{S}}{\mu+\gamma} + \frac{(1-q)\beta_A\bar{S}}{\mu+\gamma_A} \right) + \frac{q\beta\bar{S}}{\mu+\gamma} \frac{(1-q)\beta_A\bar{S}}{\mu+\gamma_A} \\ &\quad - \frac{q\beta_A\bar{S}}{\mu+\gamma_A} \frac{(1-q)\beta\bar{S}}{\mu+\gamma} \\ &= \Lambda \left(\Lambda - \frac{q\beta\bar{S}}{\mu+\gamma} - \frac{(1-q)\beta_A\bar{S}}{\mu+\gamma_A} \right). \end{aligned}$$

The smaller eigenvalue is 0, while the larger one is

$$\begin{aligned} R_0 &= \frac{q\beta\bar{S}}{\mu+\gamma} + \frac{(1-q)\beta_A\bar{S}}{\mu+\gamma_A} \\ &= qR_0^I + (1-q)R_0^A. \end{aligned}$$

Note that R_0 is the largest eigenvalue from the next-generation method, but it is NOT the largest eigenvalue from the Jacobian or characteristic equation methods. Instead, it's a *rearrangement* of the largest eigenvalue from the Jacobian or characteristic equation methods. Yes, it's confusing.

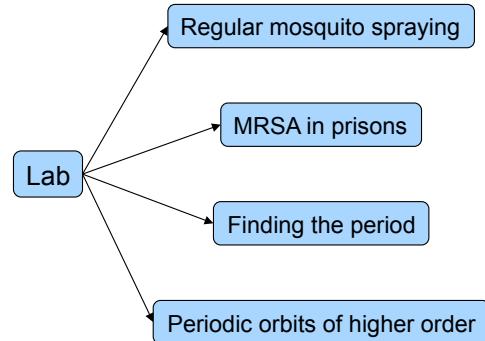
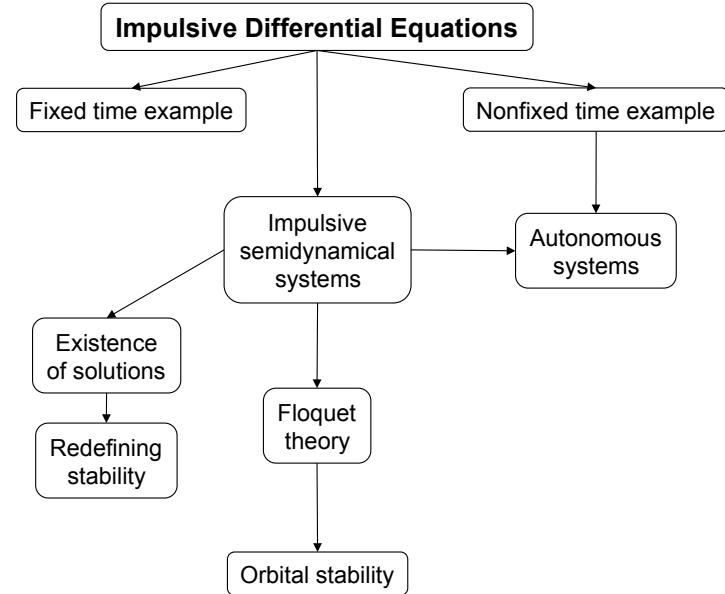
What problems arise now? Well, we got the same R_0 as with the previous method, but without the constraint of $b > 0$. So does that constraint matter? We'd never know it existed if we only used the next-generation method. We'll explore that in the lab.

14.2 Lab work

14.2.1 Exercises

1. Choose parameters that make $b > 0$ and $R_0 < 1$. Run the time series for the model in this case.

2. Now change one of your parameters so that $b < 0$ and $R_0 < 1$. Run the time series for the model now.
3. Could $\gamma < \gamma_A$ and yet $b|_{R_0=1} > 0$?
4. If $b|_{R_0=1} = 0$, find the critical value of γ_A . Is this reasonable?
5. Let $f = R_0^A/R_0^I$. What conditions are required for the outbreak to be worse as a result of both strains compared just one or the other? (Hint: substitute $R_0^A = fR_0^I$ and compare to both R_0^A and R_0^I .)



15

Impulsive Differential Equations

15.1 Introduction

Many systems undergo drastic changes in a short space of time. If you take a pill, the amount of drug in your body quickly rises to its maximal level. It takes roughly 20 minutes for this to occur (although it depends on the drug of course). This is called the time-to-peak. After this, your body starts metabolising the drug, so it decreases, approximately exponentially. If the time until you take your next pill is significantly longer than 20 minutes (say 24 hours), then we can approximate the 20-minute interval by an *instantaneous* change.

Why would we do this? Ignoring the time-to-peak part means we don't have to link up a series of differential equations (one for the time-to-peak and one for the metabolising). Since the metabolising part is usually what we're most interested in, it allows us to concentrate on this. With an exponential decay between doses, this makes the analysis quite straightforward, as we'll see shortly.

Of course, we still have two scales to deal with. So we'll approximate the time-to-peak with an instant jump. Mathematically, this consists of a difference equation, not a differential equation. That is, solutions restart at new initial conditions, except that these initial conditions are linked to the final conditions from the previous cycle.

15.1.1 Fixed-time impulse example

So our drug example would look like this:

$$\begin{aligned} x' &= -\alpha x & t &\neq t_k \\ \Delta x &= c & t &= t_k. \end{aligned}$$

That is, so long as we're not at one of the pill-popping times, solutions will decrease exponentially (the metabolising part), as described by the differential

equation. The second line is a *difference equation*, where $\Delta x \equiv x(t_k^+) - x(t_k^-)$. This is the instantaneous approximation of the time-to-peak. In our simple example, the difference between just before the pill (at time t_k^-) and just after (at time t_k^+) is a constant c , the strength of the drug in a single pill.

One thing to note here: the t_k 's don't have to be constant. Maybe we take our pill regularly, but maybe we don't. The beauty of impulsive differential equations is that they can apply to either situation. For simplicity, let's first assume that our pill is taken regularly. That is, the time between doses is $\tau = t_{k+1} - t_k$. (We'll examine the other case in the next example.)

How do we deal with something like this? First, let's solve the differential equation. Of course, most differential equations can't be solved, but this one is simple enough that it can be. We have to be a little careful here though: we're solving a differential equation, but what's the "initial" condition? Hint: it's not the initial condition from the system. That's because each impulsive cycle has its own initial condition, $x(t_k^+)$. The "initial" time for each cycle isn't zero either; it's t_k .

Thus, within each cycle, we have

$$x(t) = x(t_k^+) e^{-\alpha(t-t_k)}.$$

So far so good. We've technically solved the entire differential equation part. But each solution comes with an "initial" condition $x(t_k^+)$ that depends on the "final" condition from the previous cycle $x(t_{k-1}^-)$. The good news is that we have all the tools we need to put everything together: the difference equation from the model tells us how to get across the impulse, while the solution of the differential equation tells us how to get to the end of the next cycle.

Let's work it through from the true initial condition. (Remember that every differential equation has an initial condition, whether or not it's explicitly written down. The same is true of every impulsive differential equation.) Technical note: because there might have been an impulse at time $t = 0$, we specify the initial condition at $x(0^+)$. Let's call our initial condition x_0 .

We then have

$$\begin{aligned}
 x(0^+) &= x_0 && \text{the true initial condition} \\
 x(t_1^-) &= x_0 e^{-\alpha\tau} && \text{the "final" condition from the 1st cycle} \\
 x(t_1^+) &= x_0 e^{-\alpha\tau} + c && \text{adding the impulse} \\
 x(t_2^-) &= x(t_1^+) e^{-\alpha\tau} && \text{using the solution} \\
 &= (x_0 e^{-\alpha\tau} + c) e^{-\alpha\tau} && \text{using the previous "initial" condition} \\
 &= x_0 e^{-2\alpha\tau} + c e^{-\alpha\tau} && \text{the "final" condition for the 2nd cycle} \\
 x(t_2^+) &= x_0 e^{-2\alpha\tau} + c e^{-\alpha\tau} + c && \text{the "initial" condition for the 3rd cycle} \\
 x(t_3^+) &= x_0 e^{-3\alpha\tau} + c e^{-2\alpha\tau} + c e^{-\alpha\tau} + c && \text{and so on} \\
 &\vdots
 \end{aligned}$$

Notice that we jumped straight to the “initial” condition for the fourth cycle by recognising the pattern after a while. What we really have here is a recurrence relation; these are best solved when you can recognise patterns.

Hence the general term for the “initial” condition for the $(n+1)$ st cycle is

$$\begin{aligned} x(t_n^+) &= x_0 e^{-n\alpha\tau} + c \left[1 + e^{-\alpha\tau} + e^{-2\alpha\tau} + \cdots + e^{-(n-1)\alpha\tau} \right] \\ &= x_0 e^{-n\alpha\tau} + \frac{c(1 - e^{-n\alpha\tau})}{1 - e^{-\alpha\tau}}. \end{aligned}$$

We used the fact that a geometric series has a handy formula to collapse the sum into a nice formula. However, this is where we need the fact that the period τ is constant. If it weren’t, we couldn’t do this last line.

Okay, great. We’ve solved the system! Have we really? Yes... because we have the n th “initial” condition (from the recurrence relation), and we also have the solution during any cycle, given an “initial” condition. The solution is plotted in Figure 15.1.

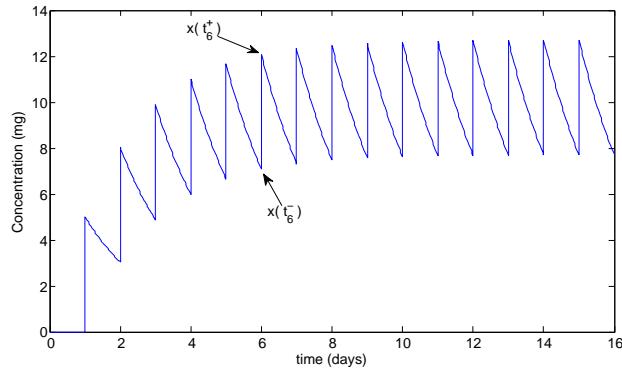


Fig. 15.1. Drug concentrations with impulses at fixed times. Parameters used were $\alpha = 0.5$ per day, $c = 5$ milligrams, and the impulse was applied daily.

So now what? Well, let’s do what we always do when we have a solution: take the limit. In particular, what happens if we have infinite cycles? We have

$$\lim_{n \rightarrow \infty} x(t_n^+) = \frac{c}{1 - e^{-\alpha\tau}}.$$

Note that this is independent of x_0 , so the solution holds for any initial conditions.

What we did here was build a recurrence relation from the endpoints of each cycle and showed they converged. The endpoints are local maxima for the system (since the metabolising always decreases the amount of drug in

your body). So each solution will decrease within each cycle, then a new pill increases the amount of drug.

What's happening at the "infinite" endpoint? Let's be a bit more down-to-earth and set this as our initial condition. We thus have

$$\begin{aligned}
 x(0^+) &= \frac{c}{1 - e^{-\alpha\tau}} && \text{the initial condition occurs after an impulse} \\
 x(\tau^-) &= \frac{c}{1 - e^{-\alpha\tau}} e^{-\alpha\tau} && \text{using the solution} \\
 x(\tau^+) &= \frac{c}{1 - e^{-\alpha\tau}} e^{-\alpha\tau} + c && \text{adding the constant amount of drug} \\
 &= \frac{ce^{-\alpha\tau} + c(1 - e^{-\alpha\tau})}{1 - e^{-\alpha\tau}} \\
 &= \frac{c}{1 - e^{-\alpha\tau}} = x(0^+).
 \end{aligned}$$

That is, if we start on this endpoint, then we return to it after one cycle. That means we have a *periodic orbit with order 1*. It's periodic because we return to it and of order 1 since there's only one impulse per cycle.

Formally, the periodic orbit is given by

$$\begin{aligned}
 x(t) &= \frac{c}{1 - e^{-\alpha\tau}} e^{-\alpha t} && 0 < t \leq \tau \\
 x(0^+) &= \frac{c}{1 - e^{-\alpha\tau}}.
 \end{aligned}$$

By convention, we assign the equality in the range to the "final" point, which means we still have to specify the true initial condition separately... just like with every differential equation. Neat how these things work out.

15.1.2 Nonfixed-time example

There's another way to formulate these things, however. What if, instead of taking your drug every day, you only took your drug when you felt sick? That is, the time at which the impulse occurs depends on the state of the system, not the time. Such systems are called *autonomous*, and we'll explore them in more detail in a bit.

For example, suppose we took our drug only when the level of drugs in our body reached some lower threshold. Then our system could be expressed as

$$\begin{aligned}
 x' &= -\alpha x && x \neq \bar{x} \\
 \Delta x &= c && x = \bar{x}.
 \end{aligned}$$

As before, the solution of the differential equation is

$$x(t) = x(t_k^+) e^{-\alpha(t-t_k)}.$$

In this case, however, we know what $x(t_k^+)$ is, but we don't know what t_k is.

There's another wrinkle. If $x(0^+) < \bar{x}$, then the impulse condition will never be reached and solutions will tend towards zero. So we need to ensure that $x_0 > \bar{x}$. (What if $x_0 = \bar{x}$? In that case, the solution has "just" undergone an impulse, so the solution then is $x(t) = \bar{x}e^{-\alpha t}$, which will approach zero.)

We thus have

$$\begin{aligned} x(0^+) &= x_0 \\ x(t_1) &= x_0 e^{-\alpha t_1} = \bar{x}. \end{aligned}$$

Solving for t_1 , we have

$$t_1 = \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}}.$$

(Note that this is positive, since $x_0 > \bar{x}$.)

Next, we have

$$\begin{aligned} x(t_1^+) &= \bar{x} + c \\ x(t_2^-) &= (\bar{x} + c)e^{-\alpha(t_2 - t_1)} = (\bar{x} + c)e^{-\alpha\tau} = \bar{x} \end{aligned}$$

We're going to call $\tau = t_2 - t_1$ for reasons that will become apparent in a moment. Solving, we have

$$\tau = \frac{1}{\alpha} \ln \frac{\bar{x} + c}{\bar{x}} = \frac{1}{\alpha} \ln \left(1 + \frac{c}{\bar{x}}\right).$$

Continuing, we have

$$\begin{aligned} x(t_2^+) &= \bar{x} + c \\ x(t_3^-) &= (\bar{x} + c)e^{-\alpha(t_3 - t_2)} = \bar{x}. \end{aligned}$$

Solving, we have

$$t_3 - t_2 = \frac{1}{\alpha} \ln \left(1 + \frac{c}{\bar{x}}\right) = \tau.$$

What's happening here? The first cycle has its own period. But thereafter, every cycle has the same period. Which makes sense, since they all have the same "initial" condition and differential equation. That is, the period has converged (after one impulse, because this was a simple example) and we have a periodic orbit. See Figure 15.2.

How do we know this periodic orbit exists? We need $x_0 > \bar{x}$, as mentioned above. Furthermore, since this is a biological system, we need the orbit to be positive. This is equivalent to requiring that the minimum value of the orbit is positive. What's the minimum value? If $x_0 > \bar{x}$, then it's \bar{x} . So our condition for existence is $\bar{x} > 0$.

(Note that if $\bar{x} = 0$, then — since that happens to be the equilibrium value of the ODE — solutions will not reach the impulsive condition in finite time. In this case, there will be no impulsive periodic orbit. So the inequality is strict.)

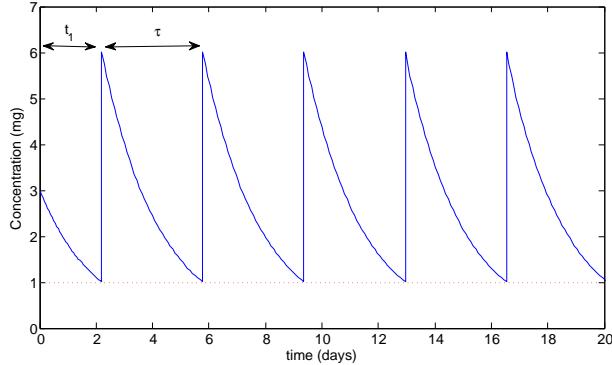


Fig. 15.2. Drug concentrations with impulses at nonfixed times. Parameters used were $\alpha = 0.5$ per day, $c = 5$ milligrams, and the impulse was applied when solutions reached $\bar{x} = 1$.

15.2 Motivation

So that's an overview. If you're not mathematically inclined, you should probably skip the rest of this chapter. But for those who care about the details, we have to worry about all sorts of things. Most mathematical theorems assume continuity at a baseline, but we don't even have that. This means that many theorems have to be reworked, while others don't have an analogue at all.

An equilibrium in the differential equation might not be an equilibrium of the system — as indeed it isn't in the fixed-time example above. The “equilibrium” $x = 0$ is fine until you reach t_1 , after which it's no longer zero, because a constant amount c has been added. This plays havoc with basic concepts like stability (although there's a fix, as we'll see below).

Mostly, unlike ODEs, we're not dealing with equilibria as our phenomena of interest, but rather periodic orbits. Or, more specifically, impulsive periodic orbits. These are cycles where the ODEs send solutions in one direction and the impulsive effect brings it back, just as we saw in the fixed-time example.

The rest of this chapter is very technical, because we need to define things from the very beginning. And the definitions can be quite finicky, because we need to account for all the complexities that arise when disturbing the solutions of differential equations (which is what we're doing). So follow along if you can, but don't sweat it if this is too much in the woods for you. Of course, if precise definitions of mathematical concepts is your kind of thing, then you're in the right place. Much of what follows is developed from the works of Bainov and Simeonov [3, 4].

15.3 Impulsive semidynamical systems

We start by expanding the definition of forward movement under differential equations. First let's state that definition.

Definition 15.1. *A triple (X, π, \mathbb{R}_+) is a semidynamical system if X is a metric space, \mathbb{R}_+ the set of all nonnegative reals and $\pi : X \times \mathbb{R}_+ \rightarrow X$ is a continuous function such that*

- i) $\pi(x, 0) = x$ for all $x \in X$, and
- ii) $\pi(\pi(x, t), s) = \pi(x, t + s)$ for all $x \in X$ and $t, s \in \mathbb{R}_+$.

What this means is we have a starting point (from the initial condition) and also a way to move forward (from the ODE).

For all $x \in X$, define $\pi_x : \mathbb{R}_+ \rightarrow X$ by $\pi_x(t) = \pi(x, t)$. π_x is continuous for all x . We call π_x the *trajectory* of x . The set

$$C^+(x) = \{\pi(x, t) : t \in \mathbb{R}_+\}$$

is called the *positive orbit* of x . Note that $x \in C^+(x)$. We also have

$$C^+(x, r) = \{\pi(x, t) : 0 \leq t \leq r\}.$$

For any $M \subseteq X$, we define the following sets: for $t \in \mathbb{R}_+$,

$$G(x, t) = \{y \in X : \pi(y, t) = x\},$$

is the attainable set of x at $t \in \mathbb{R}_+$,

$$G(x) = \bigcup_{t \in \mathbb{R}_+} G(x, t),$$

$$M^-(x) = G(x) \cap M \setminus \{x\}$$

and

$$M^+(x) = C^+(x) \cap M \setminus \{x\}.$$

We then set $M(x) = M^+(x) \cup M^-(x)$. Note that $x \notin M(x)$.

Definition 15.2. *An impulsive semidynamical system $(X, \pi; M, A)$ consists of a semidynamical system (X, π) , a nonempty closed subset M of X and a continuous function $A : M \rightarrow X$ such that*

- i) *No point $x \in X$ is a limit point of $M(x)$, and*
- ii) *$\{t \in \mathbb{R}_+ : G(x, t) \cap M \neq \emptyset\}$ is a closed subset of \mathbb{R}_+ .*

What this means is that solutions can travel forward either via the mechanism of the ODE or they can jump when they hit the set M , so long as the jump points don't accumulate. We don't want an infinite number of impulses in finite time, for instance.

We denote the image of M under the operator A by $N = A(M)$ and, for all $x \in M$, $A(x) = x^+$. That is, trajectories hit M and jump to N .

Lemma 15.3. *Let $(X, \pi; M, A)$ be an impulsive semidynamical system. Then, for any $x \in X$, there exists $r, s \in \mathbb{R}_+ \cup \{\infty\}$ such that $0 < r, s \leq \infty$ and, for $0 < t < s$ and $0 < t < r$,*

- a) $\pi(x, t) \notin M$ and if $M^+(x) \neq \emptyset$, then $\pi(x, s) \in M$
- b) $G(x, t) \cap M = \emptyset$ and if $M^-(x) \neq \emptyset$, then $G(x, r) \cap M \neq \emptyset$.

This says that solutions travel forward as usual if they're not at the set M . This is true for both forward time and backward time.

We call s the *time without impulse* of x . We define $\Phi : X \rightarrow \mathbb{R}_+ \setminus \{0\}$ such that $\Phi(x)$ is the time without impulse of x . If $\{x_n\}$ is the set of impulse points, then $\{s_n\}$ are the corresponding times without impulse. We can think of a given s_n as the time taken from the trajectory starting at x_n until x_{n+1} (the next impulse point). Naturally, if there is no further impulse point, then $s_{n+1} = \infty$.

Definition 15.4. *Let $(X, \pi; M, A)$ be an impulsive semidynamical system and $x \in X$. The (impulsive) trajectory of x is a function $\tilde{\pi}_x$ defined on a subset $[0, s)$, $s \in (0, \infty]$ as follows:*

Let $x = x_0$. If $M^+(x_0) = \emptyset$, then $\tilde{\pi}_x(t) = \pi(x, t)$ for all $t \in \mathbb{R}_+$. If $M^+(x_0) \neq \emptyset$, then, by Lemma 15.3, there exists $s_0 \in \mathbb{R}_+ \setminus \{0\}$ such that $\pi(x_0, s_0) = x_1 \in M$ and $\pi(x_0, t) \notin M$ for all $0 < t < s_0$. We define $\tilde{\pi}_x$ on $[0, s_0]$ by

$$\tilde{\pi}_x(t) = \pi(x_0, t) \quad 0 \leq t \leq s_0.$$

We then continue this process, starting at x_1^+ (which is not equal to x_1 in general). That is, if $M^+(x_1^+) = \emptyset$, then we define $\tilde{\pi}_x(t) = \pi(x_1^+, t - s_0)$ for all $t > s_0$ and $s = \infty$. If $M^+(x_1^+) \neq \emptyset$, then, by Lemma 15.3, there exists $s_1 \in \mathbb{R}_+ \setminus \{0\}$ such that $\pi(x_1^+, s_1) = x_2 \in M$ and $\pi(x_1^+, t) \notin M$ for all $0 < t < s_1$. We define $\tilde{\pi}_x$ on $(s_0, s_0 + s_1]$ by

$$\tilde{\pi}_x(t) = \pi(x_1^+, t - s_0) \quad s_0 < t \leq s_0 + s_1.$$

If $M^+(x_2^+) \neq \emptyset$, then, by Lemma 15.3, there exists $s_2 \in \mathbb{R}_+ \setminus \{0\}$ such that $\pi(x_2^+, s_2) = x_3 \in M$ and $\pi(x_2^+, t) \notin M$ for all $0 < t < s_2$. We define $\tilde{\pi}_x$ on $(s_0 + s_1, s_0 + s_1 + s_2]$ by

$$\tilde{\pi}_x(t) = \pi(x_2^+, t - s_0 - s_1) \quad s_0 + s_1 < t \leq s_0 + s_1 + s_2.$$

If $M^+(x_n^+) = \emptyset$ for some n , then the process halts. On the other hand, if $M^+(x_n^+) \neq \emptyset$ for all $n = 1, 2, \dots$ then the process continues indefinitely, with

$$\tilde{\pi}_x(t) = \pi(x_n^+, t - \sum_{i=0}^{n-1} s_i), \quad \sum_{i=0}^{n-1} s_i < t \leq \sum_{i=0}^n s_i$$

for each $n \geq 1$.

Thus the process gives rise to either a finite or infinite sequence $\{x_n\}$ of points of X such that, with each x_n , there is associated a positive real number s_n (or ∞) and, for $s_n < \infty$, an impulse x_{n+1} , where $\pi(x_n^+, s_n) = x_{n+1}$.

The interval of definition of $\tilde{\pi}_x$ is $[0, s] = [0, \sum_{i=0}^{\infty} s_i]$.

What this does is define our trajectories in a piecemeal fashion. We move forward, and if we hit M , we jump, then move forward again. Of course, if we ever reach a state where we can never hit M again, the remaining trajectory is the usual forward trajectory. Otherwise, we'll hit M infinitely often.

Definition 15.5. A trajectory $\tilde{\pi}_x$ is periodic of period r and order k if there exists $m \in \mathbb{Z}^+$ and $k \in \mathbb{Z}^+$ such that k is the smallest integer satisfying $x_m^+ = x_{m+k}^+$ and

$$r = \sum_{i=m}^{m+k-1} s_i.$$

This defines what we mean by a period: something that comes back around again (perhaps with the help of an impulse) but wasn't where it started for all time in between. Essentially, we need this to distinguish periodic orbits from equilibria (which "come back" to where they started only in the sense that they never left, so they're not really periodic).

A periodic trajectory with no impulse points can be considered to be an impulsive trajectory with one moment of impulse, such that the trajectory is continuous at the impulse point. Thus a periodic trajectory with no impulse points is a first-order periodic orbit, and the period is the time taken to travel from the impulse point back to itself; hence the period in this case corresponds to the definition of period in the non-impulsive case.

Note that the trajectory $\tilde{\pi}_x$ is continuous if either $M^+(x) = \emptyset$ or, for each n , $x_n = x_n^+$. Otherwise, the trajectory has discontinuities at a finite or infinite number of impulse points x_n . However, at any such point, $\tilde{\pi}_x$ is continuous from the left.

Trajectories of interest for impulsive semidynamical systems are those with an infinite number of discontinuities and an interval of definition of \mathbb{R}_+ . We call these *infinite trajectories*.

Example. Consider the autonomous system

$$x' = x \quad y' = \alpha y, \quad \alpha > 0,$$

the sets $M = \{(x, y) \in \mathbb{R}_+^2 : y = \frac{1}{x+1}\}$, $N = \{(x, y) \in \mathbb{R}_+^2 : x + y = 1\}$, and an operator $A : M \rightarrow N$ that associates with each point P on M the point P^+ on N which is on the ray OP . A is a continuous, bijective mapping.

We shall consider only those trajectories with initial points in the first quadrant. Note that this quadrant is invariant. We assume initial points are not on M , by convention.

Trajectories with initial points in the region $y > \frac{1}{1+x}$ do not undergo any impulsive effect. Trajectories with initial points on the x -axis also do not undergo impulsive effect, since M does not intersect the x -axis. Trajectories with initial points on the y -axis undergo impulsive effect once, at $(0,1)$, but motion is continuous, since this is a fixed point of the operator A . Both axes are invariant.

For $0 < \alpha < 1$, trajectories with initial points in the region $y < \frac{1}{1+x}$ undergo impulsive effect an infinite number of times. $(x_n^+, y_n^+) \rightarrow (1, 0)$, $s = \infty$ and $\tilde{\pi}_{(1,0)} = \pi_{(1,0)}$. See Figure 15.3.

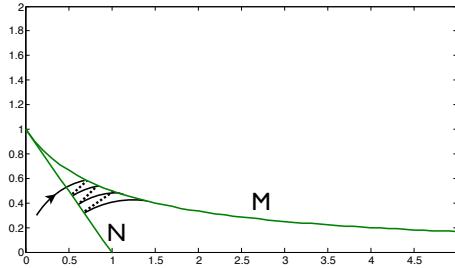


Fig. 15.3. When $\alpha < 1$, the impulsive orbits approach part of the x -axis.

Let $\alpha > 1$. Trajectories with initial points in the region $y < \frac{1}{1+x}$ are subject to impulsive effect an infinite number of times and tend towards the point $(0,1)$, which is a fixed point of the impulsive effect. See Figure 15.4.

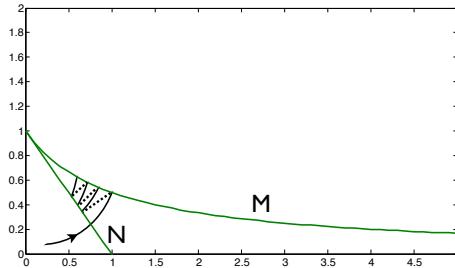


Fig. 15.4. When $\alpha > 1$, the impulsive orbits converge to a fixed point.

When $\alpha = 1$, all trajectories with initial points in the region $y < \frac{1}{1+x}$ eventually become periodic, with order 1. Motion between N and M is performed along rays $y = cx$.

15.3.1 Existence of solutions

Let $\Omega \subset \mathbb{R}^n$ be an open set. Suppose that, for each $k \in \mathbb{Z}$, the functions $\tau_k : \Omega \rightarrow \mathbb{R}$ are continuous in Ω and satisfy

$$\tau_k(x) < \tau_{k+1}(x), \text{ with } \lim_{k \rightarrow \pm\infty} \tau_k(x) = \infty$$

for $x \in \Omega$. Let $f : \mathbb{R} \times \Omega \rightarrow \mathbb{R}^n$, $I_k : \Omega \rightarrow \mathbb{R}^n$, $(t_0, x_0) \in \mathbb{R} \times \Omega$ and $\alpha < \beta$.

Consider the impulsive differential system

$$\begin{aligned} \frac{dx}{dt} &= f(t, x) & t \neq \tau_k(x) \\ \Delta x &= I_k(x) & t = \tau_k(x), \end{aligned} \tag{15.1}$$

with initial condition

$$x(t_0^+) = x_0. \tag{15.2}$$

By definition, $\Delta x \equiv x^+ - x$, so $I_k(x) = x + A_k(x)$.

Definition 15.6. *The function $\varphi : (\alpha, \beta) \rightarrow \mathbb{R}^n$ is a solution of (15.1) if*

1. $(t, \varphi(t)) \in \mathbb{R} \times \Omega$ for $t \in (\alpha, \beta)$,
2. $\varphi(t)$ is differentiable, with

$$\frac{d\varphi}{dt}(t) = f(t, \varphi(t))$$

for $t \in (\alpha, \beta)$, $t \neq \tau_k(\varphi(t))$, and

3. $\varphi(t)$ is continuous from the left in (α, β) , and if $t \in (\alpha, \beta)$, $t = \tau_k(\varphi(t))$ and $t \neq \beta$, then $\varphi(t^+) = \varphi(t) + I_k(\varphi(t))$ and, for each $j \in \mathbb{Z}$ and some $\delta > 0$, $s \neq \tau_j(\varphi(s))$ for $t < s < t + \delta$.

Definition 15.7. *A solution of the initial-value problem (15.1)–(15.2) is a function φ that is defined in an interval of the form (t_0, β) , is a solution of (15.1) and satisfies (15.2).*

These two definitions provide the framework for when a solution exists, showing that it must be able to be continued past the initial value.

15.3.2 Definitions of stability

The usual definition of stability essentially says that if two solutions start close together, then they should stay close together for all time. That's obviously not going to happen with nonfixed impulses: two solutions might stay together for a while, but if one jumps before the other, then there'll be quite far apart indeed, even if they later come back into synch. So we need to adjust our definitions of stability.

Definition 15.8. Let $x_0(t) = x(t; t_0, x_0)$ be a given solution of the initial-value problem (15.1)–(15.2), existing for $t \geq t_0$. Suppose $x_0(t)$ hits the surfaces $S_k : t = t_k(x)$ at the moments t_k such that $t_k < t_{k+1}$ and $t_k \rightarrow \infty$ as $k \rightarrow \infty$. Then the solution $x_0(t)$ of (15.1)–(15.2) is

- stable if for each $\epsilon > 0$, $\eta > 0$ and $t_0 \in \mathbb{R}_+$, there exists $\delta = \delta(t_0, \epsilon, \eta) > 0$ such that $|y_0 - x_0| < \delta$ implies $|y_0(t) - x_0(t)| < \epsilon$ for $t \geq t_0$ and $|t - t_k| > \eta$, where $y_0(t) = x(t; t_0, y_0)$ is any solution of (15.1)–(15.2) existing for $t \geq t_0$;
- uniformly stable if it is stable and δ is independent of t_0 ;
- attractive if for each $\epsilon > 0$, $\eta > 0$ and $t_0 \in \mathbb{R}_+$, there exists $\delta_0 = \delta_0(t_0) > 0$ and a $T = T(t_0, \epsilon, \eta) > 0$ such that $|y_0 - x_0| < \delta_0$ implies $|y_0(t) - x_0(t)| < \epsilon$ for $t \geq t_0 + T$ and $|t - t_k| > \eta$;
- uniformly attractive if it is attractive and δ_0 and T are independent of t_0 ;
- asymptotically stable if it is stable and attractive; and
- uniformly asymptotically stable if it is uniformly stable and uniformly attractive.

What this says is that solutions that start together will stay together for all times except for a small interval around the impulse. Even better, we can make this interval as small as we want by choosing initial conditions very close to each other. If the impulses occur at fixed times (i.e., if $\tau_k(x)$ is independent of x), then the impulse effect occurs at the same time for every solution, so the notions of stability coincide with the standard definition.

The system (15.1)–(15.2) only possesses the trivial solution if $f(t, 0) \equiv 0$ and $I_k(0) = 0$ for all k .

If there are only a finite number of impulse points, then the usual definitions of stability can be applied to the trajectories after the last impulse point.

If there are an infinite number of impulse points, then we do not want the points to accumulate at some finite value, such that $t_k \rightarrow r < \infty$. This accounts for our requiring that $t_k \rightarrow \infty$ as $k \rightarrow \infty$.

15.3.3 Autonomous systems with impulsive effect

Autonomous systems with impulsive effect are written in the form

$$\begin{aligned} \frac{dx}{dt} &= g(x) & x \notin M \\ \Delta x &= I(x) & x \in M. \end{aligned} \tag{15.3}$$

At an instant $t = t_k$ when $x(t)$ encounters the set M , it is instantaneously transferred to the point $x(t_k) + I(x(t_k))$ of the set N . The set M is sometimes given in the form $\phi(x) = 0$.

System (15.3) has the property of autonomy, so that $x(t; t_0, x_0) = x(t - t_0; 0, x_0)$. Note that systems of the form (15.1) do not possess this property, even if $f(t, x) = g(x)$ and $I_k(x) = I(x)$.

Example. The nonfixed-time example from earlier in this chapter is autonomous. To see this, we can write the solution as

$$x(t) = \begin{cases} x_0 e^{-\alpha t} & 0 \leq t \leq t_1 \\ (\bar{x} + c) e^{-\alpha(t-t_k)} & t_k < t \leq t_{k+1}. \end{cases}$$

Changing to formal notation, we can write this as $x(t; t_0, x_0)$, which consists of the following:

$$\begin{aligned} x_0 e^{-\alpha(t-t_0)} & \quad 0 \leq t \leq t_0 + \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} \\ (\bar{x} + c) e^{-\alpha(t-t_0 - \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}})} & \quad t_0 + \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} < t \leq t_0 + \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} + \frac{1}{\alpha} \ln \frac{c+\bar{x}}{c} \\ (\bar{x} + c) e^{-\alpha(t-t_0 - \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} - \frac{1}{\alpha} \ln \frac{c+\bar{x}}{c})} & \quad t_0 + \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} + \frac{1}{\alpha} \ln \frac{c+\bar{x}}{c} < t \leq t_0 + \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} \\ & \quad + \frac{2}{\alpha} \ln \frac{c+\bar{x}}{c} \\ \vdots & \quad \vdots \end{aligned}$$

We can then write $x(t - t_0; 0, x_0)$ as

$$\begin{aligned} x_0 e^{-\alpha(t-t_0)} & \quad 0 \leq t - t_0 \leq \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} \\ (\bar{x} + c) e^{-\alpha(t-t_0 - \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}})} & \quad \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} < t - t_0 \leq \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} + \frac{1}{\alpha} \ln \frac{c+\bar{x}}{c} \\ (\bar{x} + c) e^{-\alpha(t-t_0 - \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} - \frac{1}{\alpha} \ln \frac{c+\bar{x}}{c})} & \quad \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} + \frac{1}{\alpha} \ln \frac{c+\bar{x}}{c} < t - t_0 \leq \frac{1}{\alpha} \ln \frac{x_0}{\bar{x}} \\ & \quad + \frac{2}{\alpha} \ln \frac{c+\bar{x}}{c} \\ \vdots & \quad \vdots \end{aligned}$$

It follows that the system is autonomous.

Example. Consider the system

$$\begin{aligned} \frac{dx}{dt} &= x & t_k \neq k \\ \Delta x &= -\frac{x}{2} + 1 & t_k = k \\ x(0.5) &= 2. \end{aligned}$$

Solutions are given by

$$x(t) = \begin{cases} 2e^{t-0.5} & 0.5 \leq t \leq 1 \\ x_k^+ e^{t-k} & k < t \leq k+1, \quad k \geq 1. \end{cases}$$

Thus $x(1^-) = 2e^{0.5}$, so

$$\begin{aligned} x(2) &= x(1^+)e \\ &= (e^{0.5} + 1)e. \end{aligned}$$

Hence $x(2; 0.5, 2) = e^{1.5} + e$.

Conversely, consider the initial condition $x(0) = 2$. Then $x(1) = 2e$, so

$$\begin{aligned} x(1.5) &= x(1^+)e^{0.5} \\ &= (e+1)e^{0.5}. \end{aligned}$$

Thus $x(1.5; 0, 2) = e^{1.5} + e^{0.5}$, so $x(2; 0.5, 2) \neq x(1.5; 0, 2)$. Hence this system does not have the property of autonomy.

15.4 Floquet Theory in \mathbb{R}^2

15.4.1 Floquet Theory for continuous systems

Floquet theory for ODEs is a way of examining stability of periodic orbits. Since we tend to encounter a lot of periodic orbits with impulses, we need to have a way to determine their stability.

We also provide the proofs of some of the basic theorems. These proofs are straightforward but were not included in the literature, so we have included them here for completeness.

For two-dimensional systems, there is a detailed but relatively straightforward formula for calculation of the second multiplier for a periodic orbit. This allows the theory of Floquet multipliers to be applied to two-dimensional systems, or systems that can be reduced to two-dimensional systems, with ease.

15.4.2 Floquet theory

Consider the linear T -periodic system with fixed moments of impulsive effect

$$\begin{aligned} \frac{dx}{dt} &= P(t)x & t \neq t_k \\ \Delta x &= B_k x & t = t_k, \end{aligned} \tag{15.4}$$

subject to the following assumptions:

- H1 The matrix $P(\cdot) : \mathbb{R} \rightarrow \mathbb{C}^{n \times n}$ is piecewise continuous, and $P(t+T) = P(t)$ for $t \in \mathbb{R}$.
- H2 $t_k < t_{k+1}$ for $k \in \mathbb{Z}$, $B_k \in \mathbb{C}^{n \times n}$, and $\det(I + B_k) \neq 0$, where I is the $n \times n$ identity matrix.
- H3 There exists an integer $q > 0$ such that $B_{k+q} = B_k$, $t_{k+q} = t_k + T$ for $k \in \mathbb{Z}$.

Definition 15.9. Let $x_1(t), \dots, x_n(t)$ be solutions to (15.4) defined on the interval $(0, \infty)$. Let $X(t) = \{x_1(t), \dots, x_n(t)\}$ be a matrix-valued function whose columns are these solutions. Then $x_1(t), \dots, x_n(t)$ are linearly independent if and only if $\det X(0^+) \neq 0$. In this case, we say that $X(t)$ is a fundamental matrix of solutions of (15.4).

Lemma 15.10. Suppose H1–H3 hold and $\lim_{k \rightarrow \infty} t_k = \infty$. Let $X(t)$ be a fundamental matrix of solutions of (15.4) in \mathbb{R}_+ . Then

1. For any constant matrix $\bar{M} \in \mathbb{C}^{n \times n}$, $X(t)\bar{M}$ is also a solution of (15.4).
2. If $Y : \mathbb{R} \rightarrow \mathbb{C}^{n \times n}$ is a solution of (15.4), there exists a unique matrix \bar{M} such that $Y(t) = X(t)\bar{M}$. Furthermore, if $Y(t)$ is also a fundamental matrix of solutions, then $\det \bar{M} \neq 0$.

Proof. 1. $X(t)\bar{M}$ satisfies

$$\begin{aligned}\frac{d}{dt} (X(t)\bar{M}) &= \frac{dX(t)}{dt} \bar{M} \\ &= P(t)X(t)\bar{M}\end{aligned}$$

for $t \neq t_k$ and

$$\begin{aligned}\Delta(X(t_k)\bar{M}) &= X(t_k^+)\bar{M} - X(t_k)\bar{M} \\ &= [X(t_k^+) - X(t_k)]\bar{M} \\ &= [\Delta X(t_k)]\bar{M} \\ &= B_k X(t_k)\bar{M}.\end{aligned}$$

2. Since $X(t)$ is a fundamental matrix, it is invertible for each t . Let $\bar{M} = X(0^+)^{-1}Y(0^+)$, and let $Z(t) \equiv Y(t) - X(t)\bar{M}$. Then $Z(0^+) = 0$ and

$$\begin{aligned}\frac{dZ(t)}{dt} &= \frac{dY(t)}{dt} - \frac{dX(t)}{dt}\bar{M} \\ &= P(t)Y(t) - P(t)X(t)\bar{M} \\ &= P(t)Z(t) \\ \Delta Z &= Y(t_k)^+ - X(t_k)^+\bar{M} - [Y(t_k) - X(t_k)\bar{M}] \\ &= \Delta Y(t_k) - \Delta X(t_k)\bar{M} \\ &= B_k Y(t_k) - B_k X(t_k)\bar{M} \\ &= B_k Z(t_k),\end{aligned}$$

so $Z(t) \equiv 0$ is the unique solution satisfying $Z(0^+) = 0$. Hence $Y(t) = X(t)\bar{M}$.

If Y is fundamental, then

$$\begin{aligned}\det \bar{M} &= \frac{1}{\det X(0^+)} \det Y(0^+) \\ &\neq 0.\end{aligned}$$

□

Theorem 15.11. Suppose conditions H1–H3 hold. Then each fundamental matrix of (15.4) can be represented in the form

$$X(t) = \varphi(t)e^{\Lambda t} \quad t \in \mathbb{R}$$

for a non-singular, T -periodic matrix $\varphi(\cdot) \in PC^1(\mathbb{R}, \mathbb{C}^{n \times n})$ and a constant matrix $\Lambda \in \mathbb{C}^{n \times n}$.

Proof. Let $X(t)$ be a fundamental matrix for (15.4) and define $Y(t) = X(t + T)$. Then, using H1, we have

$$\begin{aligned}\frac{dy_j(t)}{dt} &= \frac{dx_j(t+T)}{dt} \\ &= P(t+T)x_j(t+T) \\ &= P(t)y_j(t)\end{aligned}$$

for $t \neq t_k$, and, using H3,

$$\begin{aligned}\Delta y_j(t_k) &= \Delta x_j(t_k + T) \\ &= \Delta x_j(t_{k+q}) \\ &= B_{k+q}x_j(t_{k+q}) \\ &= B_kx_j(t_k + T) \\ &= B_ky_j(t_k)\end{aligned}$$

for each j . Also, $\det Y(0^+) = \det X(T^+) \neq 0$, since $x_1(t), \dots, x_n(t)$ are linearly independent in the interval $(0, \infty)$ and are hence independent in the interval (T, ∞) . Thus $Y(t)$ is also a fundamental matrix.

By the lemma, there exists a unique matrix $\bar{M} \in \mathbb{C}^{n \times n}$ such that

$$X(t + T) = X(t)\bar{M}$$

for all $t \in \mathbb{R}$. Set

$$\begin{aligned}\Lambda &= \frac{1}{T} \ln \bar{M} \\ \varphi(t) &= X(t)e^{-\Lambda t}.\end{aligned}$$

Hence $\varphi(t)$ is non-singular and belongs to the class $PC^1(\mathbb{R}, \mathbb{C}^{n \times n})$. Furthermore,

$$\begin{aligned}\varphi(t + T) &= X(t + T)e^{-\Lambda T}e^{-\Lambda t} \\ &= X(t)\bar{M}e^{-\Lambda T}e^{-\Lambda t} \\ &= X(t)e^{-\Lambda t} \\ &= \varphi(t),\end{aligned}$$

since $\bar{M} = e^{\Lambda T}$, by definition of Λ . Hence φ is T -periodic. □

This is a generalisation of the idea that you can solve systems of linear differential equations by finding the eigenvalues and eigenvectors. Once you have enough solutions — and those solutions aren't multiples or other combinations of one another — then you're done. So if you need n solutions and just

happen to have n eigenvalues, then you're done. If some of those eigenvalues are repeated, it gets a bit more nuanced, but this is the essence of it. Here we're using that idea for periodic systems, where both the time-dependent matrix and the impulsive effect are periodic.

To the fundamental matrix $X(t)$, there corresponds a unique matrix \bar{M} such that $X(t+T) = \bar{M}X(t)$ for all $t \in \mathbb{R}$. The eigenvalues μ_1, \dots, μ_n of \bar{M} are called *Floquet multipliers* of (15.4). The eigenvalues $\lambda_1, \dots, \lambda_n$ of Λ are called the *characteristic exponents* of (15.4).

Corollary 15.12. *Let conditions H1–H3 hold. Then $\mu \in \mathbb{C}$ is a Floquet multiplier of (15.4) if and only if there exists a nontrivial solution $\gamma(t)$ such that $\gamma(t+T) = \mu\gamma(t)$ for all $t \in \mathbb{R}$.*

The following theorem is from Bainov and Simeonov [4].

Theorem 15.13. *Suppose conditions H1–H3 hold. Then (15.4) is*

1. *stable if and only if all multipliers μ_j satisfy $|\mu_j| \leq 1$; and for those multipliers for which $|\mu_j| = 1$, the corresponding characteristic exponent (which has zero real part) is a simple zero of the characteristic polynomial of Λ ,*
2. *asymptotically stable if and only if all multipliers satisfy $|\mu_j| < 1$, and*
3. *unstable if $|\mu_j| > 1$ for some j .*

The idea here is that we convert the (semi)continuous system into a discrete-time system. For a continuous system, we have stability if the real part of each eigenvalue is negative. For discrete-time systems, the condition is that the eigenvalues are within the unit circle in the complex plane. The great thing about impulsive differential equations is that they lend themselves to discrete-time systems extremely well.

15.4.3 Orbital stability in \mathbb{R}^2

Consider the two dimensional autonomous system

$$\begin{aligned} \frac{dx}{dt} &= P(x, y), & \frac{dy}{dt} &= Q(x, y) & (x, y) \notin M \\ \Delta x &= a(x, y), & \Delta y &= b(x, y) & (x, y) \in M \end{aligned} \quad (15.5)$$

where $t \in \mathbb{R}$, and $M \subset \mathbb{R}^2$ is the set defined by the equation $\phi(x, y) = 0$.

Let $\gamma(t)$, $t \in \mathbb{R}$ be a solution of (15.5), with instants of impulsive effect t_k , such that

$$0 < t_1 < t_2 < \dots; \quad \lim_{k \rightarrow \infty} t_k = \infty,$$

and let $L_+ = \{u \in \mathbb{R}^2 : u = \gamma(t), t \in \mathbb{R}_+\}$. Denote by $J^+(t_0, z_0)$ the maximal interval of the form (t_0, ω) in which the solution $z(t; t_0, z_0)$ of (15.5) is defined.

For $y \in \mathbb{R}^2$, let $d(y, L^+) = \min_{u \in L^+} |y - u|$ and $B_\eta(\gamma(t_1))$ be the ball of radius η centred at $\gamma(t_1)$.

Definition 15.14. *The solution $z = \gamma(t)$ of (15.5) is called*

1. *orbitally stable if for all $\epsilon > 0$, $\eta > 0$ and $t_0 \in \mathbb{R}_+$, there exists $\delta > 0$ such that $d(z_0, L^+) < \delta$ and $z_0 \notin \bar{B}_\eta(\gamma(t_k)) \cup \bar{B}_\eta(\gamma(t_k^+))$ implies $d(z(t), L^+) < \epsilon$ for $t \in J^+(t_0, z_0)$ and $|t_0 - t_k| > \eta$, where $z(t) = z(t; t_0, z_0)$ is any solution of (15.5) for which $z(t_0^+; t_0, z_0) = z_0$.*
2. *orbitally attractive if for all $\epsilon > 0$, $\eta > 0$ and $t_0 \in \mathbb{R}_+$, there exists $\delta > 0$ and $T > 0$ such that $t_0 + T \in J^+(t_0, z_0)$ and $d(z_0, L^+) < \delta$ and $z_0 \notin \bar{B}_\eta(\gamma(t_k)) \cup \bar{B}_\eta(\gamma(t_k^+))$ implies $d(z(t), L^+) < \epsilon$ for $t \geq t_0 + T$, $t \in J^+(t_0, z_0)$ and $|t_0 - t_k| > \eta$, where $z(t) = z(t; t_0, z_0)$ is any solution of (15.5) for which $z(t_0^+; t_0, z_0) = z_0$.*
3. *orbitally asymptotically stable if it is orbitally stable and orbitally attractive.*

Definition 15.15. *The solution $z = \gamma(t)$ of (15.5) has the property of asymptotic phase if for all $\epsilon > 0$, $\eta > 0$ and $t_0 \in \mathbb{R}_+$, there exists $\delta > 0$, $c > 0$ and $T > |c|$ such that $t_0 + T \in J^+(t_0, z_0)$ and $|z_0 - \gamma(t_0)| < \delta$ implies $|z(t+c) - \gamma(t)| < \epsilon$ for $t \geq t_0 + T$, $t \in J^+(t_0, z_0)$ and $|t_0 - t_k| > \eta$, where $z(t+c) = z(t; t_0 - c, z_0)$ is any solution of (15.5) for which $z(t_0^+; t_0, z_0) = z_0$.*

The definition of stability for periodic orbits is a bit trickier than for equilibria, but the idea is the same: solutions that start close need to stay close for all time, except for small windows around the impulse points. The property of asymptotic phase ensures that the cycle time of each trajectory comes into phase with the period of the periodic orbit.

Suppose (15.5) has a T -periodic solution

$$\vec{p}(t) = \begin{bmatrix} \xi(t) \\ \eta(t) \end{bmatrix},$$

with

$$\left| \frac{d\xi}{dt} \right| + \left| \frac{d\eta}{dt} \right| \neq 0.$$

Assume further that the periodic solution $\vec{p}(t)$ has q instants of impulsive effect in the interval $(0, T)$. Since we have a periodic orbit, one multiplier is equal to 1. The other is calculated according to the formula

$$\mu_2 = \prod_{k=1}^q \Delta_k \exp \left[\int_0^T \left(\frac{\partial P}{\partial x}(\xi(t), \eta(t)) + \frac{\partial Q}{\partial y}(\xi(t), \eta(t)) \right) dt \right], \quad (15.6)$$

where

$$\Delta_k = \frac{P_+ \left(\frac{\partial b}{\partial y} \frac{\partial \phi}{\partial x} - \frac{\partial b}{\partial x} \frac{\partial \phi}{\partial y} + \frac{\partial \phi}{\partial x} \right) + Q_+ \left(\frac{\partial a}{\partial x} \frac{\partial \phi}{\partial y} - \frac{\partial a}{\partial y} \frac{\partial \phi}{\partial x} + \frac{\partial \phi}{\partial y} \right)}{P \frac{\partial \phi}{\partial x} + Q \frac{\partial \phi}{\partial y}}.$$

Here $P, Q, \frac{\partial a}{\partial x}, \frac{\partial b}{\partial x}, \frac{\partial a}{\partial y}, \frac{\partial b}{\partial y}, \frac{\partial \phi}{\partial x}$ and $\frac{\partial \phi}{\partial y}$ are computed at the point $(\xi(t_k), \eta(t_k))$, and $P_+ = P(\xi(t_k^+), \eta(t_k^+)), Q_+ = Q(\xi(t_k^+), \eta(t_k^+))$.

This formula may look crazy, but it's actually incredibly beautiful when you put all the pieces together, because things tend to collapse nicely into a simple result. We then have the following theorem, from Bainov and Simeonov [4], which is an analogue of the Poincaré criterion.

Theorem 15.16. *The solution $\vec{p}(t)$ of (15.5) is orbitally asymptotically stable and has the property of asymptotic phase if the multiplier μ_2 calculated by (15.6) satisfies the condition $|\mu_2| < 1$.*

Since we have a periodic orbit, one of the multipliers already has absolute value equal to 1. That means the stability comes down to calculating the other multiplier, which is handily given by the formula above.

Example. Consider the following two-dimensional model for cells infected with HIV virus and killer cells known as cytotoxic T-lymphocytes (CTLs) [34]. If we apply a CTL vaccine at regular times that boosts the number of killer cells, the model is given by

$$\begin{aligned} T' &= \pi - dT - pCT & t \neq t_k \\ C' &= \alpha CT - \delta C & t \neq t_k \\ \Delta C &= \tilde{C} & t = t_k, \end{aligned}$$

where t_k ($k = 1, 2, \dots$) are the vaccination times and \tilde{C} is the strength of the vaccine. Here T are the infected cells, produced at a constant rate π , depleted at a death rate d or by killer CTLs C at rate p . The CTLs are produced in response to infected cells at rate α and are depleted at rate δ .

It's not possible to solve this system explicitly, but we can find an implicit solution. Define

$$T_{\text{int}} = e^{\int_0^\tau (\alpha T(u) - \delta) du}.$$

This is a measure of the difference in CTLs between the beginning and the end of an impulsive cycle. For a vaccine administered at fixed times, the period is τ .

Separating variables and integrating, we have

$$\begin{aligned} \int_0^t \frac{dC}{C} &= \int_0^t (\alpha T(u) - \delta) du \\ C(t) &= C(0) e^{\int_0^t (\alpha T(u) - \delta) du}. \end{aligned}$$

It follows that

$$\begin{aligned} C(\tau^-) &= C(0) T_{\text{int}} \\ C(\tau^+) &= C(0) T_{\text{int}} + \tilde{C}. \end{aligned}$$

Suppose $C(\tau^+) = C(0)$. Then

$$\begin{aligned} C(0) &= \frac{\tilde{C}}{1 - T_{\text{int}}} \\ C(t) &= \frac{\tilde{C} e^{\int_0^t (\alpha T(u) - \delta) du}}{1 - T_{\text{int}}} \\ C(\tau^-) &= \frac{\tilde{C} T_{\text{int}}}{1 - T_{\text{int}}} \\ C(\tau^+) &= \frac{\tilde{C} T_{\text{int}}}{1 - T_{\text{int}}} + \tilde{C} = C(0). \end{aligned} \quad (15.7)$$

It follows that there is an impulsive periodic orbit with one impulse per orbit, whose endpoints satisfy

$$\begin{aligned} C(t_n^-) &= \frac{\tilde{C} T_{\text{int}}}{1 - T_{\text{int}}} \\ C(t_n^+) &= \frac{\tilde{C}}{1 - T_{\text{int}}}. \end{aligned}$$

Next we show that $T_{\text{int}} < 1$. Suppose $T_{\text{int}} \geq 1$. Then $\int_0^\tau (\alpha T(u) - \delta) du \geq 0$. It follows from (15.7) that $C(\tau^-) \geq C(0)$. But

$$\begin{aligned} C(\tau^+) &= C(\tau^-) + \tilde{C} \\ C(0) &= C(\tau^-) + \tilde{C} \\ &\geq C(0) + \tilde{C} \\ \Rightarrow \quad \tilde{C} &\leq 0. \end{aligned}$$

This is a contradiction, so $T_{\text{int}} < 1$. It follows that the periodic orbit exists.

Finally, we examine stability of the periodic orbit by calculating the non-trivial Floquet multiplier. Define

$$\begin{aligned} P(T, C) &= \pi - dT - pCT \\ Q(T, C) &= \alpha CT - \delta C \\ a(T, C) &= 0 \\ b(T, C) &= \tilde{C}, \end{aligned}$$

with the (differentiable) function ϕ implicitly defined by $\{\phi(T(t), C(t)) = 0 : t = t_k\}$. Let the impulsive orbit be described by $(\xi(t), \eta(t))$. Then we have

$$\begin{aligned} \xi(t_k^-) &= \xi(t_k^+) = T(\tau) \\ \eta(t_k^-) &= \frac{\tilde{C} T_{\text{int}}}{1 - T_{\text{int}}} \\ \eta(t_k^+) &= \frac{\tilde{C}}{1 - T_{\text{int}}}. \end{aligned}$$

Since $\xi(t_k^-) = \xi(t_k^+)$, we can write

$$\begin{aligned} P_+ &= \pi - d\xi(t_k^-) - p\eta(t_k^+)\xi(t_k^-) & Q_+ &= \alpha\xi(t_k^-)\eta(t_k^+) - \delta\eta(t_k^+) \\ P &= \pi - d\xi(t_k^-) - p\eta(t_k^-)\xi(t_k^-) & Q &= \alpha\xi(t_k^-)\eta(t_k^-) - \delta\eta(t_k^-). \end{aligned}$$

Then we have $Q = T_{\text{int}}Q_+$ and, since $\eta(t_k^-) = \eta(t_k^+)T_{\text{int}}$,

$$\begin{aligned} P_+T_{\text{int}} &= (\pi - d\xi(t_k^-))T_{\text{int}} - p\eta(t_k^+)T_{\text{int}}\xi(t_k^-) \\ &= (\pi - d\xi(t_k^-))T_{\text{int}} - p\eta(t_k^-)T_{\text{int}}\xi(t_k^-). \end{aligned}$$

From the nullclines, $T'(t_k^-) > 0$. Thus $\pi - d\xi(t_k^-) > p\eta(t_k^-)\xi(t_k^-) > 0$. Hence, since $T_{\text{int}} < 1$, $P_+T_{\text{int}} < P$. We then have

$$\begin{aligned} \Delta_1 &= \frac{P_+\frac{\partial\phi}{\partial T} + Q_+\frac{\partial\phi}{\partial C}}{P\frac{\partial\phi}{\partial T} + Q\frac{\partial\phi}{\partial C}} \\ &< \frac{1}{T_{\text{int}}}. \end{aligned}$$

(Note that we get this result without explicitly knowing what ϕ is.)

We thus have

$$\begin{aligned} \mu_2 &< \frac{1}{T_{\text{int}}} \exp \int_0^\tau \left(\frac{\partial P}{\partial T}(\xi(t), \eta(t)) + \frac{\partial Q}{\partial C}(\xi(t), \eta(t)) \right) dt \\ &= \frac{1}{T_{\text{int}}} \exp \int_0^\tau (-d - p\eta(t) + \alpha\xi(t) - \delta) dt \\ &= \frac{1}{T_{\text{int}}} \exp \left(- \int_0^\tau d + p\eta(t) dt \right) \exp \left(\int_0^\tau (\alpha\xi(t) - \delta) dt \right) \\ &= \frac{1}{T_{\text{int}}} \exp \left(- \int_0^\tau d + p\eta(t) dt \right) T_{\text{int}} \\ &< 1. \end{aligned}$$

Thus the nontrivial impulsive Floquet multiplier lies inside the unit circle, so the τ -periodic orbit is orbitally asymptotically stable and has the property of asymptotic phase.

The Floquet theory for impulsive semidynamical systems in \mathbb{R}^n , $n \geq 3$ is also developed in Bainov and Simeonov [3, 4], but calculation of the multipliers is much more difficult.

In practice, the theory is only useful in low-dimensional systems. If we are in \mathbb{R}^2 or the system can be reduced to a two-dimensional system, then we can apply the results in this section.

15.5 Lab work

1. Consider a mosquito-spraying model, where mosquitos grow at a constant rate m , but the effect of spraying is to reduce the mosquitos by a proportion r . This can be described by the impulsive differential equation

$$\begin{array}{ll} x' = m & t \neq t_k \\ \Delta x = -rx & t = t_k, \end{array}$$

with $\tau = t_{k+1} - t_k$ and $0 < r < 1$. Determine the impulsive orbit, and show that it is stable.

2. Consider the spread of MRSA (Methicillin-Resistant Staphylococcus Aureus) in prisons, where new (susceptible) prisoners are transferred in when there is room, while infected prisoners are quarantined at regular intervals.

a) Suppose both events happen simultaneously. Then

$$\begin{array}{ll} S' = -\beta SI & S \neq C \\ I' = \beta SI - dI & S \neq C \\ \Delta S = \lambda & S = C \\ \Delta I = -\alpha I & S = C \end{array}$$

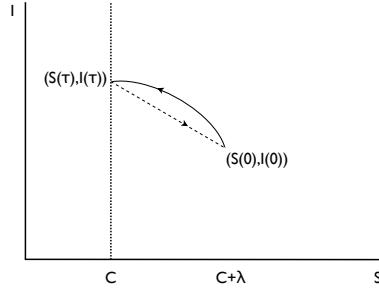


Fig. 15.5. The periodic orbit of first order.

i. Under what conditions will there be an impulsive periodic orbit?

Hint: see Figure 15.5.

ii. Show that this orbit is orbitally asymptotically stable.
 iii. Show that the period of the periodic orbit is given implicitly by

$$\tau = - \int_0^\tau \frac{1}{\beta S(I_0 + \lambda + c - S + \frac{d}{\beta} \ln \frac{S}{\lambda + c})} dS.$$

b) Let $n > 0$ be an integer and assume that new prisoners arrive at rate λ at each impulsive effect, while infected prisoners are quarantined only at the moments of impulsive effect τ_k whose ordinal number k is a multiple of n . That is,

$$\begin{aligned}
 S' &= -\beta SI & S &\neq C \\
 I' &= \beta SI - dI & S &\neq C \\
 \Delta S(\tau_k) &= \lambda & S &= C \\
 \Delta I(\tau_k) &= \begin{cases} 0 & \text{if } k \text{ is not divisible by } n \\ -\alpha I(\tau_k) & \text{if } k \text{ is divisible by } n \end{cases} & S &= C
 \end{aligned}$$

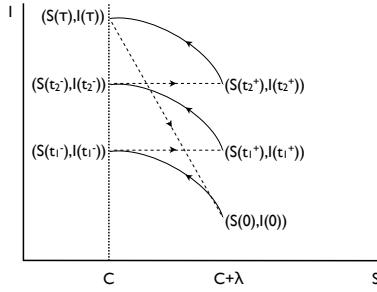


Fig. 15.6. The periodic orbit of third order.

i. Under what conditions will there be an impulsive periodic orbit?

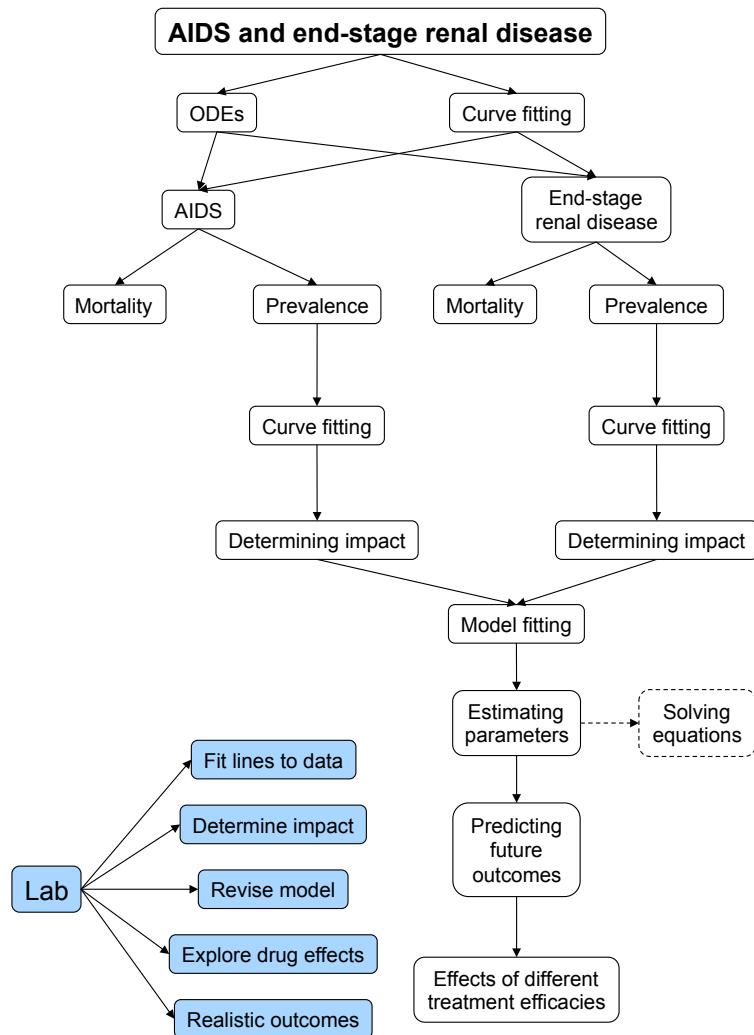
Hint: see Figure 15.6.

ii. Show that this orbit is orbitally asymptotically stable.
 iii. Show that the period of the periodic orbit is given implicitly by

$$\tau = - \sum_{i=0}^{n-1} \int_{t_i}^{t_{i+1}} \frac{1}{\beta S \left(I_i^+ + \lambda + c - S - \frac{d}{\beta} \ln \frac{S}{\lambda+c} \right)} dS.$$

Part III

Case studies



16

Application: AIDS and end-stage renal disease

AIDS is a big, bad disease. It's number eight in the top ten infectious diseases, which is pretty outstanding for a disease that's only been around since the eighties. This list measures the number of recorded deaths due to each disease and is as follows:

10. The Third plague, 12 million
9. Cocoliztli, 7–17 million
8. AIDS, 36 million
7. The Justinianic plague of 541–542AD, 15–100 million
6. Spanish flu, 50–100 million
5. The Black Death, 200 million
4. Measles, 200 million (in the past 165 years)
3. Smallpox, 300–500 million (in the 20th century alone)
2. TB, 1 billion (in the past 200 years)
1. Malaria, 5–50 billion

How many people have there been in total? About 110 billion, which means malaria has killed up to 45% of everyone who ever lived, making it the all-time biggest killer of humans. Check out my popular-science book “The Top Ten Diseases of All Time” for the details [32].

Many patients with AIDS go on to develop end-stage renal disease, which is basically kidney failure. AIDS itself doesn't kill you, but it creates a climate for opportunistic infections that will. In the United States, end-stage renal disease is particularly prevalent in African American patients and the numbers have risen alarmingly throughout the early years of the 21st century. Thus we'll focus on this subset of the population.

The introduction of antiretroviral drugs has drastically changed the face of HIV. While there is still no cure, ART (antiretroviral therapy, a.k.a. the triple-drug “cocktail”) has reduced the number of AIDS deaths and made HIV a disease that it's possible to live with. What's not so clear is whether ART has had an effect on the overall prevalence of AIDS in general or end-stage renal disease in particular.

In this chapter, we want to combine the epidemic models of Chapters 5 through 7 with the data-fitting methods from Chapters 3 and 4. (We're not using chapters 8 or 9, since we're only looking temporally, not spatially.) We'd like to answer the following questions:

1. Has ART had an impact on the prevalence rate of AIDS?
2. Has ART had an impact on the prevalence rate of end-stage renal disease?
3. If aggressive treatment is initiated now, with different effects, what will the long-term outcome be?

Thus we'll need to a) formulate a model, b) fit parameters to data, c) draw conclusions and d) predict the future. This combines the various strands of applying models to biological problems and will incorporate some real-world data.

16.1 Determining AIDS prevalence

Has ART had an impact on the number of deaths from AIDS? First we need some data, courtesy of the US CDC (United States Centers for Disease Control):

1991	10475
1992	11946
1993	15460
1994	17844
1995	18971
1996	15909
1997	10333
1998	8744
1999	9097
2000	8723
2001	9085
2002	8927
2003	9077
2004	9302
2005	8562

These are the number of deaths due to AIDS for African Americans in the United States. To see it a bit more clearly, take a look at Figure 16.1. Seems pretty obvious, doesn't it? The death rate was increasing until 1995, then underwent a massive decrease before plateauing out. We don't need to fit curves to the data to see that there's a very obvious change after the introduction of ART.

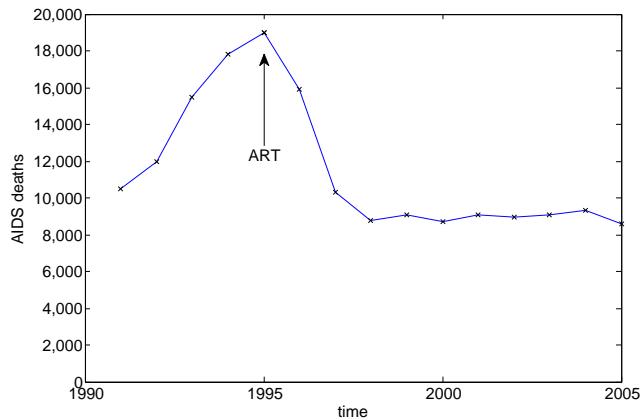


Fig. 16.1. The annual number of deaths due to AIDS for African Americans in the United States

What about prevalence? Since not as many people are dying, but other people are progressing from HIV to AIDS, we expect prevalence to be increasing, but perhaps not as sharply as it did before ART. The CDC prevalence data is:

1991	14561
1992	15897
1993	60649
1994	71847
1995	81317
1996	92319
1997	105464
1998	117890
1999	112483
2000	121903
2001	181475
2002	193814
2003	204466
2004	214017
2005	225270

This is shown in Figure 16.2.

It's tough to draw any conclusions from these data. Prevalence has continued to rise, but has it continued to rise as quickly, or has ART slowed down the rate of increase? Let's be a bit formal about this and construct a *null hypothesis*. This is a hypothesis that states that there's been no change. If there has been a change, then we'll reject this hypothesis. Our null hypothesis

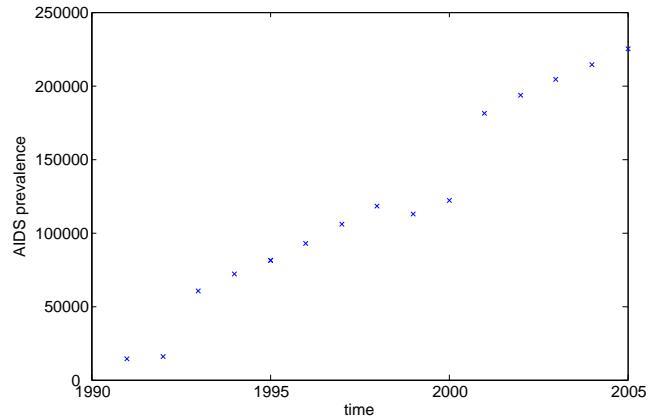


Fig. 16.2. Prevalence of AIDS among African Americans in the United States.

is thus: ART has had no effect on the prevalence of AIDS. Note that we don't specify anything about whether ART has increased or decreased the prevalence; it may have done either (which would be interesting, regardless) or it may have had no effect.

How could we test our hypothesis? One way would be to fit curves to the pre-ART and post-ART data separately and see what happens. The data are approximately linear, so we only have to fit straight lines. If we do that, we have the situation in Figure 16.3.

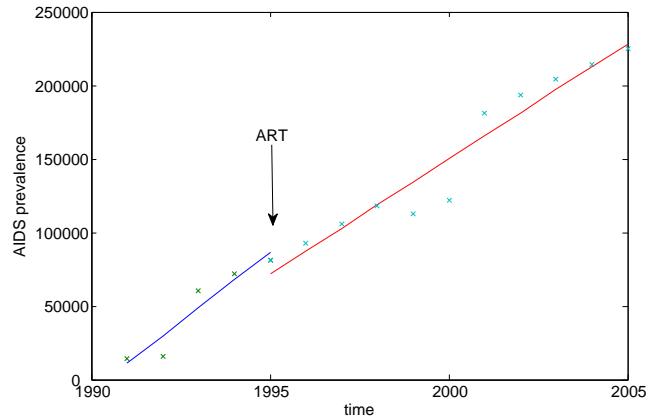


Fig. 16.3. Linear fit to pre-ART and post-ART AIDS data.

What do you think? Looks somewhat convincing. The slope of the first line is steeper than the slope of the second line, so maybe ART has slowed the rate of increase of the prevalence. Of course, at this point we should be asking ourselves how good the fit is, but the two lines have regression coefficients of 0.949 and 0.967, respectively. So these lines are good fits to these data.

However, that's only half the story. We need to compare this to our null hypothesis of no effect on the prevalence. To do this, we could fit a line to the entire data set. See Figure 16.4.

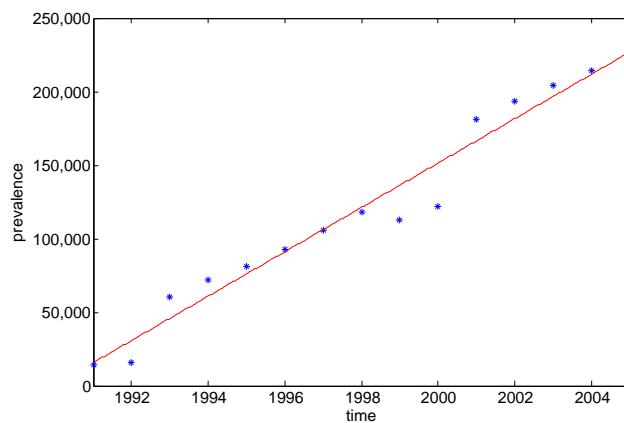


Fig. 16.4. Linear fit to the entire AIDS data set.

Now what do you think? Is this a better or worse fit? The eye can't tell, so we have to rely on the regression coefficient. In this case $r = 0.981$, so this is actually a better fit than either of the two lines in Figure 16.3! This means we can't reject the null hypothesis. In other words, ART has had no significant impact on the prevalence of AIDS among African Americans in the United States. So we've answered our first question.

16.2 Determining end-stage renal disease prevalence

Let's look at end-stage renal disease data, courtesy of the United States Renal Data System.¹ The mortality data, showing the number of deaths in end-

¹ The data reported here have been supplied by the United States Renal Data System (USRDS). The interpretation and reporting of these data are the responsibility of the author and in no way should be seen as an official policy or interpretation of the U.S. government.

stage renal disease patients, with AIDS nephropathy as primary cause of renal failure, is

1991	88
1992	126
1993	159
1994	176
1995	255
1996	185
1997	120
1998	141
1999	149
2000	131
2001	126
2002	143
2003	135
2004	128

This is shown in Figure 16.5.

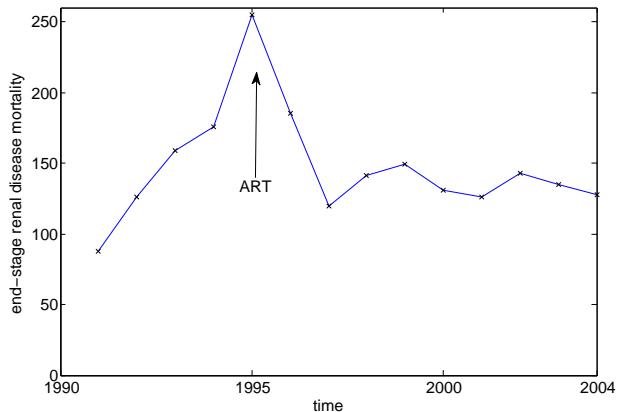


Fig. 16.5. The number of deaths in African American end-stage renal disease patients with AIDS nephropathy as primary cause of renal failure in the United States.

Once again, we don't need mathematical models to tell us that ART has had an impact on the mortality.

Next, we turn to prevalence. The prevalence of end-stage renal disease data is:

1991	350
1992	473
1993	644
1994	709
1995	934
1996	1038
1997	1287
1998	1521
1999	1774
2000	1989
2001	2181
2002	2296
2003	2414
2004	2449

This is shown in Figure 16.6.

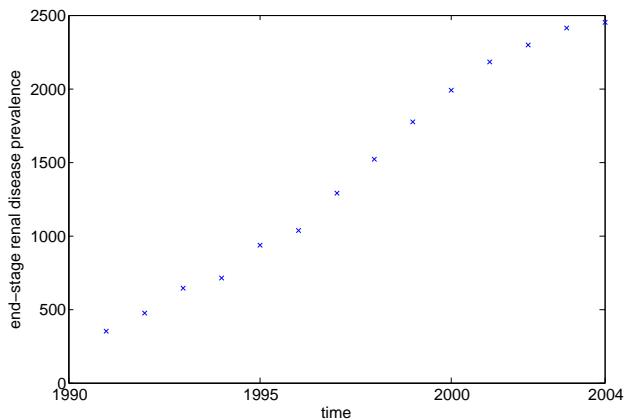


Fig. 16.6. Prevalence of end-stage renal disease among African Americans in the United States.

Once again, it's hard to tell if ART has had any impact on the prevalence. Let's do the same trick again: fit linear curves to pre-ART and post-ART data separately and then to the entire data set. See Figures 16.7 and 16.8, respectively.

This time around, it looks as though ART may have actually increased the prevalence of end-stage renal disease. That's not necessarily out of the question, given that many more people are alive because of ART than would be otherwise, so it's possible that the prevalence could increase. This is why we're

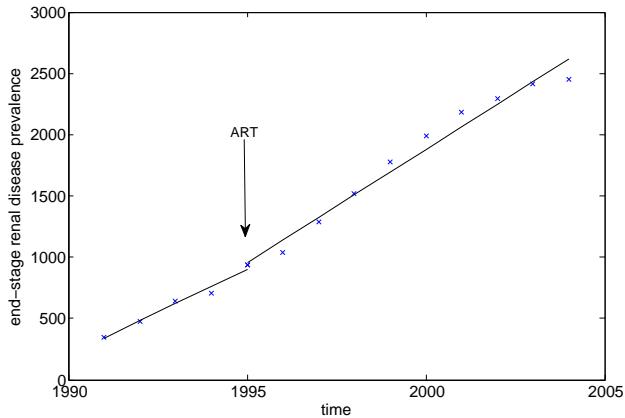


Fig. 16.7. Linear fit to pre-ART and post-ART end-stage renal disease data.

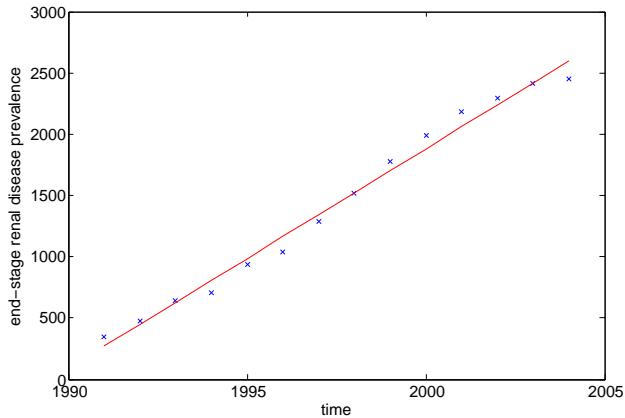


Fig. 16.8. Linear fit to the entire end-stage renal disease data.

interested in asking whether ART has had an impact, rather than specifying what that impact is.

The regression coefficients for the two linear fits in Figure 16.7 are $r = 0.98887$ and $r = 0.98683$. The regression coefficient for the linear fit in Figure 16.7 is $r = 0.99352$. Thus the line that fits the overall data is a better fit than each of the lines which fit the two subsets. Once again we have to conclude that ART has had no impact on the prevalence of end-stage renal disease. We've answered our second question.

16.3 Model fitting

What have we done in the last two sections? We've made choices about which models best fit the data. True, these were simple, linear models, but they're models nonetheless. We can use these linear fits to estimate parameters and construct a more complex ODE model. Note that we could have simply looked at the data and concluded, with pretty reasonable confidence, that the fits were linear, regardless of the ART issue. So, if the first half of this chapter wasn't to your liking, you can just start here. The only thing we need to know is that the data are approximately linear (which should be pretty clear).

Let's think about how to construct such a model. We need to know what form the model will take, or else we won't know what parameters we need to estimate. There are many things we could try, but let's start with the simplest version. We only have two variables of interest: the prevalence of AIDS and the prevalence of end-stage renal disease. The latter doesn't cause the former, so it means we can consider AIDS prevalence in isolation.

Let's start with the easier case and deal with the prevalence of AIDS. Since we now have some faith that it's a linear fit, we can use that to construct a simple linear differential equation. The prevalence is clearly increasing (see Figure 16.4), so the derivative will be positive. Thus we could assume the prevalence of AIDS cases satisfies the differential equation

$$\frac{dA}{dt} = g. \quad (16.1)$$

This is just about the simplest differential equation ever and one we could easily solve. But let's leave that for a moment and figure out the parameters we need. Clearly we need g , but we'll also need an initial condition, A_0 . From Figure 16.4, the intercept is $A_0 = 14959.04166$ and the slope is $g = 15133.20357$. (Of course, this assumes that time starts at 1991, so we're really transposing the x axis by 1991.)

The second case is a bit trickier. A proportion s of the population with AIDS will progress to end-stage renal disease; those with end-stage renal disease will die at rate δ , proportional to the end-stage renal disease prevalence. This leads to the second equation

$$\frac{dN}{dt} = sA - \delta N. \quad (16.2)$$

We can't solve this directly, of course, because it depends on the solution to equation (16.1). However, we don't need to solve the equations to estimate the parameters. From Figure 16.8, we have $N_0 = 268.11$ (remembering that we're starting at 1991), and the slope of the graph is 179.18.

How can we use the slope of the line to estimate parameters? The slope is the derivative, so this says that

$$\frac{dN}{dt} \approx 179.18$$

$$sA - \delta N = 179.18.$$

Since A and N are variables, we can pick whichever values of them we like. However, remember that the linear fit is just an approximation, so we should pick values that are closest to the line. The years 1997 and 1998 are the two years with values closest to the line, in both Figures 16.4 and 16.8. We need data from two years, because we're solving for two parameters, s and δ .

We thus have

$$s(105464) - \delta(1287) = 179.18$$

$$s(117890) - \delta(1521) = 179.18.$$

We can convert this to matrix form and use Matlab to solve it:

$$\begin{bmatrix} s \\ \delta \end{bmatrix} = \begin{bmatrix} 105464 & -1287 \\ 117890 & -1521 \end{bmatrix}^{-1} \begin{bmatrix} 179.18 \\ 179.18 \end{bmatrix}$$

$$= \begin{bmatrix} 0.0048 \\ 0.2563 \end{bmatrix}.$$

We now have estimates for all three parameters (g , s and δ), plus the two initial conditions (A_0 and N_0). Notice that we didn't solve either equation, even though the first is straightforward and the second is doable, once you've solved the first. This is the nice thing about fitting parameters to models: you can do it without having to find the solution. Can we actually find the analytic solution? Yes, but the details are a bit technical. See Appendix L if you're interested.

16.4 Using the model to predict future outcomes

Now that we have our model, what can we do with it? The best thing that models can do for us is to predict the future. The downside, of course, being that we can't be too confident about such predictions without actual evidence... but if we wait for such evidence to arrive, it won't be the future any more.

One way to deal with this is to make a range of predictions, depending on what you want from the model. In this case, we'll consider the effects of what happens if aggressive treatment is initiated. Currently, treatment hasn't done very much to slow the prevalence of end-stage renal disease, but treatment also hasn't been applied as widely as it could, especially in disadvantaged groups like African Americans. So let's see if treatment can eliminate end-stage renal disease.

We can represent the effects of treatment blocking by the term $(1 - h)$, where $h = 0$ specifies no change in progression to end-stage renal disease while

$h = 1$ represents therapy that completely blocks the progression to end-stage renal disease. Thus equation 16.2 becomes

$$\frac{dN}{dt} = s(1 - h)A - \delta N.$$

If $h = 0$, we have the same equation as (16.2), so the prevalence of end-stage renal disease would continue to grow. If $h = 1$, then the prevalence of end-stage renal disease would be uniformly decreasing, so eventually the disease would be eliminated.

Of course, we're likely to be somewhere in between the two. We don't know where, so let's consider a number of possibilities: $h = 0.38, 0.65, 0.80, 0.95, 1$. This way, we'll get an idea of the effect of different treatment options.

To calculate this, we have two possibilities. We could use the explicit solution found in Appendix L, or we could use `ode45` in Matlab. Either is fine, although in general the equations won't be solvable, so the numerical solver is probably more useful. The results are shown in Figure 16.9.

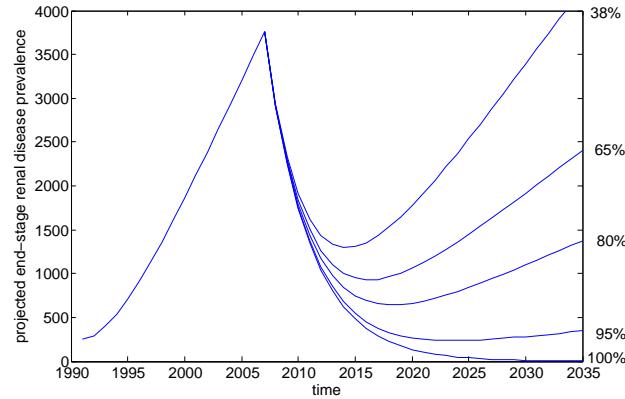


Fig. 16.9. ART-blocking effects on disease progression.

What does this tell us? Most strikingly, it tells us that the only way to eliminate end-stage renal disease is to have 100% effective therapy. All other therapies will have an initial dip, then rise in numbers again. Even 95% effective therapy will eventually lead to an increase in prevalence. Only perfectly efficacious therapy will eliminate the disease. Unfortunately, that's impossible to achieve, in practical terms.

However, that doesn't mean our model doesn't tell us anything useful. For one thing, even therapy that was only 38% effective would not result in an increased prevalence for about 25 years. While it's obviously best to have the most effective therapy possible, our model tells us that there are short-term

gains to be made, even from less-effective therapy. At the least, instituting therapy that's 38% effective now means we've got 25 years to come up with something new before we see a return to the current prevalence levels. If we can institute therapy that's moderately effective (say 65%), then we've got about 40 years to develop new strategies. That's certainly something worth attempting, even if we can't actually eradicate the disease.

This sort of thing is done quite a lot in disease control, especially for diseases like HIV where there is no cure. We may not be able to "solve" the problem of curing the disease, but we can at least hold back the disease in the interim, while we work on other treatments. Part of the issue is expectations: if we're hoping that our therapy will eradicate the disease, then we'd be in for a shock. On the other hand, modelling tells us that putting all our efforts into perfecting therapy might not be the best strategy, since we'll never achieve perfect therapy anyway, while even fairly ineffective therapy can do a lot of good in the meantime.

In this way, we use our models to make important decisions about whether to proceed or not, knowing the likely outcomes. And we can do it with nothing more sophisticated than linear regression and simple ODEs.

16.5 Lab work

By this point, you should have all the tools to do this lab. So this is a good chance to practice everything we've learnt in previous labs. If you're rusty, have another look at the labs in Chapters 5 and 3.

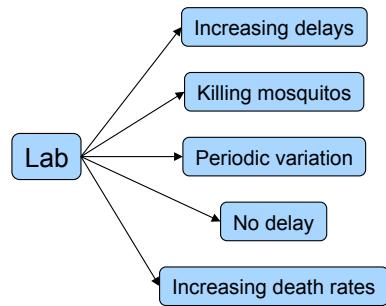
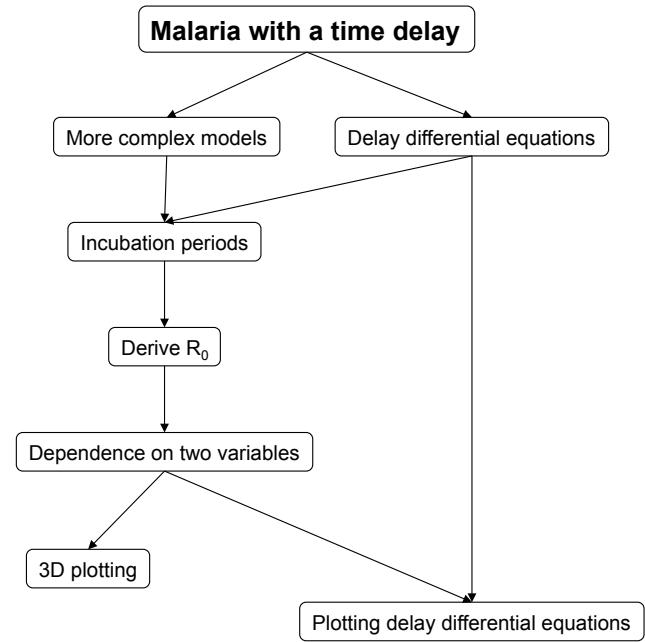
Here, we'll fit curves to the mortality data detailed earlier in this chapter, in order to justify our "obvious" conclusions. Then we'll see what happens when we tweak the ODE model a bit, to add another level of realism.

16.5.1 Exercises

1. Using the mortality data for both AIDS and end-stage renal disease (on pages 200 and 204, respectively), use the "Tools → Basic Plotting" command to fit straight lines to a) the pre-ART and b) post-ART data (Hint: enter the data in two separate blocks, counting the 1995 entry in each). Find the regression coefficient r for each.
2. Now fit a straight line to the entire data set for each (be sure not to count 1995 twice, though) and find the regression coefficients, r .
3. Which is a better fit? You might want to compare r^2 values instead. (Can you guess why?) What does this tell us?
4. When considering the effects of treatment, we assumed that the prevalence of end-stage renal disease would be affected by treatment, which meant we had to adjust equation (16.2) by a factor $(1 - h)$. However, we made

no adjustment to the prevalence of AIDS. What if treatment reduces the prevalence of AIDS as well? Adjust equation (16.1) in a similar way to include the effect of treatment j .

5. Now we'll explore the effects of this new model on the long-term outcome. Use Matlab to plot the prevalence of end-stage renal disease with no treatment change from 1991 to 2005 ($h = j = 0$). Then explore different values of j for different values of h ; ie $h = 0.38, 0.65, 0.8, 0.95, 1$ for each value $j = 0.2, 0.5, 0.8$. (You should have three separate graphs, each with five curves like Figure 16.9.) Note: you can use `ode45` or adjust the solution found in Appendix L.
6. What happens as $j \rightarrow 1$? Explain in words.
7. What about $j > 1$? Explain how this might occur biologically and plot some scenarios.
8. What happens eventually for $j > 1$? Plot some scenarios. How realistic is this?



Application: Malaria with a time delay

Models of malaria can take many forms, as with most diseases. Many biological factors were ignored in our model in Chapter 6, so there are many possibilities to complexify the basics. For example, malaria typically incubates in both humans and mosquitos (but for different lengths of time), which was something we'd ignored. So one adjustment to the model might be to include this factor.

17.1 A delay differential equation model of malaria

Consider the Ross–MacDonald malaria model with a delay, as given in Ruan *et al.* (2008):

$$\begin{aligned}\frac{dx}{dt} &= -rx(t) + abm[1 - x(t - \tau_1)]y(t - \tau_1)e^{-r\tau_1} \\ \frac{dy}{dt} &= -\mu y(t) + acx(t - \tau_2)[1 - y(t - \tau_2)]e^{-\mu\tau_2},\end{aligned}$$

with variables and parameters given in Table 17.1.

This model is a system of delay differential equations, which we saw briefly in Chapter 5. This makes things harder... but we're also only modelling infected humans and infected mosquitos, so that makes things easier again. Modelling is often a trade-off between factors you want to focus on and the complexity of the overall system: put in more details of one factor and you might have to ignore other factors, if you want to keep your models reasonable.

Let's derive an R_0 threshold, not analytically (which is tedious), but by heuristically thinking about it. For a primary human with recovery rate r , the average time spent in an infectious state is $1/r$. During this time, since the incubation period in humans has duration τ_1 , the average number of mosquito bites received from m susceptible mosquitos each with biting rate a gives a total of $R_H = acme^{-r\tau_1}/r$ mosquitos infected by the primary human case.

Each of these mosquitos survives for an average time $1/\mu$ and, with another incubation period τ_2 in mosquitos, makes a total of $R_M = abe^{-\mu\tau_2}/\mu$ infectious

Parameter	Description	Value
$x(t)$	proportion of infected humans	(variable)
$y(t)$	proportion of infected mosquitos	(variable)
m	ratio of mosquitos to humans	2
a	biting rate on a human per mosquito	0.2–0.5/day
b	infected mosquito to human transmission efficiency	0.5
c	infected human to mosquito transmission efficiency	0.5
r	per capita human recovery rate	0.01–0.05/day
μ	per capita mortality rate of mosquitos	0.05–0.5/day
τ_1	incubation period in humans	10–100 days
τ_2	incubation period in mosquitos	5–15 days

Table 17.1. Variables and parameters for the Ross–Macdonald model of malaria with delay.

bites. The total number of secondary cases is then

$$R_0 = R_H \times R_M.$$

See Figure 6.1 on Page 75 for an example illustrating this. Thus

$$R_0 = \frac{a^2 b c m e^{-r\tau_1} e^{-\mu\tau_2}}{r\mu}.$$

Note that a appears twice in this expression, since the mosquito-biting rate controls transmission from humans to mosquitos and from mosquitos to humans.

Notice that R_0 depends on both incubation periods τ_1 and τ_2 . To see how it depends on them, we could either plot R_0 against each one individually... or we could plot R_0 against both simultaneously. To do this, we'll need to plot in 3D, but Matlab can handle that quite well, as we saw in Chapter 9. The following code will produce Figure 17.1:

```

clear all
a=0.2;
b=0.5;
c=0.5;
m=2;
r=0.05;
mu=0.05;
tau1=0:0.5:40;
tau2=0:0.5:20;
[X,Y]=meshgrid(tau1,tau2);
R0=a.^2.*b.*c.*m.*exp(-r.*X).*exp(-mu.*Y)./(r.*mu);
Z=R0./R0;
mesh(X,Y,R0)

```

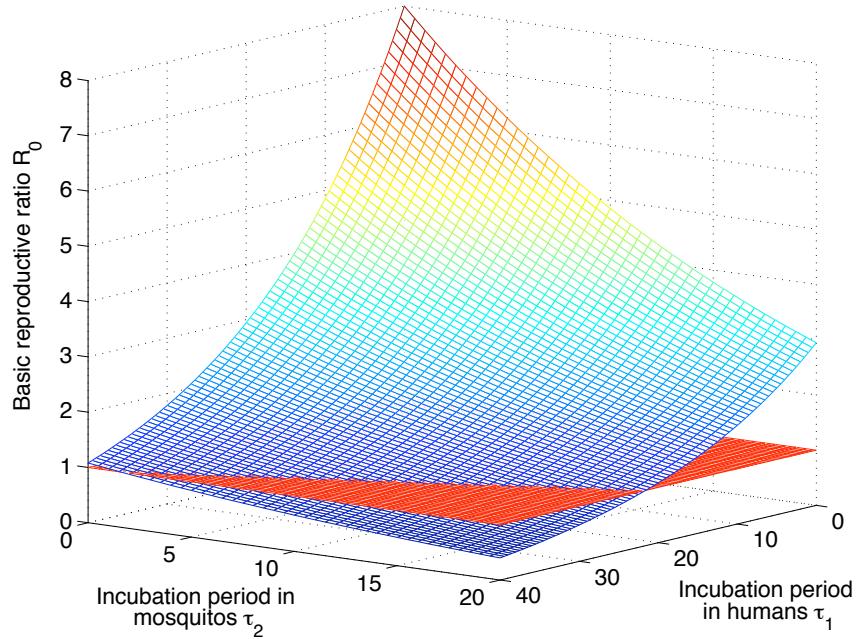


Fig. 17.1. A example of a two-stage basic reproductive ratio.

```

hold on
mesh(X,Y,Z)
hold off
xlabel('Incubation period in humans \tau_1')
ylabel('Incubation period in mosquitos \tau_2')
zlabel('Basic reproductive ratio R_0')

```

Go to **Tools** → **Rotate3D** (as we did in Chapter 9). Play with this until it looks like Figure 17.1. The flat surface is when $R_0 = 1$. This means that if both τ_1 and τ_2 are very large, then the disease would die out. This makes sense: if the disease incubates for a very long time, that would slow down transmission. Likely delays, however, are something like $\tau_1 = 15$ and $\tau_2 = 9$. Looking at our graph, this would suggest that $R_0 \approx 3$.

Is this good or bad? Well, it's not great, because we still have disease (since $R_0 > 1$). But it's a lot better than $R_0 = 8$, which is what we'd have if $\tau_1 = \tau_2 = 0$. So it seems that incubation is a good thing, because, although it doesn't lead to eradication, it nevertheless makes the disease less virulent.

How do we plot delay differential equations in Matlab? It's similar to regular ODEs; it just requires an extra function file to specify the history before

time began (necessary because of the delay). First we specify our differential equations:

```
function dydt = malariadelayf(t,y,Z)
a=0.2;
b=0.5;
c=0.5;
m=2;
r=0.05;
mu=0.05;
tau1=15;
tau2=9;
ylag1=Z(:,1);
ylag2=Z(:,2);
dydt=[-r.*y(1)+a.*b.*m.*(1-ylag1(1)).*ylag1(2).*exp(-r.*tau1)
      -mu.*y(2)+a.*c.*ylag2(1).*(1-ylag2(2)).*exp(-mu.*tau2)];
```

(Note that the `dydt=` term is actually two lines. You need a hard return at the end of the penultimate line in the code, because there are two components, not one.)

Here `ylag1` and `ylag2` are the two delays. We have to deal with them as separate entities. The following code specifies the initial conditions:

```
function s=malariadelayhist(t)
s=[0.01;0.01];
```

Again, this is a function file, so save it under the same name you give it (`malariadelayhist` in this example). This function specifies that the initial conditions for $-15 \leq u \leq 0$ are $x(u) = 0.01$ and $y(u) = 0.01$ (we need to give initial conditions backwards in time because the solution has a delay; thus, at time zero, the solution needs to know what happened at $t = 0 - \tau_1 = -15$ and $t = 0 - \tau_2 = -9$). Finally, this is the code we'll run:

```
clear all
sol=dde23(@malariadelayf,[15, 9],@malariadelayhist,[0, 500]);
figure;
plot(sol.x,sol.y)
xlabel('time t');
ylabel('solution y');
```

The delays are given here in the `[15, 9]` vector, and we have to call both of our function files. This should produce the solution in Figure 17.2.

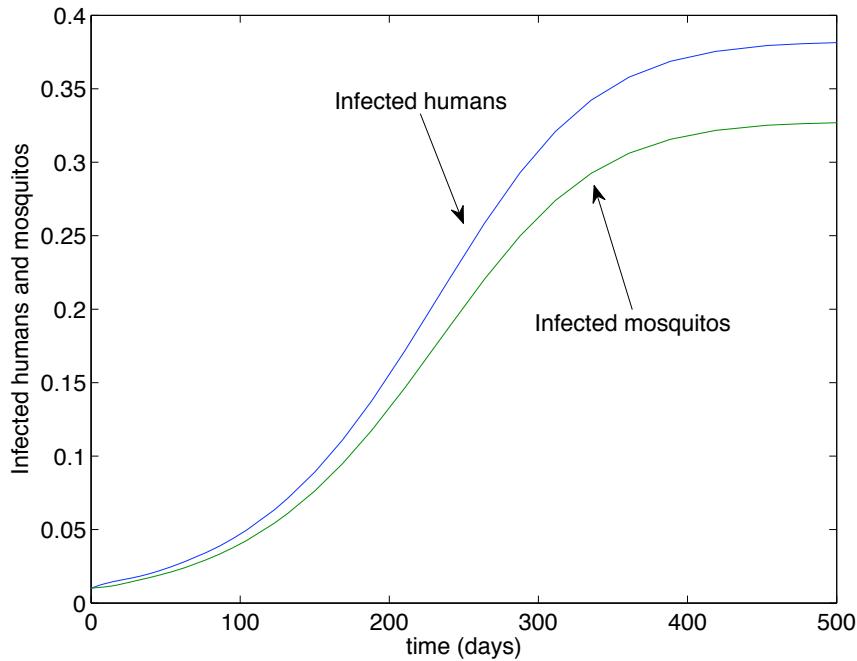


Fig. 17.2. Infected humans and mosquitos.

What's happening here is that there are two equilibria: an unstable disease-free equilibrium and a stable endemic equilibrium. Notice, however, that we didn't find these analytically. This is because things get much more difficult with delay differential equations. However, the nice part is that the Matlab coding doesn't change very much.

17.2 Lab work

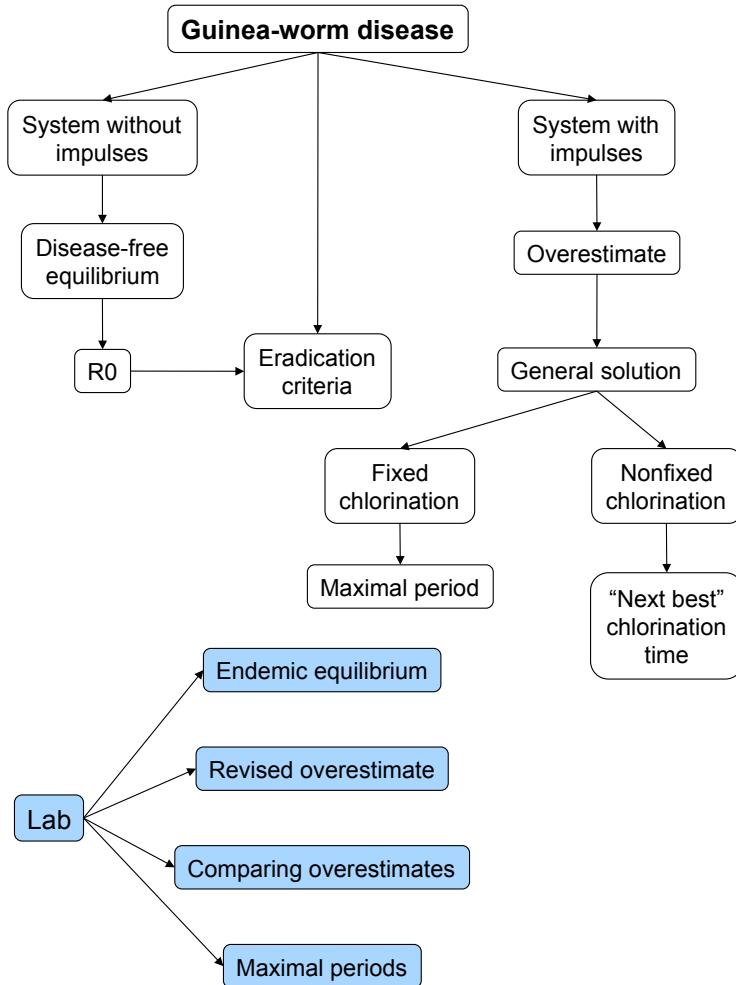
17.2.1 Exercises

1. Run the code again with $\tau_1 = 15, 18, 21, 24$ and $\tau_2 = 9$.
2. Run the code again with $\tau_1 = 15$ and $\tau_2 = 9, 12, 15, 18$.
3. Does this fit in with your expectations, based on Figure 17.1? Which incubation delay has more effect on reducing the disease?
4. Suppose you treated malaria by changing the death rate of mosquitos to $\mu = 0.105$. What happens now?

5. Can you explain the “bump” that you see?
6. Now suppose the number of mosquitos per human is a periodic function $m = 1.5 + \sin(2\pi t)$. What happens now? (You may want to run the code for longer; say 1500 days).
7. Show that if there is no incubation period for either humans or mosquitos, then the delay model reduces to the classic Ross–Macdonald model

$$\begin{aligned}\frac{dx}{dt} &= -rx(t) + abm[1 - x(t)]y(t) \\ \frac{dy}{dt} &= -\mu y(t) + acx(t)[1 - y(t)].\end{aligned}$$

8. Find the two equilibria of this system.
9. Show that, in this case, $R_0 = a^2bcm/\mu r$.
10. What happens if the death rates for humans and mosquitos are large (ie $\mu r > a^2bcm$)? Explain the biological implications of this.



Application: Guinea-worm disease

18.1 Introduction

One of the most fascinating diseases that almost nobody has heard of (but hopefully soon will!) is Guinea-worm disease (GWD). This is a neglected tropical disease, one that's spread via drinking water (and only via drinking water, which matters). Unfortunately, there is no drug to treat GWD, and there is no vaccine either. Miraculously, however, GWD is about to be eradicated, making it the first parasitic disease to be eradicated and the first to be eradicated without biomedical interventions. This is largely thanks to the efforts of one man. (See if you can guess who before the reveal.) So how can you eradicate a disease without a drug, vaccine or immunity?

GWD has been with us since antiquity: it's mentioned in the Bible, and Egyptian mummies suffered from it. Essentially, the parasite attaches itself to a water flea; you drink the flea; and your stomach acid dissolves the flea, leaving the parasite free to invade your body. Because of gravity, it usually makes its way to the foot, where it lives for an entire year.

After a year, your foot is burning and itching, so you put it in the water. And if your village only has one water source, then that often ends up being the drinking water. At this point, the fully grown worm bursts out of your foot, spraying forth 100,000 parasites and hence restarting the process.

In the 1950s, GWD affected 50 million people across most of Africa, Asia and the Middle East. Yet today it's on the verge of being eradicated, with less than 25 human cases reported in 2016, in just three African countries. This ancient scourge is almost gone. So what happened?

Before we reveal the answer, let's think about how you might eradicate a water-borne disease (i.e., a disease transmitted through contaminated water). Possibilities are a vaccine, drugs that treat symptoms, chemicals that kill the parasite, better hygiene or education that changes people's behaviour.

This chapter shows how to build a simple model, drawing on Chapter 5, and illustrates how impulses can be used to make predictions, a la Chapter 15.

18.2 The model

To create our mathematical model, we need to keep track of what comes in and what goes out. In the case of Guinea-worm disease, we divide the population of humans into three subcategories. GWD is not lethal, so each time we speak of death rate, it is the usual background death rate. The first category is that of susceptible individuals. Three things can happen to them: they are born, become infected or die. Similarly, infected individuals become infectious or die. Infectious individuals recover or die. We also have a population of worms: the parasite is born when infectious individuals put their foot in the drinking water (because fresh water produces relief) and dies shortly thereafter.

Denote susceptible individuals by S , exposed individuals by E and infected individuals by I . The number of larvae in the water is denoted W . The human birth rate is Π , the infection rate is β , the rate of worm emergence is α , the recovery rate is κ , and the death rate is μ . Infected individuals produce new larvae at rate γ , and the larvae are naturally cleared from the water at rate μ_W . Although water fleas act as an intermediate host, carrying the nematode until human digestion, we conflate the larvae and the fleas, in order to keep the model tractable.

Interventions include filtration, education or chlorination of the water supply. Although “education” is a complex term, encompassing a multitude of interventions, we consider education to refer directly to teaching people not to put their infected limbs in the water supply, in line with established behaviour-change programs for tackling GWD. Thus, we consider that an increase in education will have the direct effect of reducing the parasite birth rate, hence reducing γ . Likewise, by “filtration”, we mean a method that reduces the ability of the parasite to infect a human host, thus reducing β . Chlorination has the effect of increasing the death rate of the parasite, thus increasing μ_W .

Our mathematical model is thus

$$\begin{aligned} S' &= \Pi - \beta SW - \mu S + \kappa I & t \neq t_k \\ E' &= \beta SW - \alpha E - \mu E & t \neq t_k \\ I' &= \alpha E - \kappa I - \mu I & t \neq t_k \\ W' &= \gamma I - \mu_W W & t \neq t_k \\ \Delta W &= -rW & t = t_k. \end{aligned}$$

See Figure 18.1. Note that S , E , I and W are nonnegative. Furthermore, since these quantities are averages, we do not assume that individuals are necessarily infected with only one worm at a time. People have been observed to have up to seven worms at once.

We use mass-action transmission, since the interaction between parasites in the water and humans involves drinking parasite-laden water. Thus, since everyone in the village usually drinks from a single source, each human has roughly equal chance of encountering the parasite.

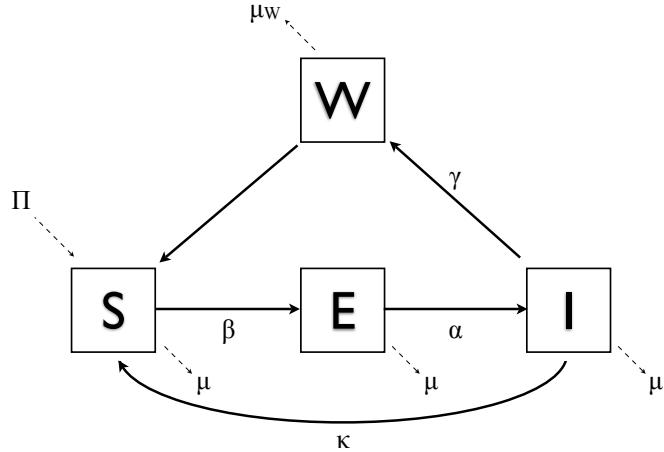


Fig. 18.1. The Guinea-worm disease model

18.3 The system without impulses

First, we shall analyse the corresponding system of ODEs. Note that

$$S' + E' + I' = \Pi - \mu(S + E + I).$$

Thus

$$S + E + I \leq \frac{\Pi}{\mu}.$$

Hence

$$\begin{aligned} I' &\leq \frac{\alpha \Pi}{\mu} - (\kappa + \mu)I \\ I &\leq \frac{\alpha \Pi}{\mu(\kappa + \mu)} + \left[I(0) - \frac{\alpha \Pi}{\mu(\kappa + \mu)} \right] e^{-(\kappa + \mu)t}. \end{aligned}$$

Since κ is large, the exponential term is small, so we have

$$I \leq \frac{\alpha \Pi}{\mu(\kappa + \mu)}.$$

It follows that

$$W' \leq \frac{\alpha \Pi \gamma}{\mu(\kappa + \mu)} - \mu_w W \quad (18.1)$$

and thus

$$W \leq \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)}. \quad (18.2)$$

These inequalities overestimate the parasite levels in the water, but they allow us to estimate these levels without solving the original system of differential equations. This will be useful in the next section.

The disease-free equilibrium satisfies

$$(\bar{S}, \bar{E}, \bar{I}, \bar{W}) = \left(\frac{\Pi}{\mu}, 0, 0, 0 \right).$$

Using the next-generation method, we have

$$\begin{aligned} \mathcal{F} &= \begin{bmatrix} \beta SW \\ 0 \\ 0 \end{bmatrix} & \mathcal{V} &= \begin{bmatrix} \alpha E + \mu E \\ -\alpha E + \kappa E + \mu I \\ -\gamma I + \mu_W W \end{bmatrix} \\ F &= \begin{bmatrix} 0 & 0 & \beta \bar{S} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} & V &= \begin{bmatrix} \alpha + \mu & 0 & 0 \\ -\alpha & \kappa + \mu & 0 \\ 0 & -\gamma & \mu_W \end{bmatrix}. \end{aligned}$$

Hence we have

$$\begin{aligned} FV^{-1} &= \begin{bmatrix} 0 & 0 & \beta \bar{S} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{\alpha + \mu} & 0 & 0 \\ \frac{\alpha}{(\alpha + \mu)(\kappa + \mu)} & \frac{1}{\kappa + \mu} & 0 \\ \frac{\alpha \gamma}{\mu_W (\alpha + \mu)(\kappa + \mu)} & \frac{\gamma}{\mu_W (\kappa + \mu)} & \frac{1}{\mu_W} \end{bmatrix} \\ &= \begin{bmatrix} \frac{\beta \bar{S} \alpha \gamma}{\mu_W (\alpha + \mu)(\kappa + \mu)} & \frac{\beta \bar{S} \gamma}{\mu_W (\kappa + \mu)} & \frac{\beta \bar{S}}{\mu_W} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}. \end{aligned}$$

Since this matrix is upper triangular, the eigenvalues are on the diagonal. So, from the properties of the next-generation process, the largest eigenvalue is thus

$$R_0 = \frac{\Pi \alpha \gamma \beta}{\mu (\alpha + \mu) (\kappa + \mu) \mu_W}. \quad (18.3)$$

Thus, if $R_0 < 1$, then the disease-free equilibrium is stable and is the only equilibrium (since $\hat{I} < 0$ in this case). If $R_0 > 1$, then the disease-free equilibrium is unstable and the endemic equilibrium exists. Note that R_0 is increasing with Π , α , γ and β , and decreasing with μ , μ_W and κ .

Education will discourage infected individuals from putting infected limbs into the drinking water. This will decrease γ . Filtration of drinking water using cloth filters will decrease β . Continuous chlorination of the water will increase μ_W . All of these interventions will result in R_0 decreasing.

However, continuous chlorination is neither possible nor desirable, so we shall assume chlorination occurs at distinct (not necessarily fixed) times t_k . At these times, the number of larvae in the water are reduced by some proportion r . This results in a system of impulsive differential equations, as discussed in Chapter 15.

18.4 The system with impulses

In this section, we use inequality (18.1) to overestimate the number of larvae in the water. This allows us to solve the corresponding impulsive differential equation, in order to derive sufficient controls.

Suppose we have maximum growth of larvae in the water, so that we have equality in (18.1). Then we have the one-dimensional impulsive differential equation

$$\begin{aligned} W' &= \frac{\alpha II\gamma}{\mu(\kappa + \mu)} - \mu_W W & t \neq t_k \\ \Delta W &= -rW & t = t_k. \end{aligned} \quad (18.4)$$

It follows that, for a single impulsive cycle $t_k \leq t \leq t_{k+1}$, the solution is

$$W(t_{k+1}^-) = W(t_k^+)e^{-\mu_W(t_{k+1}-t_k)} + \frac{\alpha II\gamma}{\mu\mu_W(\kappa + \mu)} \left[1 - e^{-\mu_W(t_{k+1}-t_k)} \right],$$

where $W(t_k^-)$ is the value immediately before the impulse and $W(t_k^+)$ is the value immediately after. For simplicity of notation, we can denote $W_k^+ = W(t_k^+)$ and $W_k^- = W(t_k^-)$. The degree of overestimation in (18.4) is shown in Figure 18.2.

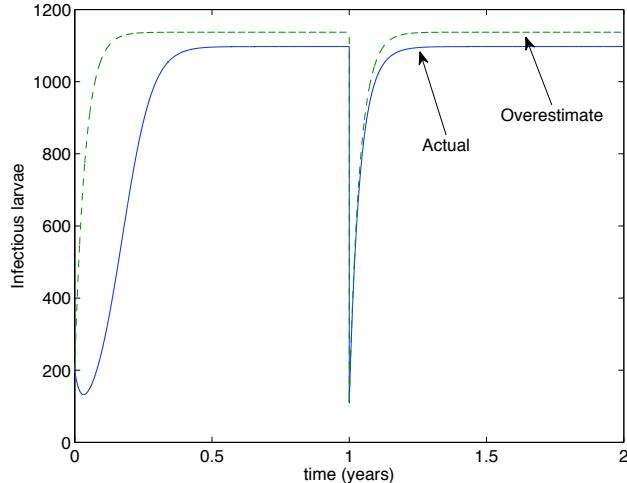


Fig. 18.2. Comparison of the actual W with the overestimate used when the growth rate is assumed to be maximal.

If we start on the endemic equilibrium, then the parasite values at the impulse times satisfy

$$\begin{aligned}
W_1^- &= \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} \\
W_1^+ &= (1-r) \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} \\
W_2^- &= (1-r) \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} e^{-\mu(t_2-t_1)} + \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} (1 - e^{-\mu_W(t_2-t_1)}) \\
W_2^+ &= (1-r)^2 \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} e^{-\mu(t_2-t_1)} \\
&\quad + (1-r) \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} (1 - e^{-\mu_W(t_2-t_1)}) \\
W_3^- &= (1-r)^2 \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} e^{-\mu_W(t_3-t_1)} + (1-r) \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} e^{-\mu_W(t_3-t_2)} \\
&\quad - (1-r) \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} e^{-\mu_W(t_3-t_1)} + \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} \\
&\quad - \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} e^{-\mu_W(t_3-t_2)}.
\end{aligned}$$

Thus, the general solution satisfies

$$\begin{aligned}
W_n^- &= \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} \left[(1-r)^{n-1} e^{-\mu_W(t_n-t_1)} + (1-r)^{n-1} e^{-\mu_W(t_n-t_2)} + \dots \right. \\
&\quad + (1-r) e^{-\mu_W(t_n-t_{n-1})} + 1 - (1-r)^{n-2} e^{-\mu_W(t_n-t_1)} \\
&\quad \left. - (1-r)^{n-3} e^{-\mu_W(t_n-t_2)} - \dots - e^{-\mu_W(t_n-t_{n-1})} \right]. \tag{18.5}
\end{aligned}$$

We have thus derived a general solution for the maximal number of parasites in the water. This occurs immediately before chlorination is applied and was derived from the overestimate (18.2). Note that this solution does not depend on the time between chlorinations being fixed.

18.4.1 Fixed chlorination

If chlorination occurs at fixed times, then $t_n - t_{n-1} = \tau$ is constant. We thus have

$$\begin{aligned}
W_n^- &= \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} \left[1 + (1-r) e^{-\mu_W \tau} + (1-r)^2 e^{-2\mu_W \tau} + \dots \right. \\
&\quad + (1-r)^{n-1} e^{-(n-1)\mu_W \tau} - e^{-\mu_W \tau} \left(1 + (1-r) e^{-\mu_W \tau} + \dots \right. \\
&\quad \left. \left. + (1-r)^{n-2} e^{-(n-2)\mu_W \tau} \right) \right] \\
&= \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} \left[\frac{1 - (1-r)^n e^{-n\mu_W \tau}}{1 - (1-r) e^{-\mu_W \tau}} - e^{-\mu_W \tau} \frac{1 - (1-r)^{n-1} e^{-(n-1)\mu_W \tau}}{1 - (1-r) e^{-\mu_W \tau}} \right]
\end{aligned}$$

Hence

$$\lim_{n \rightarrow \infty} W_n^- = \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} \left[\frac{1 - e^{-\mu_W \tau}}{1 - (1 - r)e^{-\mu_W \tau}} \right].$$

This is the long-term maximum value of the infected water (since the effect of the impulse is to immediately reduce the level of infection). To keep this below a desired threshold \tilde{W} , we thus require

$$\tau < \frac{1}{\mu_W} \ln \left[\frac{\alpha \Pi \gamma - (1 - r)\tilde{W} \mu \mu_W (\kappa + \mu)}{\alpha \Pi \gamma - \tilde{W} \mu \mu_W (\kappa + \mu)} \right] \equiv \tau_{\max}. \quad (18.6)$$

This is the maximum period between water treatments required to keep the infection below \tilde{W} .

Note that \tilde{W} must satisfy

$$\tilde{W} < \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} \quad (18.7)$$

from (18.2) in order ensure we are not taking the logarithm of a negative.

It follows that, in the case of fixed chlorination, we can derive a maximal (fixed) period of chlorination that will keep the parasite level strictly below a threshold of our choosing.

18.4.2 Nonfixed chlorination

In resource-constrained regions, regular disease control may be difficult, due to limited resources and infrastructure. In particular, chlorinating water at fixed intervals may be difficult or impossible. In order to determine the “next best” chlorination time under these circumstances using (18.5), the entire history of chlorination would need to be known. This is unlikely to be the case, so we assume that only the two most recent chlorination events are known. Specifically, we assume that

$$e^{-\mu_W (t_n - t_k)} \approx 0 \text{ for } k > 2$$

We thus have

$$W_n^- \approx \frac{\alpha \Pi \gamma}{\mu \mu_W (\kappa + \mu)} \left[(1 - r)^2 e^{-\mu_W (t_n - t_{n-2})} + (1 - r)e^{-\mu_W (t_n - t_{n-1})} + 1 - (1 - r)e^{-\mu_W (t_n - t_{n-2})} - e^{-\mu_W (t_n - t_{n-1})} \right].$$

To keep this below \tilde{W} , we thus require

$$1 - r(1 - r)e^{-\mu_W (t_n - t_{n-2})} - (2 - r)e^{-\mu_W (t_n - t_{n-1})} < \frac{\tilde{W} \mu \mu_W (\kappa + \mu)}{\alpha \Pi \gamma}. \quad (18.8)$$

Hence, if the previous two chlorination times are known, then the “next best” chlorination time satisfies

$$t_n < \frac{1}{\mu_W} \ln \left[\frac{2 - r^2}{1 - r(1 - r)e^{\mu_W t_{n-2}} - (2 - r)e^{\mu_W t_{n-1}} - \frac{\tilde{W}\mu\mu_W(\kappa+\mu)}{(\alpha\pi\gamma)}} \right].$$

It follows that, when chlorination is not fixed, we can derive the “next best” chlorination time, assuming the previous two chlorination times are known.

18.5 Numerical simulations

To examine the three crucial control parameters in more detail, we fixed all other parameters at their sample values and set $R_0 = 1$. We can solve equation (18.3) for β (which is trivial) and then use γ and μ_W as our independent variables. This allows us to create a 3D surface.

```

clear all
close all
Pi=37;
mu=1/70;
kappa=365*24;
alpha=1;
gamma=[100:100:1000 10000:10000:100000];
muW=0:0.1:200;
r=0.9;
[X,Y]=meshgrid(gamma,muW);
beta=mu.*(alpha+mu).*(kappa+mu).*Y./(Pi.*alpha.*X);
mesh(X,Y,log(beta))
xlabel('Parasite birth rate')
ylabel('Parasite death rate')
zlabel('log(Transmissibility)')
hold on
plot3(100000, 26, log(0.02555),'*r')
plot3(1000, 26, log(0.02555),'*g')

```

The resulting surface is plotted in Figure 18.3. Parameter combinations under the surface will lead to eradication, while those above will maintain disease persistence. The outcome is significantly dependent on changes in γ . Even if μ_W were increased a hundredfold, it is still unlikely to lead to eradication, while β would have to be reduced to extremely low levels because of the log scale.

The code for solving impulsive differential equations is here:

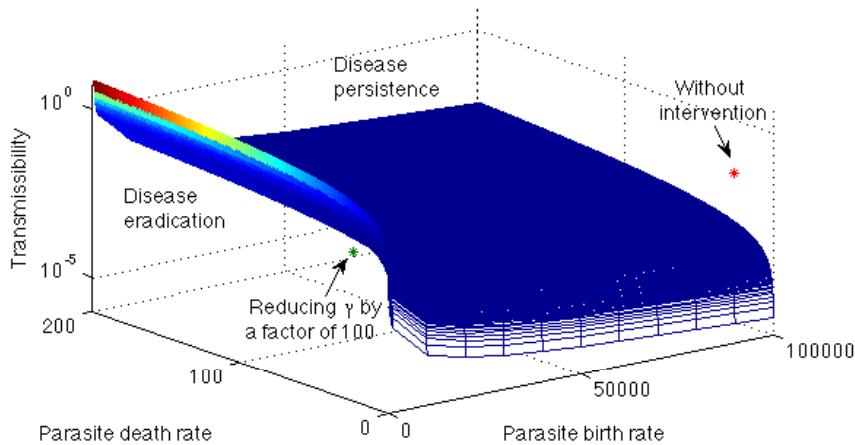


Fig. 18.3. Eradication threshold for the three parameters with the greatest influence on R_0 . Eradication will occur if the infection rate is reduced to a tiny fraction of its current value (through filtration of drinking water) or the parasite death rate is increased more than a hundredfold (through chlorination) or if the parasite birth rate is reduced to approximately a 1% of its current size (through education).

```

clear all
global Pi alpha mu kappa beta gamma muW

Pi=37;
alpha=1;
mu=0.0142;
kappa=365*24;
beta=0.0255;
gamma=100000;
muW=26;
r=0.9;
t0=0;
tq=[];
yq=[];
tfinal=12;
x0=[Pi/mu 0 0 200 200];
tq=[];
xq=[];
t0=0;

```

```

reps = 6;
tau=1;

for i=1:reps
    tspan=[t0 t0+tau];
    [t,x] = ode45(@GWDf,tspan,x0);
    n=length(x);
    x0=x(n,:);
    x0(4)=(1-r).*x(n,4);
    x0(5)=(1-r).*x(n,5);
    tf=t(n);
    tq=[tq;t(1:n)];
    xq=[xq;x((1:n),:)];
    t0=tf;
end

%plot(tq,xq(:,5))
plot(tq,xq(:,1),tq,xq(:,2),tq,xq(:,3),tq,xq(:,4))
xlabel('time (years)')
ylabel('Population')

```

```

function yp=GWDf(t,y)
global Pi alpha mu kappa beta gamma muW
yp(1,:)=Pi-beta.*y(1).*y(4)-mu.*y(1)+kappa.*y(3);
yp(2,:)=beta.*y(1).*y(4)-alpha.*y(2)-mu.*y(2);
yp(3,:)=alpha.*y(2)-kappa.*y(3)-mu.*y(3);
yp(4,:)=gamma.*y(3)-muW.*y(4);
yp(5,:)=gamma.*Pi./((mu.*(kappa+mu))-muW.*y(5));

```

The idea behind this formulation is that we run the differential equations for a while (`tau`), then reset the initial conditions in accordance with the impulse instructions, using the final conditions from the previous cycle. We do this a number of times (`reps`). The `tq` and `xq` parts are simply record-keeping: we need to track the solution every cycle. Note that we also keep track of the overestimate here in the fifth variable.

Using this code, the effect of annual chlorination is illustrated in Figure 18.4. In this case, despite significant reductions in the larval population immediately after chlorination, the population returns to high levels quite quickly. The number of susceptibles remains low, while almost all individuals remain infected.

The effect of reducing the parasite birth rate by 99% is illustrated in Figure 18.5. In this case, the number of exposed and infectious individuals approaches zero and the entire population becomes uninfected.

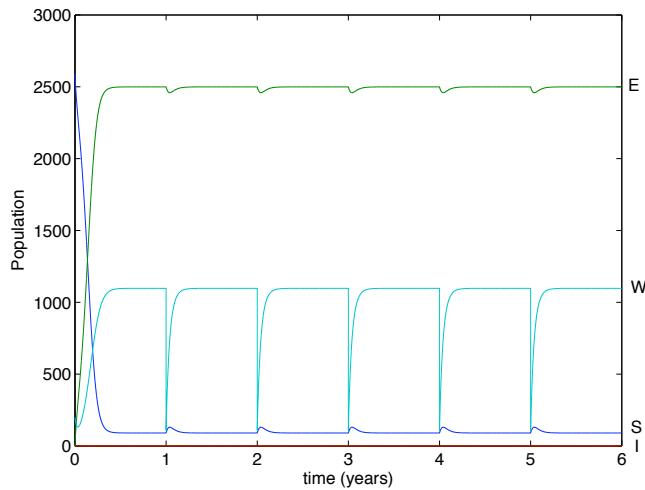


Fig. 18.4. Persistence of the disease under annual chlorination. Chlorination was assumed 90% successful and applied annually. All parameters were their sample value in Table 18.1. Note that infection levels are low, since individuals are infectious for only a brief time (the amount of time they physically submerge their foot in water). However, the burden of the disease is expressed through the exposed class, where individuals are infected with Guinea worms for months at a time.

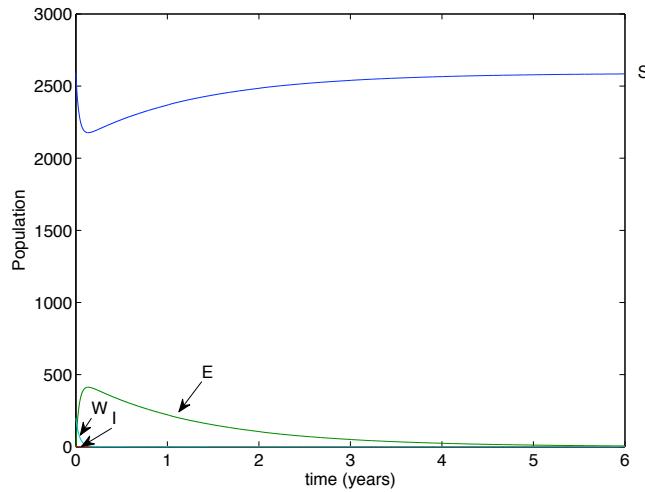


Fig. 18.5. Eradication of the disease when the parasite birth rate is decreased. Parameters used were the same as the previous scenario, except that $\gamma = 1000$ (see Figure 18.3).

Table 18.1. Parameter values. The average transmissibility β was derived from $(7 \text{ drinks of water per day}) \times (365 \text{ days}) / (100,000 \text{ larvae}) = 0.02555$. This represented the ratio of total yearly water ingested to number of parasites. The average lifespan $1/\mu$ was set to 70 years. The average infectious time $1/\kappa$ was set to be 1 hour (the length of time that an infected foot is actually submerged in the water), so that $\kappa = 24 \times 365 = 8760 \text{ years}^{-1}$. The birth rates per 1000 people in the four endemic countries are 46.09 (Mali), 43.34 (Ethiopia), 33.25 (Sudan) and 28.09 (Ghana), giving an average of 37.

Parameter	Definition	Sample value	units
S	Susceptible individuals	$S(0) = \Pi/\mu$	people
E	Exposed individuals	$E(0) = 0$	people
I	Infectious individuals	$I(0) = 0$	people
W	Water-borne larvae	$W(0) = 200$	larvae
Π	birth rate	37	people years $^{-1}$
β	transmissibility	0.0255	larvae $^{-1} \cdot$ years $^{-1}$
μ	death rate	0.0142	years $^{-1}$
κ	recovery rate	8760	years $^{-1}$
α	rate of worm emergence	1	years $^{-1}$
γ	parasite birth rate	100,000	larvae \cdot people $^{-1} \cdot$ years $^{-1}$
μ_W	parasite death rate	26	years $^{-1}$
r	chlorine effectiveness	90%	—

18.6 Discussion

So who was responsible for the near-eradication of this disease? Former President Jimmy Carter, who in 1986 set his sights on removing a disease from the planet — and he’s almost succeeded. He did this through the unglamorous but important work of mobilising public–private partnerships, delivering education messages to remote populations and even negotiating a “Guinea-worm ceasefire” in the Sudan civil war so that NGOs could go in and educate those most at risk.

We stand at the brink of eradicating one of humanity’s oldest scourges. There are three criteria for the eradication of an infectious disease: 1. biological and technical feasibility; 2. costs and benefits; and 3. societal and political considerations. GWD satisfies all three. While eradication efforts have been immensely successful thus far, the final phase of eradication will occur in resource-poor and underfunded areas of the world. Knowing which strategies may be optimal will be of enormous benefit.

Smallpox remains the only disease we have completely eradicated, despite eradication hopes for malaria, yaws and yellow fever in the twentieth century and current eradication programs, such as polio and leprosy. Measles, rubella and hepatitis A and B are biologically and technically feasible candidates for eradication.

A critical tool for smallpox eradication, in addition to an extremely effective vaccine, was photographic disease-recognition cards, demonstrating that

non-biomedical interventions were also important. Barriers to smallpox eradication included cultural traditions, a lack of societal support and religious beliefs. Despite strong biological, technical and cost-benefit arguments for eradication of many infectious diseases, securing societal and political commitment has been recognised as a substantial challenge.

The most effective way to eradicate GWD is to reduce the parasite birth rate. This can be achieved via education; specifically, teaching people not to put infected limbs into the drinking water. Although behaviour changes are, in general, notoriously difficult, GWD eradication programs have had significant success in altering people's behaviour. If 99% of people can be persuaded not to put their infected feet in the drinking water, then eradication is assured. While chlorination can theoretically control the disease and we have provided estimates for the necessary frequency and strength of chlorination, numerical simulations demonstrate that education is far more effective. Thus, our results here are not advocating for something untested but rather point to the importance that one of the three existing intervention methods — namely, persuading people not to put infected limbs in the drinking water — will have in the final push towards complete eradication.

The final steps towards eradication of GWD should take place within the next few years. Our modelling shows that education is the most effective intervention method, but a combination of education, chlorination and filtration will likely be required to achieve the final steps in the long journey to eradication. By mustering both scientific and cultural resources, we can successfully defeat one of the oldest diseases in human history.

18.7 Lab work

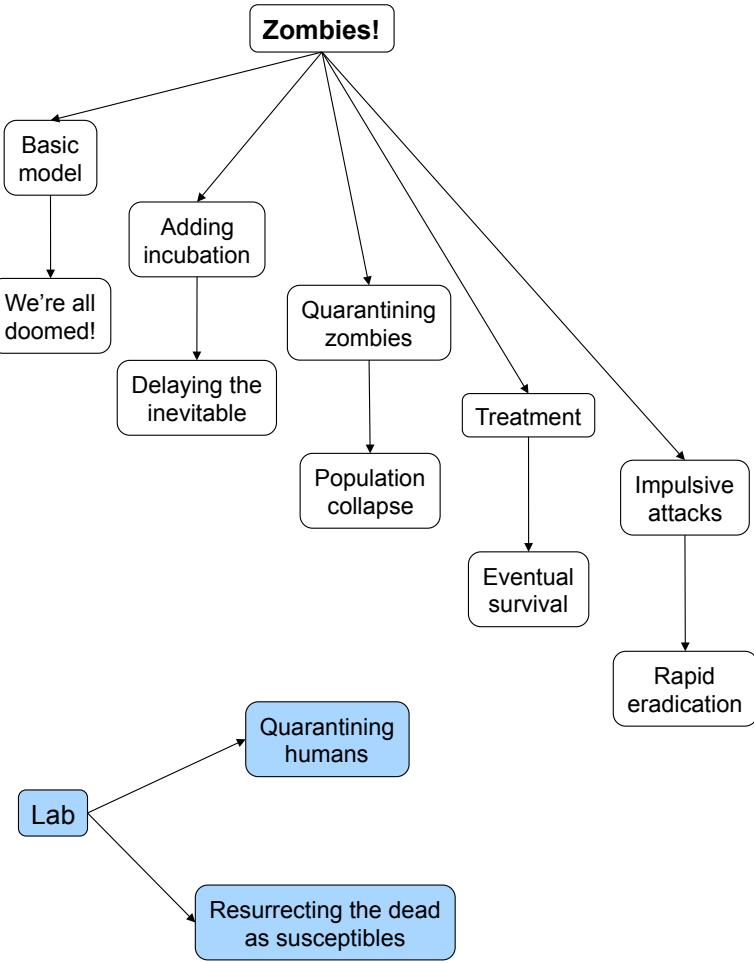
1. Show that the endemic equilibrium of the system without impulses is given by

$$\begin{aligned}\hat{S} &= \frac{\mu_W}{\beta\gamma} \left(\kappa + \mu + \frac{\kappa\mu}{\alpha} + \frac{\mu^2}{\alpha} \right) \\ \hat{E} &= \frac{\kappa + \mu}{\alpha} \hat{I} \\ \hat{W} &= \frac{\gamma}{\mu_W} \hat{I} \\ \hat{I} &= \frac{\Pi\beta\gamma\alpha - \mu\mu_W(\kappa\alpha + \mu\alpha + \kappa\mu + \mu^2)}{(\alpha + \kappa + \mu)\beta\gamma\mu}.\end{aligned}$$

2. Instead of (18.1), suppose we made a simpler estimate: namely, that $I \leq \frac{\Pi}{\mu}$. That is, the number of infected individuals is less than the total number of humans. In this case, our overestimate would be

$$\begin{aligned}W' &= \frac{\gamma\Pi}{\mu} - \mu_W W & t \neq t_k \\ \Delta W &= -rW & t = t_k.\end{aligned}\tag{18.9}$$

- a) Solve equation (18.9) for fixed chlorination.
- b) Plot this overestimate and the actual solution on the same graph using the parameters in Table 18.1. What do you think of this overestimate?
- c) Find the maximal period for the parameters in Table 18.1 and a parasite threshold of $\tilde{W} = 1000$.
- d) Solve (18.6) for the same parameters. Which is more reasonable?
(Note: you might want to change the scale once you've found the answer to help with interpretation.)



Application: Zombies! (Aargh!)

19.1 Introduction — of peril!

A zombie is a reanimated human corpse that feeds on living human flesh. They are mindless monsters who do not feel pain and who have an immense appetite for human flesh, making them one of our natural predators. Their aim is to kill, eat or infect us. The “undead” move in small, irregular steps and show signs of physical decomposition, such as rotting flesh, discoloured eyes and open wounds. Modern zombies are often related to an apocalypse, where civilisation could collapse due to a plague of the undead.

Major outbreaks of zombies have been recorded since 1968, primarily in the US and the UK. Various historical records have described zombies overwhelming such important places as isolated farmhouses, shopping malls and British pubs. When a susceptible individual is bitten by a zombie, it leaves an open wound. The wound created by the zombie has the zombie’s saliva in and around it. This bodily fluid mixes with the blood, thus infecting the (previously susceptible) individual.

The origins of the zombie outbreak are particularly murky. The historical records that we have access to have speculated that the causes might include radiation (*Night of the Living Dead*), exposure to airborne viruses (*Resident Evil*) or stings from genetically altered bees (*Dead Rising*). However, such theories remain untested, as the individuals in question were soon overwhelmed by activities such as running in terror, fighting hordes of the undead and, shortly afterwards, attempting to eat their friends.

Instead, focus has largely been on methods to defeat the zombies. Successful tools at our disposal include guns, the army, eventual starvation — and, of course, Dire Straits records. However, there is one tool in the zombie arsenal that has not been utilised: mathematics!

One of the confounding factors in the zombie apocalypse is that previously living friends and relatives will return to life in the form of the living dead, who must then be disposed of. So if you feel comfortable shooting grandma in the

face, then congratulations: a) you stand a chance of surviving the apocalypse and b) the world will hence consist entirely of sociopaths.

The biggest problem with the undead is their sheer numbers. They don't need to eat, sleep, go to work or stand in line at the bank (although many will, if undisturbed). So while one zombie may be defeatable, good luck facing a thousand of them with a limited supply of bullets and no knowledge of how to forge metal for replacements because Google is long gone at this point. Yes, that's how serious this is.

Happily, we have an advantage that the zombies don't: our braaaiiinnnss. The very thing the zombies want to eat is also our greatest weapon. Coincidence? You make the call. We can do all sorts of things the zombies can't: build moats, electrify fences, operate vehicles with an internal combustion engine and draw cartoons. The latter may not be so useful in the apocalyptic collapse of civilisation as we know it, but you never know.

In order to model a zombie outbreak, it is important to observe the action of zombie infection, preferably from some considerable distance. Humans can be infected by direct contact with a zombie. However, humans may also die of natural causes, whereupon they can be resurrected as a zombie. zombies can also be killed in an encounter with humans, possibly permanently.

19.2 The model — of doom!

For the basic model, we consider three classes: Susceptible (S), Zombies (Z) and Removed (R). The removed class represents the temporarily deceased, whether through natural death surrounded by weeping loved ones or through the (temporary) death of a zombie. This event may also be accompanied by weeping, although for entirely different reasons. Zombies may also be killed permanently, upon severance of the brain stem through methods too gory to describe here. However, this is not as easy as most people imagine (and don't even ask about how we know this for a fact), so only a percentage of zombie deaths will be permanent.

Given the sheer numbers of the undead, we will model infection through mass action (fortunately, sexual transmission of the zombie virus has only been documented in a very few cases, about which we do not want to know or speculate further). We assume a constant birth rate (Π) and a linear death rate (δ), the latter of which does not remove individuals from the system of course. The transmission rate is β and the reanimation rate is ζ . When humans kill zombies, we assume that some proportion p are permanently killed, but the remainder move to the Removed class, where they can subsequently be reanimated. However, the rate of effectively killing zombies is likely to be inefficient, so we shall make the assumption that $\alpha p < \beta$. That is, we cannot kill zombies faster than they can infect us.

The basic SZR model is thus given by

$$\begin{aligned} S' &= \Pi - \beta SZ - \delta S \\ Z' &= \beta SZ + \zeta R - \alpha SZ \\ R' &= \delta S + \alpha(1-p)SZ - \zeta R. \end{aligned}$$

See Figure 19.1.

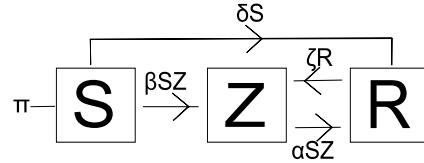


Fig. 19.1. The basic zombie model.

19.3 Analysis — of anguish!

Since we are usually dealing with only a short timescale (on the matter of days or weeks rather than decades), the contributions of human birth and natural death will be negligible. So we will assume $\Pi = \delta = 0$.

In this case, the equilibria satisfy

$$\begin{aligned} -\beta SZ &= 0 \\ \beta SZ + \zeta R - \alpha SZ &= 0 \\ \alpha(1-p)SZ - \zeta R &= 0. \end{aligned}$$

From the first equation, either $S = 0$ or $Z = 0$. When $Z = 0$, we have the DFE

$$(\bar{S}, \bar{Z}, \bar{R}) = (N, 0, 0),$$

where N is the entire population. It follows from $S = 0$ that we get the “doomsday” equilibrium

$$(\bar{S}, \bar{Z}, \bar{R}) = (0, N, 0),$$

with every human converted and every dead individual reanimated. These equilibrium points suggest that human–zombie coexistence is impossible. You should probably cancel any well-intentioned dinner parties, unless you want to be on the menu.

The Jacobian is

$$J = \begin{bmatrix} -\beta Z & -\beta S & 0 \\ \beta Z - \alpha Z & \beta S - \alpha S & \zeta \\ \alpha(1-p)Z & \alpha(1-p)S & -\zeta \end{bmatrix}.$$

The Jacobian at the DFE is

$$J(N, 0, 0) = \begin{bmatrix} 0 & -\beta N & 0 \\ 0 & \beta N - \alpha N & \zeta \\ 0 & \alpha(1-p)N & -\zeta \end{bmatrix}.$$

We have $\det(J - \lambda I) = -\lambda [\lambda^2 + \lambda(\alpha N - \beta N + \zeta) + \zeta N(\alpha p - \beta)]$. With the assumption that $\alpha p - \beta < 0$, it follows that the characteristic equation always has a root with positive real part. It follows that the DFE is always unstable, meaning it only takes a small perturbation (say, a zombie or two) and the infection will spread. This is not good.

Next we have

$$J(0, N, 0) = \begin{bmatrix} -\beta N & 0 & 0 \\ \beta N - \alpha N & 0 & \zeta \\ \alpha(1-p)N & 0 & -\zeta \end{bmatrix}.$$

In this case, the characteristic equation is $\det(J - \lambda I) = -\lambda(-\beta N - \lambda)(-\zeta - \lambda)$. Since all eigenvalues of the doomsday equilibrium are negative, it is asymptotically stable. Hence, in a short outbreak, zombies will likely infect everyone.

In Figure 19.2, we simulated a mid-size city of a million people, starting with a single zombie on Day 0. Parameters used were $\beta = 5 \times 10^{-6}$ per people per day (representing 5 initial bites for each zombie in a population of 1 million), $\zeta = 1$ per day (representing a 24 hour period before reanimation), $\alpha = \beta/10$ (representing a kill-to-bite ratio of 10%) and $p = 0.9$ (representing 90% effectiveness at permanently killing zombies). The result is that the human race is effectively eradicated after four days.

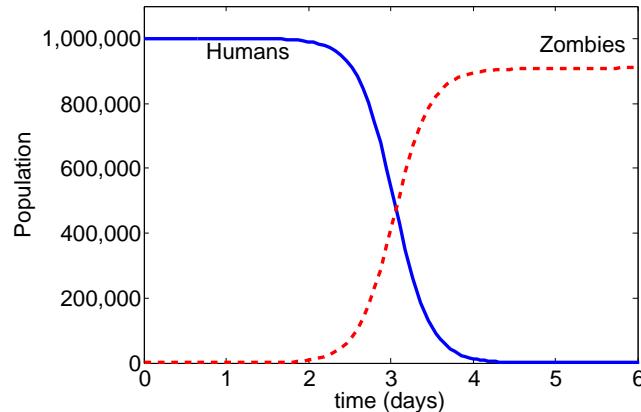


Fig. 19.2. Basic model outbreak scenario. Susceptibles are quickly eradicated and zombies take over, infecting everyone.

Since ODEs model populations, not individuals, this does not preclude a few survivors holed up in a shopping mall. However, these individuals represent a negligible fraction of the population. Furthermore, a) it's not looking too good outside and b) they can survive for precisely the length of time it takes until the stupidest member of their group opens the door. So not terribly long, then.

The code for the basic zombie model is given by:

```

clear all
x0=[1e6-1 1 0];
t0=0;
tau=7;
tspan=[t0 t0+tau];
[t,x]=ode45(@zombiesf,tspan,x0);
plot(t,x(:,1),'b',t,x(:,2),'r--')
xlabel('time (days)')
ylabel('Population')

```

```

function yprime=zombiesf(t,y)
beta=5./1e6;
zeta=1;
alpha=beta./10;
p=0.9;
yprime(1,:)=-beta.*y(1).*y(2);
yprime(2,:)=beta.*y(1).*y(2)+zeta.*y(3)-alpha.*y(1).*y(2);
yprime(3,:)=alpha.*((1-p).*y(1).*y(2)-zeta.*y(3));

```

19.4 Incubation — of destiny!

We now revise the model to include more realism, which is very important when dealing with zombies. As we all know, there is a period of time after the susceptible human gets bitten before they succumb to their wounds and become a zombie. During this time, any infected human must on no account reveal their status to their friends and instead should sweat, shiver uncontrollably and generally act suspiciously, while their friends remain hilariously unaware of the forthcoming transformation.

Susceptibles will first move to an infected class (I) and remain there for some time until the infection activates and they transform into a zombie. Infected individuals can still die a “natural” death before activation.

The revised model is thus

$$\begin{aligned}
S' &= \Pi - \beta SZ - \delta S \\
I' &= \beta SZ - \rho I - \delta I \\
Z' &= \rho I + \zeta R - \alpha SZ \\
R' &= \delta S + \alpha(1-p)SZ - \zeta R.
\end{aligned}$$

See Figure 19.3.

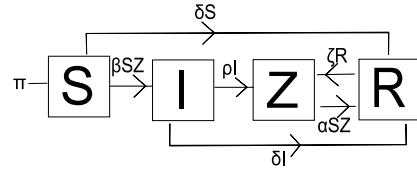


Fig. 19.3. The basic model with latent infection.

As before, we shall assume a short timescale and take $\Pi = \delta = 0$. In this case, we have two equilibria:

$$(\bar{S}, \bar{I}, \bar{Z}, \bar{R}) = (N, 0, 0, 0) \quad \text{and} \quad (\bar{S}, \bar{I}, \bar{Z}, \bar{R}) = (0, 0, N, 0).$$

For the DFE, we can use the next-generation method to find the reproduction number. We have

$$\begin{aligned}
\mathcal{F} &= \begin{bmatrix} \beta SZ \\ 0 \end{bmatrix} & \mathcal{V} &= \begin{bmatrix} \rho I \\ -\rho I - \zeta R + \alpha SZ \end{bmatrix} \\
F &= \begin{bmatrix} 0 & \beta S \\ 0 & 0 \end{bmatrix} & V &= \begin{bmatrix} \rho & 0 \\ -\rho & \alpha S \end{bmatrix} \\
FV^{-1} &= \begin{bmatrix} \beta/\alpha & \beta/\alpha \\ 0 & 0 \end{bmatrix}
\end{aligned}$$

Hence the reproduction number is $R_0 = \beta/\alpha$, which is greater than 1. It follows that the DFE is always unstable.

The Jacobian matrix is

$$J = \begin{bmatrix} -\beta Z & 0 & -\beta S & 0 \\ \beta Z & -\rho & \beta S & 0 \\ -\alpha Z & \rho & -\alpha S & \zeta \\ \alpha(1-p)Z & 0 & \alpha(1-p)S & -\zeta \end{bmatrix}.$$

At the doomsday equilibrium, we have

$$J(0, 0, N, 0) = \begin{bmatrix} -\beta N & 0 & 0 & 0 \\ \beta N & -\rho & 0 & 0 \\ -\alpha N & \rho & 0 & \zeta \\ \alpha(1-p)N & 0 & 0 & -\zeta \end{bmatrix}.$$

The eigenvalues of this matrix are $-\beta N$, $-\rho$, 0 and $-\zeta$. Since we have a zero eigenvalue, we cannot determine the stability. We cannot apply the next-generation method either, since the matrix V is not invertible if $S = 0$.

What happens if this equilibrium is unstable? Since the DFE is also unstable, solutions must persist. If both equilibria are unstable, then there must be some stable object, such as a periodic orbit or chaos, to which solutions converge. For our purposes, however, what we really care about is whether the zombies exist or not, so the stability of the DFE is enough to tell us that we're in trouble. Whether the zombie population oscillates or settles down into an equilibrium is fairly inconsequential.

Using the same parameters as before and adding $\rho = 2$ (so that it takes half a day for the infection to activate), we have the following function file:

```
function yprime=latentzombiesf(t,y)
beta=5./1e6;
zeta=1;
alpha=beta./10;
p=0.9;
rho=2;
yprime(1,:)=-beta.*y(1).*y(3);
yprime(2,:)=beta.*y(1).*y(3)-rho.*y(2);
yprime(3,:)=rho.*y(2)+zeta.*y(4)-alpha.*y(1).*y(3);
yprime(4,:)=alpha.*(1-p).*y(1).*y(3)-zeta.*y(4);
```

The only change to the script file is that the initial conditions need to have four components, not three. So with $I(0) = 0$, we get the situation shown in Figure 19.4.

Zombies still take over the population, but adding a 12-hour delay to each infection has bought us an additional four days. We're still doomed, but it gives us a bit more time to say goodbye to any surviving loved ones and perhaps catch a few last episodes of our favourite TV shows, assuming there's still electricity, television and, indeed, actors.

19.5 Quarantine — of terror!

It's time to strike back. Of course, we've already been fighting the zombies at rate α , but we can be more strategic. This is where our braaaiiinnnss come in. So let's quarantine at least some of the zombies. The great thing about quarantine, mathematically speaking, is that it removes individuals from the infectious pool. So any zombies who are in quarantine cannot infect new humans while they remain locked up.

Of course, we'd probably want to quarantine infected humans as well, assuming we can find them. Your average zombie might be a hideous shambling

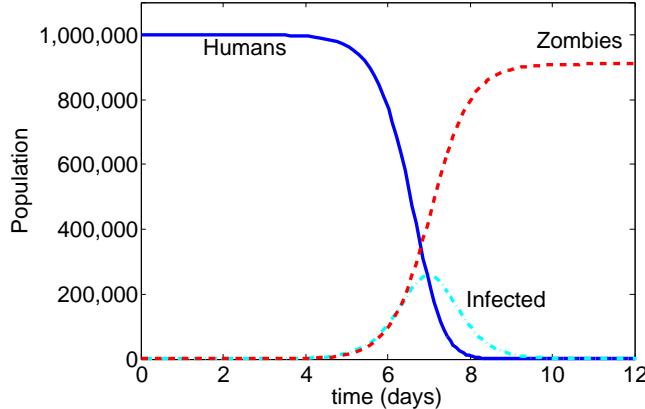


Fig. 19.4. An outbreak with latent infection.

creature of doom, but at least we can recognise one when we see it. Finding zombies to quarantine isn't a problem; of course, we run the risk of them trying to convert us when we do find them. Infected individuals present the reverse problem: they're relatively easy to capture, but we usually don't know what they look like. (And who can blame them for wanting to hide their symptoms if they're about to be tossed into a cage full of zombies?)

So we'll quarantine both the infected and also the zombies but at different rates (κ and σ , respectively). What happens after quarantine? We could assume nothing happens, but a) that's not very interesting and b) mathematically, we need to deal with this. Some of the quarantined zombies or infected humans might try to escape, but any that tried to would be killed (at rate γ) before finding their "freedom". If we post snipers around the edge, then there's a good chance that the snipers could kill escaping individuals, but not a great chance that any such death would be permanent. So we'll assume a proportion q are permanently killed, with the possibility that q may be zero. The rest will enter the removed class and may later become reanimated as "free zombies".

The model with quarantine is thus

$$\begin{aligned}
 S' &= \Pi - \beta SZ - \delta S \\
 I' &= \beta SZ - \rho I - \delta I - \kappa I \\
 Z' &= \rho I + \zeta R - \alpha SZ - \sigma Z \\
 R' &= \delta S + \alpha(1-p)SZ - \zeta R + \gamma(1-q)Q \\
 Q' &= \kappa I + \sigma Z - \gamma Q
 \end{aligned}$$

See Figure 19.5.

For a short outbreak, we have the DFE as usual:

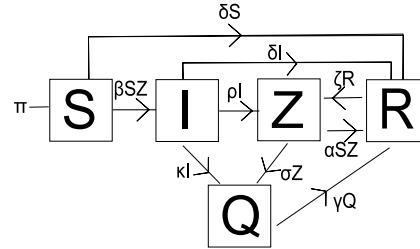


Fig. 19.5. The model with quarantine.

$$(\bar{S}, \bar{I}, \bar{Z}, \bar{R}, \bar{Q}) = (N, 0, 0, 0, 0).$$

There is also a coexistence equilibrium

$$(\bar{S}, \bar{I}, \bar{Z}, \bar{R}, \bar{Q}) = \left(0, 0, \bar{Z}, \frac{\sigma \bar{Z}}{\zeta}, \frac{\sigma \bar{Z}}{\gamma} \right)$$

that only exists if $q = 0$. If $q \neq 0$, then the second equilibrium is $(0, 0, 0, 0, 0)$.

Using the next-generation method to analyse the DFE, we consider the infective classes I and Z . We have

$$F = \begin{bmatrix} 0 & \beta N \\ 0 & 0 \end{bmatrix} \quad V = \begin{bmatrix} \rho + \kappa & 0 \\ -\rho & \alpha N + \sigma \end{bmatrix}$$

$$FV^{-1} = \begin{bmatrix} \frac{\rho \beta N}{(\rho + \kappa)(\alpha N + \sigma)} & \frac{\beta N}{\alpha N + \sigma} \\ 0 & 0 \end{bmatrix}.$$

This gives us

$$R_0 = \frac{\rho \beta N}{(\rho + \kappa)(\alpha N + \sigma)}.$$

The DFE is only stable if $R_0 < 1$. We have no control over β , N or ρ . Assuming that α is already maximised (i.e., we are already killing zombies as best we can), the two parameters we can control are κ and γ , the rates of quarantining infected and zombified individuals, respectively. If the population is large, then

$$R_0 \approx \frac{\rho \beta}{(\rho + \kappa)\alpha}.$$

The only way to keep $R_0 < 1$ and prevent an outbreak would be to increase κ . The result thus depends critically on finding the very individuals who do not want to be found. This is a problem.

Furthermore, quarantining a large percentage of infected individuals is unrealistic, due to infrastructure limitations, so we do not expect κ or σ to be large. We thus expect $R_0 > 1$ in practice. Hence zombies can invade, even with quarantine. Figure 19.4 illustrates the situation for parameters as before, with additionally $\kappa = 0.01$, $\sigma = 0.1$, $q = 0.1$ and $\gamma = 1/7$. In this case, the zombies are in fact eventually controlled, but honestly it's a bit of a pyrrhic victory.

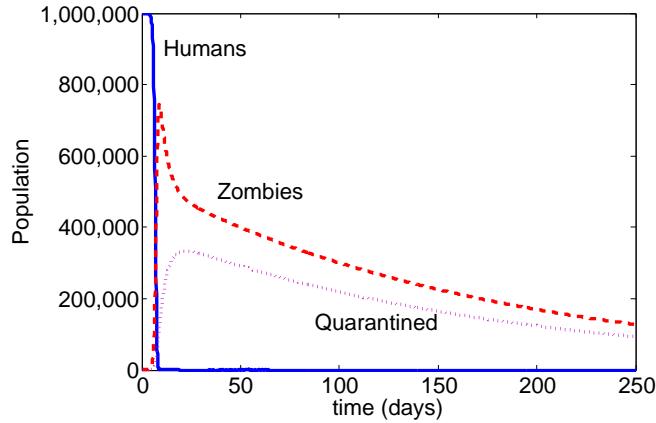


Fig. 19.6. With quarantine, the entire population collapses, and we all lose.

19.6 A cure — of fear!

Suppose we can produce a treatment for zombie-ism that returns the undead to life. Okay, stop laughing. Sure, it's a bit ridiculous, but stay with me here. The only thing we need to assume is that the treatment does not provide immunity, meaning that the zombies will become susceptible humans.

Since quarantine didn't work, let's drop that. And hey, who cares about quarantine because *we have a cure!* See Figure 19.7.

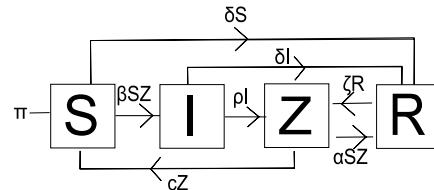


Fig. 19.7. The model with treatment.

The model is given by

$$\begin{aligned}
 S' &= \Pi - \beta SZ - \delta S + cZ \\
 I' &= \beta SZ - \rho I - \delta I \\
 Z' &= \rho I + \zeta R - \alpha SZ - cZ \\
 R' &= \delta S + \alpha(1-p)SZ - \zeta R.
 \end{aligned}$$

For a short timescale, there are two equilibria, the DFE

$$(\bar{S}, \bar{I}, \bar{Z}, \bar{R}) = (N, 0, 0, 0)$$

and a survivors' equilibrium

$$(\bar{S}, \bar{I}, \bar{Z}, \bar{R}) = \left(\frac{c}{\beta}, 0, 0, 0 \right).$$

Using the next-generation method, we find

$$R_0 = \frac{\beta \bar{S}}{\alpha \bar{S} + c}.$$

For $\bar{S} = c/\beta$, we have $R_0 = \frac{\beta}{\alpha + \beta} < 1$ always. So the survivors' equilibrium is always stable.

For $\bar{S} = N$, we have $R_0 = \frac{\beta N}{\alpha N + c}$. The critical threshold for vaccination is thus $c^* = (\beta - \alpha)N$. If $c < c^*$, then the DFE is unstable. If $c > c^*$, the DFE is stable. In this case we have bistability, so the results depend on the initial conditions. However, for our parameters, we find $c^* = 4.5$ per day. That is, the average zombie must be infectious for a maximum $1/4.5=0.222$ days before being cured. This means we need to cure zombies within 5 hours and 20 minutes of their activation. Good luck with that. You're going to need it.

For a more realistic value, say curing zombies within 5 days of activation, we have $c = 0.2$ and hence there are 40,000 survivors. This is better than nothing, but it's a lot smaller than the 1,000,000 people we started with. The results are shown in Figure 19.8.

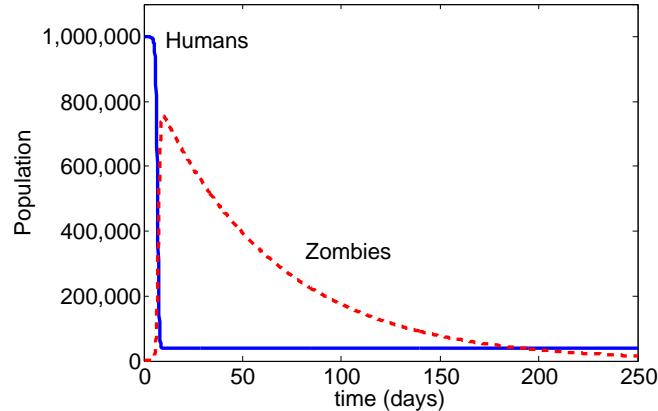


Fig. 19.8. With treatment, humans eventually survive, but it takes a long time.

A cure would be nice, but a) we can't sit around waiting for someone to invent one when there are zombies at the door, and b) do you really want to be part of that tiny portion of survivors fighting and curing zombies aggressively for the better part of a year? Well, it's better than the alternative, I suppose.

19.7 Using our braaaaiinnnsss! (impulsively)

Let's seize the major advantage we have over zombies: no, not our fashion sense, our intellect. Suppose we try to strategically destroy zombies at such times that our resources permit. This may be through targeted attacks by government forces or coordinating among survivors. Each time we attack the zombies, we can learn lessons, plan better for next time and then strike back harder, destroying more zombies with each attack. So it's a careful, thoughtful balance between education and mindless violence. A bit like high school, then.

This results in an impulsive effect. At each attack, the zombies are reduced. There's a baseline level of zombie-killing competence, the kill ratio $k \in (0, 1]$. We'll hit the zombies n times, where n is the number of attacks we need until they're eradicated. That is, we'll hit the zombies at distinct times, with ever-increasing force.

We'll add this effect to the basic model. Yes, the very worst model of all, the one that gave us only four days to survive.

$$\begin{aligned} S' &= \Pi - \beta SZ - \delta S & t \neq t_n \\ Z' &= \beta SZ + \zeta R - \alpha SZ & t \neq t_n \\ R' &= \delta S + \alpha(1-p)SZ - \zeta R & t \neq t_n \\ \Delta Z &= -knZ & t = t_n \end{aligned}$$

We used a kill ratio of 20% and applied the impulsive effect every 24 hours. The results are illustrated in Figure 19.9.

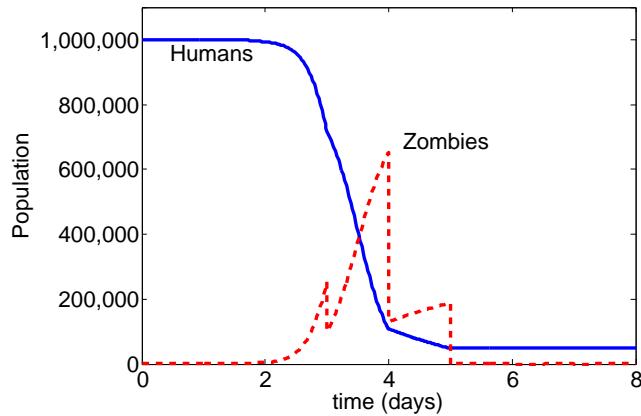


Fig. 19.9. Impulsive attacks can control the epidemic.

The impulsive attacks were carried out once per day. For the first two days, the effects are not noticeable. However, by the third day, the attack

strength has increased (due to lessons learned) and is considerable by Day 4. The zombies are finally eradicated on Day 5, leaving a small but non-negligible population of surviving humans. Note that the sharp turns in the human populations are a result of the impulsive effect; the solutions here are continuous, but their derivatives are not, producing non-smooth corners at the impulsive times.

The code for the impulsive attacks is given as follows:

```

clear all
x0=[1e6-1 1 0];
t0=0;
tq=[];
xq=[];
k=0.2;
tau=1;
for i=1:8
    tspan=[t0 t0+tau];
    [t,x]=ode45(@zombiesf,tspan,x0);
    xq=[xq;x];
    tq=[tq;t];
    t0=t(length(t));
    x0=x(length(t),:);
    x0(2)=(1-k.*i).*x0(2);
end
plot(tq,xq(:,1),tq,xq(:,2),'r--')
xlabel('time (days)')
ylabel('Population')

```

There's one further tweak we have to do, however: we need to ensure that solutions do not become negative (since we're subtracting potentially more zombies than there actually are). In the function file, we solve that problem by adding this:

```

if y(2)<0
    y(2)=0;
end

```

19.8 Discussion — of dread!

An outbreak of zombies infecting humans is likely to be disastrous, unless extremely aggressive tactics are employed against the undead. Without intervention, the collapse of civilisation would happen within days. The key

difference between the models presented here and other models of infectious disease is that the dead can come back to life.

Quarantine has the potential to contain the epidemic, but sufficient numbers of infected individuals need to be located or else the likely outcome is the catastrophic elimination of the entire population. Treatment can result in some humans surviving, but it relies upon the development of a cure either in advance or within the first few days of the epidemic. Only sufficiently frequent attacks, with increasing force, will result in eradication, assuming the available resources can be mustered in time.

Furthermore, these results assumed that the timescale of the outbreak was short, so that the natural birth and death rates could be ignored. If the timescale of the outbreak increases, then the result is the doomsday scenario: an outbreak of zombies will result in the collapse of civilisation, with every human infected or dead. This is because human births and deaths will provide the undead with a limitless supply of new bodies to infect, resurrect and convert. Thus, if zombies arrive, we must act quickly and decisively to eradicate them before they eradicate us.

In summary, a zombie outbreak is likely to lead to the collapse of civilisation, unless it is dealt with quickly. The most effective way to contain the rise of the undead is to hit hard and hit often. It is imperative that zombies are dealt with quickly, or else we are all in a great deal of trouble.

19.9 Lab work — of horror!

One of the lessons we saw here was going around the modelling cycle (see Figure 2.1 on Page 11) several times. In the lab, we'll do that some more and examine some potential scenarios.

1. One possibility might be to quarantine the humans instead of the zombies. In this case, these humans will be immune to infection while they remain quarantined. When they are released from quarantine, they will be susceptible humans again.
 - a) Draw the model diagram.
 - b) Write down the differential equations.
 - c) For a short-term outbreak, find the DFE and determine its stability.
 - d) Under what circumstances will there be a second equilibrium? Determine its stability when it exists.
 - e) Choose some likely parameter values. Include units for each parameter.
 - f) Plot the time series.
 - g) Determine the long-term outcome.
 - h) Describe any interesting or unexpected features of this disease not included in the above list.

2. What if the dead returned to life as susceptible individuals? That is, zombies will still infect the living, but when the dead rise they rejoin the human race, at least until bitten again. Perform the same steps as in the previous question for this option.

Part IV

Appendices

A

Solving the time series directly

Let's solve the system

$$\frac{dS}{dt} = bI - aSI \quad (\text{A.1})$$

$$\frac{dI}{dt} = aSI - bI \quad (\text{A.2})$$

directly. First, note that the two right-hand sides are negatives of each other. So if we add the equations together, we have

$$\begin{aligned} S' + I' &= 0 \\ S + I &= N, \text{ a constant.} \end{aligned}$$

That is, we've *deduced* that the entire population is constant. (Note that this isn't true for most models, but we're dealing with a very simple one here.) That is, everyone in the town is either susceptible or infected, and no one enters, leaves, is born or dies. We could thus substitute $S = N - I$ into equation (A.2) to get

$$\begin{aligned} \frac{dI}{dt} &= a(N - I)I - bI \\ &= (aN - b - aI)I. \end{aligned}$$

Since $aN - b$ is just some constant number, we can represent it by a constant A . Thus

$$\frac{dI}{dt} = (A - aI)I.$$

In order to solve this, we are going to have to put all the " I "s (including dI) on one side and all the " t "s on the other (although all we have is dt in this case). Thus, we can rearrange to get

$$\frac{dI}{(A - aI)I} = dt. \quad (\text{A.3})$$

To solve this, we integrate both sides and use partial fractions (see Appendix B) to get

$$\begin{aligned} \int \frac{1}{A} \left(\frac{a}{A - aI} + \frac{1}{I} \right) dI &= \int dt \\ \frac{1}{A} [-\ln(A - aI) + \ln I] &= t + C \\ \frac{1}{A} \ln \left(\frac{I}{A - aI} \right) &= t + C \quad \text{since } \ln u - \ln v = \ln \frac{u}{v} \\ \ln \left(\frac{I}{A - aI} \right) &= At + AC \\ \frac{I}{A - aI} &= e^{At + AC} \quad \text{since } e \text{ and } \ln \text{ are inverses} \\ &= e^{At} e^{AC} \quad \text{since } e^{x+y} = e^x e^y \\ &= Be^{At} \quad (B = e^{AC} \text{ is just a constant}) \\ I &= (A - aI)Be^{At} \\ (1 + aBe^{At})I &= ABe^{At} \\ I &= \frac{ABe^{At}}{1 + aBe^{At}}. \end{aligned}$$

We now solve for B (our integration constant). When $t = 0$, $I = I_0$. Substituting this in and remembering that $A = aN - b$, we have

$$\begin{aligned} I_0 &= \frac{AB(1)}{1 + aB(1)} \quad \text{since } e^0 = 1 \\ I_0 &= \frac{(aN - b)B}{1 + aB} \\ (1 + aB)I_0 &= (aN - b)B \\ B[(aN - b) - aI_0] &= I_0 \\ B &= \frac{I_0}{(aN - b) - aI_0} \end{aligned}$$

Substituting and rearranging to simplify our equation somewhat, we get

$$\begin{aligned} I &= \frac{\frac{(aN - b)I_0 e^{(aN - b)t}}{(aN - b) - aI_0}}{1 + \frac{aI_0 e^{(aN - b)t}}{(aN - b) - aI_0}} \\ &= \left[\frac{(aN - b)I_0 e^{(aN - b)t}}{(aN - b) - aI_0} \right] \left[\frac{(aN - b) - aI_0}{(aN - b) + aI_0 [e^{(aN - b)t} - 1]} \right] \\ &= \frac{(aN - b)I_0 e^{(aN - b)t}}{(aN - b) + aI_0 [e^{(aN - b)t} - 1]}. \end{aligned}$$

Lastly, since $S = N - I$, we have

$$S = N - \frac{(aN - b)I_0 e^{(aN-b)t}}{(aN - b) + aI_0[e^{(aN-b)t} - 1]}.$$

We've thus (finally!) derived the time-series equations. That is, we have explicitly found $S(t)$ and $I(t)$ for every set of initial conditions, depending on parameters inherent to the model. These tell us how the system changes over time.

In this case, our system was very simple. In fact, it was degenerate, which is what allowed us to turn the two-dimensional system of differential equations into a single differential equation (using the fact that N was constant, which isn't usually true). Most ODE models aren't so simple and aren't solvable.

B

Partial fractions

In order to solve

$$\frac{dI}{(A - aI)I} = dt,$$

we need to split the fraction $\frac{1}{(A - aI)I}$ into simpler fractions (because we don't know how to integrate the left-hand side as it stands).

These fractions will be some combination of $\frac{1}{A - aI}$ and $\frac{1}{I}$. We don't yet know what kind of combination, so let's label the unknowns by G and H :

$$\frac{G}{A - aI} + \frac{H}{I} = \frac{1}{(A - aI)I}.$$

We're trying to solve for G and H so let's multiply everything by the common denominator (which is $(A - aI)(I)$). This means we get

$$IG + (A - aI)H = 1.$$

Because G and H are constants, this equation must hold true no matter what values of I we pick. So let's be clever and pick $I = 0$ (since that will eliminate the G part) and $I = \frac{A}{a}$ (since that will eliminate the H part). Thus

$$\begin{aligned} I = 0 : \quad 0 + AH = 1 &\quad \Rightarrow \quad H = \frac{1}{A} \\ I = \frac{A}{a} : \quad \frac{A}{a}G + 0 = 1 &\quad \Rightarrow \quad G = \frac{a}{A}. \end{aligned}$$

Substituting our values for G and H into the original fractions, we have

$$\frac{1}{(A - aI)I} = \frac{a}{A(A - aI)} + \frac{1}{AI}.$$

This is much, much easier, because we know how to integrate the right-hand side.

C

Eigenvalues

Eigenvalues are numbers that “represent” a matrix; if we have an $n \times n$ matrix A and can find a number λ and a nonzero vector x such that $Ax = \lambda x$, then λ is an eigenvalue and x is an eigenvector. Thus

$$\begin{aligned} Ax - \lambda x &= 0 \\ Ax - \lambda Ix &= 0, \end{aligned}$$

where I is the $n \times n$ identity matrix. We put this in so that λI is a matrix of the same size as the matrix A (the expression $A - \lambda$ would make no sense). Hence

$$(A - \lambda I)x = 0$$

Clearly we want $x \neq 0$ (or else this is all trivial). But if $(A - \lambda I)^{-1}$ exists (i.e. $\det(A - \lambda I) \neq 0$), then the only solution is $x = 0$. So there will only be eigenvalues when $\det(A - \lambda I) = 0$.

Thus, for the matrix

$$J|_{(N,0)} = \begin{bmatrix} 0 & b - aN \\ 0 & aN - b \end{bmatrix}$$

in Chapter 6, we have

$$\begin{aligned} 0 &= \det(J - \lambda I) = \det \left(\begin{bmatrix} 0 & b - aN \\ 0 & aN - b \end{bmatrix} - \lambda \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right) \\ &= \det \begin{bmatrix} -\lambda & b - aN \\ 0 & aN - b - \lambda \end{bmatrix} \\ &= -\lambda(aN - b - \lambda). \end{aligned}$$

How did we get this last line? Eigenvalues of a 2×2 or a 3×3 matrix have a formula. For the former, we have

$$\det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = ad - bc,$$

and for the latter we have

$$\det \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & j \end{bmatrix} = aej + bfg + cdh - ceg - afh - bdj.$$

In general eigenvalues are quite hard... unless we have a row or column where all but one entry is zero. In this case we're allowed to reduce the size of the matrix by extracting that entry. But not only do we get to extract the entry, we get to eliminate everything else in that row and column!

Thus

$$\det \begin{bmatrix} a & b & c & 0 \\ d & e & f & 0 \\ g & h & j & 0 \\ k & m & n & p \end{bmatrix} = p \det \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & j \end{bmatrix}.$$

So not only does the p come out of the determinant (because everything else in the last column was zero), reducing the remaining determinant to a much more manageable 3×3 determinant, but the last row is simply gone. That is, k, m and n are out of the picture.

(Technical note: If you're extracting anything that's not one of the entries along the main diagonal then you may or may not need an extra minus sign when you extract it. We don't do any such extracting here, so you don't need to worry about it, but if you're interested, check out any undergraduate linear algebra textbook.)

Assuming you have matrices with lots of zeros, you can reduce very high order matrices down to 3×3 or 2×2 matrices using this method. Fortunately, the Jacobian matrix almost always has lots of zeros and the things that aren't are usually on the diagonals anyway, so life is a lot easier than it otherwise would be.

D

The R_0 sleight of hand

There's nothing actually incorrect, but there are a couple of dodgy bits. Moving "the negatives" to one side isn't as obvious as it might seem, since there's nothing inherent about positive or negative values. For example, we could apply the same reasoning if we simply added and subtracted 5 to equation (6.1):

$$\begin{aligned}aN + 5 - 5 - b &< 0 \\aN + 5 &< 5 + b \\ \frac{aN + 5}{5 + b} &< 1\end{aligned}$$

and then we could define an " R_0^{SIS2} " to be $\frac{aN+5}{5+b}$. This would have the same threshold properties (i.e., if $R_0^{\text{SIS2}} < 1$ then the disease dies out, whereas if $R_0^{\text{SIS2}} > 1$ then the disease will become endemic), but it clearly isn't the same value. Furthermore, it is highly unlikely to be the average number of secondary infections (since adding and subtracting 5 was pretty arbitrary).

We could obviously define an infinite number of threshold parameters in this way. However, there's more to it than that. When we have the condition

$$\frac{aN}{b} < 1,$$

it's by no means clear that we must necessarily define R_0^{SIS} the way we did. For instance, we could just as easily define

$$R_0^{\text{SIS3}} = \left[\frac{aN}{b} \right]^2$$

and we'd *still* have a threshold parameter with the right properties that is again unlikely to be the average number of secondary infections (since we arbitrarily squared).

In fact, how do we even know that our original value R_0^{SIS} is the average number of secondary infections? Answer: we don't. The whole question of matching the R_0 values derived from ODE models to the “true” R_0 is a fascinating and lengthy one, and we've only just scratched the surface.

E

Finding eigenvalues for the case of permanent immunity

This is a much more complicated system than we've seen previously. The first step we'd like to do is find the equilibria, but even that is going to be a lot of work (feel free to try it for yourself, though). However, one of the equilibrium points is quite easy: the disease-free equilibrium occurs when $\bar{H}_I = 0$. (Since it's a *disease-free* equilibrium, there has to be no disease at this equilibrium.) From the last equation, we immediately conclude that $\bar{H}_R = 0$ (why?). Similarly, from the second equation, $\bar{M}_I = 0$ (which we'd expect anyway, given that we're looking for the disease-free equilibrium; we could have started with this assumption just as easily). From the first equation, we have $\bar{M}_S = \frac{\lambda^M}{\mu^M}$, and from the third equation, we have $\bar{H}_S = \frac{\lambda^H}{\mu^H}$. So the disease-free equilibrium is thus

$$(\bar{M}_S, \bar{M}_I, \bar{H}_S, \bar{H}_I, \bar{H}_R) = \left(\frac{\lambda^M}{\mu^M}, 0, \frac{\lambda^H}{\mu^H}, 0, 0 \right).$$

Why is this so useful? Answer: because if we can calculate the Jacobian matrix, we only need one equilibrium, namely the disease-free equilibrium, in order to determine the long-term behaviour. So our next step is to calculate the Jacobian. As before, we differentiate every element with respect to every other element. The first column will be the partial derivatives with respect to M_S , the second column will be the partial derivatives with respect to M_I and so on. Thus

$$J = \begin{bmatrix} -\mu^M - \beta^M H_I & 0 & 0 & -\beta^M H_S & 0 \\ \beta^M H_I & -\mu^M & 0 & \beta^M M_S & 0 \\ 0 & -\beta^H H_S & -\mu^H - \beta^H M_I & 0 & 0 \\ 0 & \beta^H H_S & \beta^H M_I & -\mu^H - \gamma^H - \nu^H & 0 \\ 0 & 0 & 0 & \nu^H & -\mu^H \end{bmatrix}.$$

At the disease-free equilibrium, we have

$$J \Big|_{(\bar{M}_S, 0, \bar{H}_S, 0, 0)} = \begin{bmatrix} -\mu^M & 0 & 0 & -\beta^M H_S & 0 \\ 0 & -\mu^M & 0 & \beta^M M_S & 0 \\ 0 & -\beta^H H_S & -\mu^H & 0 & 0 \\ 0 & \beta^H H_S & 0 & -\mu^H - \gamma^H - \nu^H & 0 \\ 0 & 0 & 0 & \nu^H & -\mu^H \end{bmatrix}.$$

To calculate the eigenvalues, we thus evaluate the determinant

$$\det \begin{bmatrix} -\mu^M - \Lambda & 0 & 0 & -\beta^M H_S & 0 \\ 0 & -\mu^M - \Lambda & 0 & \beta^M M_S & 0 \\ 0 & -\beta^H H_S & -\mu^H - \Lambda & 0 & 0 \\ 0 & \beta^H H_S & 0 & -\mu^H - \gamma^H - \nu^H - \Lambda & 0 \\ 0 & 0 & 0 & \nu^H & -\mu^H - \Lambda \end{bmatrix}.$$

(We use Λ for the eigenvalues since λ was already taken.)

Finding the determinant of a high-order matrix like this is usually very messy... unless we have a row or column where every entry except one is zero. In this case, it's rather elegant: we take the sole remaining factor of that row or column out as a product and then evaluate the determinant of the remaining matrix when we eliminate *that entire row and column*. See Appendix C. So in our matrix, this becomes

$$(-\mu^M - \Lambda) \det \begin{bmatrix} -\mu^M - \Lambda & 0 & \beta^M M_S & 0 \\ -\beta^H H_S & -\mu^H - \Lambda & 0 & 0 \\ \beta^H H_S & 0 & -\mu^H - \gamma^H - \nu^H - \Lambda & 0 \\ 0 & 0 & \nu^H & -\mu^H - \Lambda \end{bmatrix}.$$

Applying the same trick twice more, to the second and last columns (both are columns in which all but one entry is zero, so we can extract that entry and eliminate everything else in their respective rows), we get

$$(-\mu^M - \Lambda)(-\mu^H - \Lambda)(-\mu^H - \Lambda) \det \begin{bmatrix} -\mu^M - \Lambda & \beta^M \bar{M}_S \\ \beta^H \bar{H}_S & -\mu^H - \gamma^H - \nu^H - \Lambda \end{bmatrix}.$$

(See Appendix C again if you're confused; we did two steps here.)

Rearranging the minus signs and remembering the definition of the determinant of a 2×2 matrix from Appendix C, we have

$$-(\mu^M + \Lambda)(\mu^H + \Lambda)^2 [(\mu^M + \Lambda)(\mu^H + \gamma^H + \nu^H + \Lambda) - \beta^M \beta^H \bar{M}_S \bar{H}_S] = 0.$$

F

Integrating factors

An integrating factor is a factor you multiply in order to convert two terms into the derivative of a single term. It's just the product rule for derivatives, but in reverse. It only works for linear equations, by which we mean equations that are linear in the variable to be differentiated. We don't care about nonlinearity in the other variable.

So to solve

$$\frac{dx}{dt} = a(t)x + b(t),$$

we have to make sure that this equation is linear in x (which it is; we don't care about nonlinearity in the functions $a(t)$ or $b(t)$). If there were something in front of the $\frac{dx}{dt}$ term, we'd divide that out first.

We then rewrite this as

$$\frac{dx}{dt} - a(t)x = b(t).$$

The integrating factor is

$$I(t) = e^{-\int a(t)dt}.$$

So we multiply every term by our integrating factor:

$$e^{-\int a(t)dt} \frac{dx}{dt} - e^{-\int a(t)dt} a(t)x = b(t)e^{-\int a(t)dt}.$$

The left-hand side can now be collapsed into a single term:

$$\frac{d}{dt} \left[e^{-\int a(t)dt} x \right] = b(t)e^{-\int a(t)dt}.$$

Why is this? Well, if we "undid" this by using the product rule for derivatives, then we'd have the previous line. What's so great about this is we now integrate both sides... the left-hand integral is dead easy, because we don't

have to do anything and hopefully the right-hand integral isn't too hard (but it usually isn't).

$$\begin{aligned}\int \frac{d}{dt} \left[e^{-\int a(t)dt} x \right] dt &= \int b(t) e^{-\int a(t)dt} dt \\ e^{-\int a(t)dt} x &= \int b(t) e^{-\int a(t)dt} dt + C \\ x &= e^{\int a(t)dt} \int b(t) e^{-\int a(t)dt} dt + C e^{\int a(t)dt}.\end{aligned}$$

This isn't nearly as hard as it looks, because each of the integrals is often quite straightforward, and we do them as we go, not in a big formula like this.

So to solve

$$\frac{dM}{dt} + \mu^M M = \lambda^M,$$

the integrating factor is

$$\begin{aligned}I(t) &= e^{\int \mu^M dt} \\ &= e^{\mu^M t}.\end{aligned}$$

Multiply everything by the integrating factor:

$$e^{\mu^M t} \frac{dM}{dt} + e^{\mu^M t} \mu^M M = \lambda^M e^{\mu^M t}.$$

We can now collapse the left-hand side into a single derivative:

$$\frac{d}{dt} \left(e^{\mu^M t} M \right) = \lambda^M e^{\mu^M t}.$$

Now integrate both sides:

$$\begin{aligned}e^{\mu^M t} M &= \frac{\lambda^M}{\mu^M} e^{\mu^M t} + C \\ M &= \frac{\lambda^M}{\mu^M} + C e^{-\mu^M t}.\end{aligned}$$

As $t \rightarrow \infty$, $e^{-\mu^M t} \rightarrow 0$, so

$$\begin{aligned}\lim_{t \rightarrow \infty} M(t) &= \frac{\lambda^M}{\mu^M} \\ &= \bar{M}_S.\end{aligned}$$

G

Trivial solutions for nonnegative constants

We want to examine the two differential equations represented by

$$\frac{\dot{T}}{DT} = \frac{X''}{X} = C.$$

with boundary conditions

$$U(0, t) = U(L, t) = 0 \quad (\text{for } L > 0)$$

and nontrivial initial conditions. (We made a specific choice in Chapter 9, but the result is the same for any nonzero initial conditions.) That is $U(x, 0) \not\equiv 0$.

We have no idea what C is. It could be positive, negative or zero. So let's try all three.

Case i) $C = 0$.

This case is easy: the left-hand side implies that $\dot{T} = 0$, so that $T(t) = T_0$, a constant. Likewise, $X'' = 0$ so $X' = X_0$, which means the solution is $X = X_0x + X_1$.

Multiplying our two sub-solutions together, the solution to the PDE is thus

$$U(x, t) = T_0(X_0x + X_1).$$

Next we apply the boundary conditions:

$$\begin{aligned} U(0, t) &= T_0X_1 = 0 \\ U(L, t) &= T_0(X_0L + X_1) = T_0X_0L = 0. \end{aligned}$$

Since $L \neq 0$ (i.e., we actually have a corridor of some length), then either $T_0 = 0$ or $X_1 = X_0 = 0$. Either way, the solution is $U(x, t) \equiv 0$. But this is not possible with our initial conditions. Hence $C \neq 0$.

Case ii) $C > 0$

Solving the T equation, we have

$$\begin{aligned}\dot{T} &= CDT \\ \int \frac{\dot{T}}{T} dt &= CD \int dt \\ \ln \frac{T}{T_0} &= CDt \\ T &= T_0 e^{CDt}.\end{aligned}$$

Solving the X equation, we have

$$\begin{aligned}X'' &= CX \\ X'' - CX &= 0 \\ \left(\frac{d}{dx}\right)^2 X - CX &= 0 \\ \left[\frac{d}{dx} + \sqrt{C}\right] \left[\frac{d}{dx} - \sqrt{C}\right] X &= 0.\end{aligned}$$

We have two cases here, either of which could be zero. So let's look at them both.

$$\begin{array}{ll}\frac{dX}{dx} + \sqrt{C}X = 0 & \frac{dX}{dx} - \sqrt{C}X = 0 \\ X' = -\sqrt{C}X & X' = \sqrt{C}X \\ \frac{X'}{X} = -\sqrt{C} & \frac{X'}{X} = \sqrt{C} \\ \ln \frac{X}{X_0} = -\sqrt{C}x & \ln \frac{X}{X_0} = \sqrt{C}x \\ X = X_0 e^{-\sqrt{C}x} & X = X_0 e^{\sqrt{C}x}.\end{array}$$

Hence the general X solution is any combination of the above two solutions:

$$X = A e^{-\sqrt{C}x} + B e^{\sqrt{C}x}.$$

The full solution is thus

$$U(x, t) = T_0 e^{CDt} (A e^{-\sqrt{C}x} + B e^{\sqrt{C}x}).$$

Applying the first boundary conditions, we have

$$U(0, t) = T_0 e^{CDt} (A + B) = 0.$$

If $T_0 = 0$, then $U(x, t) \equiv 0$, and we know that won't satisfy the initial condition. Hence we have $A + B = 0$.

Applying the second boundary condition, we have

$$\begin{aligned}
U(L, t) &= T_0 e^{CDt} (A e^{-\sqrt{C}L} + B e^{\sqrt{C}L}) = 0 \\
A e^{\sqrt{C}L} - A e^{-\sqrt{C}L} &= 0 \quad (\text{since } B = -A) \\
e^{-\sqrt{C}L} - e^{\sqrt{C}L} &= 0 \quad (\text{since if } A = 0, \text{ then } U \equiv 0) \\
1 &= e^{2\sqrt{C}L} \\
2\sqrt{C}L &= 0,
\end{aligned}$$

which implies that either $C = 0$ or $L = 0$, neither of which can be true. It follows that C cannot be positive or zero and is hence negative.

H

Taylor's Theorem

Taylor's theorem is used to approximate a function (say $\sin x$ or e^x , but in fact it can handle just about any function) about a given value of x (which we'll call x_0) by an infinite series of polynomials. Taylor's theorem is less accurate the further one strays from this value of x_0 .

Suppose we wish to approximate our function $f(x)$ about $x = 0$; that is, $x_0 = 0$. Between our "starting point" $f(0)$ and some nearby point $f(b)$, we have a curve that we must try to duplicate mathematically to give us an approximation of $f(x)$ about 0 that we will call $F(x)$.

Our first attempt, F_1 , is a linear approximation: $F_1 = a_0 + a_1x$. This is a pretty simple (and not very good) approximation.

Our second attempt, F_2 , is a quadratic approximation (a parabola): $F_2 = a_0 + a_1x + a_2x^2$. F_2 is a better approximation of "the real thing" than F_1 (though by no means perfect). We can guess that a function of the form $F = a_0 + a_1x + a_2x^2 + a_3x^3 + \dots + a_nx^n + \dots$ would represent the given curve precisely so that $F = f$.

Assuming, then, that $f = a_0 + a_1x + a_2x^2 + a_3x^3 + \dots + a_nx^n + \dots$, we now evaluate f at 0.

$$\begin{aligned} f(0) &= a_0 + a_1(0) + a_2(0)^2 + a_3(0)^3 + \dots + a_n(0)^n + \dots \\ \Rightarrow f(0) &= a_0. \end{aligned}$$

Differentiating,

$$\begin{aligned} f'(x) &= a_1 + 2a_2x + 3a_3x^2 + \dots + na_nx^{n-1} + \dots \\ \Rightarrow f'(0) &= a_1 \\ f''(x) &= 2a_2 + (3)(2)a_3x + \dots + n(n-1)a_nx^{n-2} + \dots \\ \Rightarrow f''(0) &= 2a_2. \end{aligned}$$

(Note that the 2 is due to differentiation of x^2 .)

$$f'''(0) = (3)(2)a_3 = 3!a_3.$$

Remember that $m! = m(m-1)(m-2) \cdots (3)(2)(1)$ and $0! = 1$.

You can see the pattern here:

$$f^{(k)}(0) = k(k-1)(k-2) \cdots (3)(2)(1)a_k = k!a_k.$$

Dividing by $k!$, we have

$$a_k = \frac{f^{(k)}(0)}{k!}.$$

We now have a formula for the coefficients a_k , as long as we know $f(x)$ or have information about its derivatives at 0. Thus

$$f(x) = f(0) + \frac{f'(0) \cdot x}{1!} + \frac{f''(0) \cdot x^2}{2!} + \frac{f'''(0) \cdot x^3}{3!} + \cdots + \frac{f^{(k)}(0) \cdot x^k}{k!} + \cdots$$

This is Taylor's theorem.

For example, if $f(x) = e^x$,

$$\begin{array}{ll} f(x) = e^x & f(0) = 1 \\ f'(x) = e^x & f'(0) = 1 \\ f''(x) = e^x & f''(0) = 1 \\ & \text{etc} \end{array}$$

Thus

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \frac{x^4}{4!} + \cdots = \sum_{n=0}^{\infty} \frac{x^n}{n!}.$$

Similarly, if $f(x) = \sin x$,

$$\begin{array}{ll} f(x) = \sin x & f(0) = 0 \\ f'(x) = \cos x & f'(0) = 1 \\ f''(x) = -\sin x & f''(0) = 0 \\ f'''(x) = -\cos x & f'''(0) = -1 \\ f^{(iv)}(x) = \sin x & f^{(iv)}(0) = 0 \\ f^{(v)}(x) = \cos x & f^{(v)}(0) = 1 \\ & \text{etc} \end{array}$$

Thus

$$\sin x = 0 + x + 0 - \frac{x^3}{3!} + 0 + \frac{x^5}{5!} + \cdots = \sum_{n=0}^{\infty} \frac{(-1)^n x^{2n+1}}{(2n+1)!}.$$

I

Stability of periodic orbits in the logistic equation

Equilibria of $g(x)$ occur when $g(x) = x$. Thus

$$\begin{aligned} g(x) &= r^2x(1-x)[1-rx(1-x)] = x \\ r^2x(1-x)[1-rx(1-x)] - x &= 0 \\ r^3x^3 - 2r^3x^2 + (r^3 + r^2)x + r^2 - 1 &= 0 \end{aligned}$$

(where divided by x in the second-last line, because we are only looking for nontrivial equilibria).

This is a cubic, so it has three roots. However, if \bar{x} is a fixed point of f , then it must be a fixed point of g :

$$\begin{aligned} g(\bar{x}) &= f[f(\bar{x})] \\ &= f(\bar{x}) \\ &= \bar{x}. \end{aligned}$$

It follows that $x - \bar{x} = x - \frac{r-1}{r}$ must be a factor of the cubic.

Using long division to factor out the $x - \bar{x} = x - \frac{r-1}{r}$ term, we arrive at the quadratic for the remaining two roots:

$$rx^2 - (1+r)x + 1 + \frac{1}{r} = 0. \quad (\text{I.1})$$

Substituting into the quadratic formula, we have

$$\begin{aligned} x &= \frac{(1+r) \pm \sqrt{(1+r)^2 - 4r(1+\frac{1}{r})}}{2r} \\ &= \frac{(1+r) \pm \sqrt{1+2r+r^2 - 4r - 4}}{2r} \\ &= \frac{(1+r) \pm \sqrt{r^2 - 2r - 3}}{2r} \\ &= \frac{(1+r) \pm \sqrt{(r-3)(r+1)}}{2r}. \end{aligned}$$

Note that these two roots are only real when $r < -1$ (\bar{x}_1 is unstable) or when $r > 3$ (\bar{x}_2 is unstable). We ignore the range $r < 0$, but this indicates that Period 2 points only exist beyond $r = 3$, which is what we discovered from cobwebbing in Chapter 10's lab.

We thus have

$$\begin{aligned}|g'(w_1)| &= |f'(w_2) \cdot f'(w_1)| \\ &= |[r(1 - 2w_1)] \cdot [r(1 - 2w_2)]| \\ &= |r^2[1 - 2(w_1 + w_2)] + 4w_1w_2|.\end{aligned}\quad (\text{I.2})$$

Useful trick: When y and z are roots of a quadratic $x^2 + bx + c = 0$, then $y + z = -b$ and $yz = c$. If we divide equation (I.1) by r and substitute these in to (I.2), we have

$$\begin{aligned}|g'(w_1)| &= \left| r^2 \left[1 - \frac{2}{r}(1+r) + \frac{4}{r^2}(1+r) \right] \right| \\ &= |r^2 - 2r(1+r) + 4(1+r)| \\ &= |4 + 2r - r^2|.\end{aligned}$$

J

One-dimensional discrete stability condition

First suppose $|f'(\bar{x})| < 1$. Then, since f is continuous, there must be some (possibly small) interval $[\bar{x} - \epsilon, \bar{x} + \epsilon]$ around \bar{x} such that $|f'(x)| < c < 1$ on this interval.

The mean-value theorem says that if a function f is continuous and differentiable on an interval $[a, b]$, then there is at least one number ξ satisfying $a < \xi < b$ such that

$$f'(\xi) = \frac{f(b) - f(a)}{b - a}.$$

That is, there's at least one intermediate point whose tangent is parallel to the line joining the endpoints. (It doesn't sound that exciting, but it's one of the most important theoretical tools in calculus.)

By the mean-value theorem, with $x_0 \in [\bar{x} - \epsilon, \bar{x} + \epsilon]$, we have

$$\begin{aligned} |\bar{x} - f(x_0)| &= |f(\bar{x}) - f(x_0)| \\ &= |f'(\xi_1)| |\bar{x} - x_0| \\ &\leq c |\bar{x} - x_0|, \end{aligned}$$

where ξ_1 is between x_0 and \bar{x} . Since $c < 1$, it follows that $f(x_0)$ is also in the interval $[\bar{x} - \epsilon, \bar{x} + \epsilon]$.

Now we apply the mean-value theorem to $f(x_0)$ instead of x_0 :

$$\begin{aligned} |\bar{x} - f^2(x_0)| &= |f^2(\bar{x}) - f^2(x_0)| \\ &= |f'(f(\xi_2))f'(\xi_2)| |\bar{x} - f(x_0)| \\ &\leq c |\bar{x} - f(x_0)| \\ &\leq c^2 |\bar{x} - x_0| \end{aligned}$$

where ξ_2 is between \bar{x} and $f(x_0)$. By induction, it follows that $|\bar{x} - f^j(x_0)| \leq c^j |\bar{x} - x_0|$ for $j = 1, 2, \dots$

Thus $\lim_{j \rightarrow \infty} f^j(x_0) = \bar{x}$ (since $c < 1$). That is, repeated iterations of f will converge to the equilibrium for any initial condition sufficiently close to the equilibrium. It follows that \bar{x} is locally asymptotically stable.

Conversely, suppose that $|f'(\bar{x})| > 1$. Then, by continuity, there exists $\epsilon > 0$ such that, for $x \in [\bar{x} - \epsilon, \bar{x} + \epsilon]$, $|f'(x)| > c > 1$.

For $0 < |x_0 - \bar{x}| < \epsilon$, by the mean-value theorem, we have

$$\begin{aligned} |\bar{x} - f(x_0)| &= |f'(\xi_1)| |\bar{x} - x_0| \\ &\geq c |\bar{x} - x_0|, \end{aligned}$$

where ξ_1 is between \bar{x} and x_0 .

We can apply the argument again: if $|\bar{x} - f(x_0)| < \epsilon$, then $|\bar{x} - f^2(x_0)| \geq c^2 |\bar{x} - x_0|$. However, because $c > 1$, this argument cannot be continued indefinitely. There must exist a k such that

$$c^k |\bar{x} - x_0| > \epsilon.$$

Hence there exists k such that $|\bar{x} - f^k(x_0)| > \epsilon$. That is, no matter how close the initial condition is to the equilibrium, eventually it will be pushed away from it. It follows that \bar{x} is unstable.

K

The lower bound for λ_5

First note that, since $p > 0$, we have

$$\lambda_5 > g(\alpha B) = \frac{2 - \alpha B + \sqrt{(\alpha B)^2 - 4\alpha B}}{2}$$

$$g(4) = -1.$$

Differentiating, we have

$$g'(\alpha B) = -\frac{1}{2} + \frac{\alpha B - 2}{\sqrt{(\alpha B)^2 - 4\alpha B}}$$

$$= -\frac{1}{2} + \frac{\alpha B - 2}{\sqrt{\alpha B(\alpha B - 4)}}.$$

Differentiating again, we have

$$g''(\alpha B) = \frac{\sqrt{\alpha B(\alpha B - 4)} - \frac{(\alpha B - 2)(2\alpha B - 4)}{\sqrt{\alpha B(\alpha B - 4)}}}{\alpha B(\alpha B - 4)}$$

$$= \frac{(\alpha B)(\alpha B - 4) - (\alpha B - 2)(2\alpha B - 4)}{[\alpha B(\alpha B - 4)]^{3/2}}$$

$$= -\frac{(\alpha B)^2 - 4\alpha B + 8}{[\alpha B(\alpha B - 4)]^{3/2}}$$

$$= -\frac{(\alpha B - 2)^2 + 4}{[\alpha B(\alpha B - 4)]^{3/2}} < 0 \text{ for } \alpha B > 4.$$

It follows that g is concave down for $\alpha B > 4$.

Finally, we have

$$\begin{aligned}
\lim_{\alpha B \rightarrow \infty} g(\alpha B) &= \lim_{\alpha B \rightarrow \infty} \frac{2 - \alpha B + \sqrt{(\alpha B)^2 - 4\alpha B}}{2} \\
&= \lim_{\alpha B \rightarrow \infty} \frac{2 - \alpha B + \sqrt{(\alpha B)^2 - 4\alpha B}}{2} \cdot \frac{2 - \alpha B - \sqrt{(\alpha B)^2 - 4\alpha B}}{2 - \alpha B + \sqrt{(\alpha B)^2 - 4\alpha B}} \\
&= \frac{1}{2} \lim_{\alpha B \rightarrow \infty} \frac{(2 - \alpha B)^2 - (\alpha B)^2 + 4\alpha B}{2 - \alpha B + \sqrt{(\alpha B)^2 - 4\alpha B}} \\
&= \frac{1}{2} \lim_{\alpha B \rightarrow \infty} \frac{4}{2 - \alpha B + \sqrt{(\alpha B)^2 - 4\alpha B}} \\
&= 0.
\end{aligned}$$

A function that is initially negative, is concave down thereafter and has final limit at zero cannot be lower than its initial value. See Figure K.1. It follows that

$$\lambda_5 > g(\alpha B) \geq -1.$$

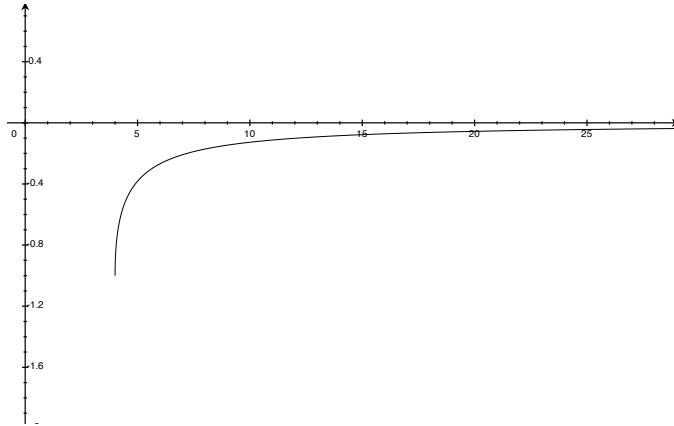


Fig. K.1. A function that is concave down, has negative initial value and horizontal asymptote at zero.

L

Solving the end-stage renal disease equations

We'd like to solve the equations for our model of end-stage renal disease:

$$\begin{aligned}\frac{dA}{dt} &= g \\ \frac{dN}{dt} &= sA - \delta N,\end{aligned}$$

with initial conditions $A(0) = A_0$, $N(0) = N_0$.

Luckily, we can solve the first equation independently of the second. We'll just do what we did before: separate variables, integrate both sides, use our initial condition and rearrange.

$$\begin{aligned}dA &= gdt \\ \int_0^t dA &= \int_0^t gdt \\ A - A(0) &= gt \\ A &= A_0 + gt.\end{aligned}$$

For the second equation, we can use this solution to write

$$\frac{dN}{dt} = s(A_0 + gt) - \delta N.$$

Does this look familiar? If not, have another look at Appendix F. What we need here is an integrating factor. So let's do what we did then: put the N s on one side, multiply by an integrating factor, integrate and rearrange.

$$\begin{aligned}
\frac{dN}{dt} + \delta N &= sA_0 + sgt \\
e^{\delta t} \frac{dN}{dt} + \delta e^{\delta t} N &= sA_0 e^{\delta t} + sgt e^{\delta t} \quad (\text{the integrating factor is } e^{\delta t}) \\
\frac{d}{dt} (e^{\delta t} N) &= sA_0 e^{\delta t} + sgt e^{\delta t} \quad (\text{the product rule in reverse}) \\
\int_0^t \frac{d}{dt} (e^{\delta t} N) dt &= \int_0^t (sA_0 e^{\delta t} + sgt e^{\delta t}) dt \\
e^{\delta t} N - N(0) &= \frac{sA_0}{\delta} e^{\delta t} + sg \int_0^t te^{\delta t} dt.
\end{aligned}$$

To do this last integral, we need integration by parts, with $u = t$, $u' = 1$, $v' = e^{\delta t}$ and $v = \frac{e^{\delta t}}{\delta}$. Thus, we have

$$\begin{aligned}
e^{\delta t} N - N_0 &= \frac{sA_0}{\delta} e^{\delta t} + sg \left[\frac{t}{\delta} e^{\delta t} - \int_0^t \frac{e^{\delta t}}{\delta} dt \right] \\
&= \left[\frac{sA_0}{\delta} e^{\delta t} + \frac{sgt}{\delta} e^{\delta t} - \frac{sg e^{\delta t}}{\delta^2} \right]_0^t \\
e^{\delta t} N &= N_0 + \frac{sA_0}{\delta} e^{\delta t} + \frac{sgt}{\delta} e^{\delta t} - \frac{sg e^{\delta t}}{\delta^2} - \frac{sA_0}{\delta} + \frac{sg}{\delta^2}.
\end{aligned}$$

Finally, we multiply everything by $e^{-\delta t}$:

$$N = N_0 e^{-\delta t} + \frac{sA_0}{\delta} + \frac{sgt}{\delta} - \frac{sg}{\delta^2} - \frac{sA_0}{\delta} e^{-\delta t} + \frac{sg}{\delta^2} e^{-\delta t}.$$

This is the explicit solution for the prevalence of end-stage renal disease.

M

Matlab overview

Matlab is a matrix-based, high-performance language for technical computing. It integrates computation, visualization and programming in an easy-to-use environment, where problems and solutions are expressed in familiar mathematical notation. The basic data element is an array that does not require dimensioning. This allows you to solve many technical computing problems, especially those with matrix and vector formulations, in a fraction of the time it would take to write a program in a scalar noninteractive language such as C or Fortran. The name Matlab stands for *matrix laboratory*.

In Matlab, a matrix is a rectangular array of numbers. Special meaning is sometimes attached to 1-by-1 matrices, which are scalars, and to matrices with only one row or column, which are vectors. Matlab has other ways of storing both numeric and nonnumeric data, but, in the beginning, it is usually best to think of everything as a matrix. The operations in Matlab are designed to be as natural as possible. Where other programming languages work with numbers one at a time, Matlab allows you to work with entire matrices quickly and easily.

There are typically three windows: the Command Window (where the main typing can be done), an editor (where you can type programs) and figures (where the figures can be displayed and edited).

At any time, you can type “help topic” (e.g., `help plot`) in the command window to find help on a particular topic. You can also press the help button for a comprehensive overview.

M.1 Variables and operators

Variable names consist of a letter, followed by any number of letters, digits or underscores. Matlab uses only the first 31 characters of a variable name. Matlab is case sensitive; it distinguishes between uppercase and lowercase letters. Thus `A` and `a` are not the same variable. To view the matrix assigned to any variable, simply enter the variable name.

For example, if you type `a=[1 2; 3 4]`, immediately the output will be a 2×2 matrix. You can see how to type in any matrix of any dimension you want from this example. To stop the output showing, end with a semicolon; Thus if you type `b=[5 6 7; 8 9 10];` there won't be any output (because the semi-colon has prevented it), but any any time you can type `b` to see your matrix.

You can also generate a vector (which, remember, is a matrix with only one row or column) by using a colon: `x=0:10;` will generate the vector `[1 2 3 4 5 6 7 8 9 10]`. The default gap between numbers is 1, but you can make it finer by using two colons: `x=0:0.1:10;` will generate a vector starting at 0 and ending at 10, with intervals of 0.1 (don't forget to end with the semi-colon or else you'll get a hideously long output, because such a vector is quite big).

You can use the regular function operators ($+, -, *, /, ^$), but remember that Matlab is matrix-based. So if `A` and `B` are matrices (or vectors) then `A*B` will also be a matrix. What if you just want to multiply the elements of two vectors together? In this case, use a period before the operator. Leaving out the period is the most common mistake for novice Matlab users. If `x` and `y` are vectors of the same size, then `x.*y` is the product. So `[1 2 3].*[4 5 6]=[4 10 18]`

Comments are indicated by a % sign.

M.2 M-Files

You can create your own matrices using M-files, which are text files containing Matlab code. Use the Matlab editor or another text editor to create a file containing the same statements you would type at the Matlab command line. Save the file under a name that ends in `.m`. For example, create a file containing these five lines.

```
A = [16.0      3.0      2.0      13.0
      5.0      10.0     11.0      8.0
      9.0      6.0      7.0      12.0
      4.0      15.0     14.0      1.0];
```

Store the file under the name `magik.m`. Then the statement `magik` reads the file and creates a variable, `A`, containing our example matrix.

M-files can be either scripts or functions. Scripts are simply files containing a sequence of Matlab statements. Functions make use of their own local variables and accept input arguments. The name of an M-file begins with an alphabetic character and has a filename extension of `.m`. The M-file name, less its extension, is what Matlab searches for when you try to use the script or function. A line at the top of a function M-file contains the syntax definition. The name of a function, as defined in the first line of the M-file, should be

the same as the name of the file without the .m extension. For example, the existence of a file on disk called **stat.m** with

```
function [mean,stdev] = stat(x)
n = length(x);
mean = sum(x)/n;
stdev = sqrt(sum((x-mean).^2/n));
```

defines a new function called **stat** that calculates the mean and standard deviation of a vector *x* that you've inputted in another program. It's imperative that you save the function with the same name as the function (in this case **stat.m**) or else Matlab won't know how to find it when you call it later. You cannot run function files themselves, only call them within another program.

Another program can call the function file like this:

```
x=[1 2 3 4 5 6];
[u v]=stat(x)
```

(or you could do this in the Command Window).

M.3 Figures

The function **plot** is a linear 2D plot. If you have vectors *x* and *y* that are the same size, then **plot(x,y)** will plot the *x* vector on the horizontal (*x*) axis and the *y* vector on the vertical (*y*) axis.

So to plot the function $y = x^2$, use this code:

```
x=-5:0.1:5;
y=x.^2;
plot(x,y)
```

Notice the period before the power: we want to square every element of the vector *x* not multiply the vector by itself (which isn't legal). You can experiment with coarser or finer intervals than 0.1 if you want (e.g., try *x*=-5:0.9:5 and see what happens).

You can also plot more than one thing at a time. E.g., if we want to plot $y = x^2$ and $z = x^3$ on the same graph, use this code:

```

x=-5:0.1:5;
y=x.^2;
z=x.^3;
plot(x,y,x,z)

```

See how we did that? First we plot the horizontal axis variable x , then the first vertical axis variable y , then x again, then the second vertical axis variable z .

M.4 Loops

When programming in Matlab, we can use **if**, **then**, **else** loops or **for** loops to create conditions under which a situation holds in order to execute certain actions. The syntax is:

```

if expression
    statements
end

```

When you are nesting **ifs**, each **if** must be paired with a matching **end**. When using **elseif** and/or **else** within an **if** statement, the general form of the statement is

```

if expression1
    statements1
elseif expression2
    statements2
else
    statements3
end

```

For loops repeat statements a specific number of times.

```

for variable = scalar:scalar
    statement
    ...
    statement
end

```

where **scalar:scalar** is a count from one number to another. The scope of the **for** statement is always terminated with a matching **end**.

M.5 Solving ODEs

To solve ODEs, we need to specify an input variable `y0` and a range of times `tspan` (which will be of the form `[t0 tf]`). We also need a function that evaluates the right-hand side of the differential equations, $y' = f(y)$. We'll save this function as a function file, in the form

```
function pdot=odefunction(t,p)
pdot(1,:)=...
pdot(2,:)=...
```

(and, again, we need to save the function file as `odefunction.m`). For a matrix, `x(1,:)` picks out the entire first row, whereas `x(:,1)` picks out the entire first column and `x(i,j)` picks out the scalar entry in the *i*th row and the *j*th column.

To put it all together, in our executable M-file, we specify the initial conditions and the range of times and call the function

```
[t,y] = ode23(@odefunction,tspan,y0);
```

(You can use `ode23` or `ode45` as you like.) This produces a matrix whose first column is the range of times and whose remaining columns define the solution of the ODE (in as many spatial variables as specified in the initial condition). So if you type `plot(t,y)` you'll have a graph of the solution.

M.6 Solving delay differential equations

Solving delay differential equations is almost exactly like solving ODEs, except that we need to specify the “prehistory” as well. This is because delay differential equations rely on what happened in the past, so, at time zero, we need to know what happened in negative time. The syntax for the prehistory function (assuming constant values in the prehistory) is

```
function s = histfunction(t)
s=...
```

(Don't forget to save this under the same name you give the function, in this case `histfunction.m`.) There's also a function file for the ODE function, of course, which also depends on the delay.

```
function pdot = odefunction(t,y,Z)
ylag1=Z(:,1);
ylag2=Z(:,2);
```

Thus, to call these two function files, we need to call the function files, specify the delay(s), call the prehistory function and give the time span.

```
sol = dde23(@odefunction,[ylag1 ylag2],@histfunction,tspan);
```

M.7 Glossary of other terms

Here are a list of other terms we've used. Remember you can always type `help term` to find out more about it.

- `axis` controls axis scaling and appearance. `axis([xmin xmax ymin ymax])` sets scaling for the x- and y-axes on the current plot.
- `clear all` clears any existing variables.
- `figure(h)` does one of two things, depending on whether or not a figure with handle `h` exists. If `h` is the handle to an existing figure, `figure(h)` makes the figure identified by `h` the current figure, makes it visible and raises it above all other figures on the screen. The current figure is the target for graphics output. If `h` is not the handle to an existing figure, but is an integer, `figure(h)` creates a figure and assigns it the handle `h`.
- `global` keeps track of variables across different files.
- `gtext` displays a text string in the current figure window after you select a location with the mouse.
- `input`: The command `userentry = input('prompt')` displays `prompt` as a prompt on the screen, waits for input from the keyboard and returns the value entered in `userentry`.
- `length(x)` gives the length of the vector `x` (i.e., the number of elements).
- `meshgrid` transforms the domain specified by vectors `x` and `y` into arrays `X` and `Y`, which can be used to evaluate functions of two variables and three-dimensional mesh/surface plots.
- `mesh` draws a 3D wireframe mesh from the `X-Y` meshgrid with colour determined by `Z` so that colour is proportional to surface height.
- `pause` by itself causes M-files to stop and wait for you to press any key before continuing. `pause(n)` pauses execution for `n` seconds before continuing, where `n` can be any nonnegative real number.
- `polyfit(x,y,n)` finds the coefficients of a polynomial $p(x)$ of degree n that fits the data, $p(x(i))$ to $y(i)$, in a least squares sense.
- `polyval(p,x)` returns the value of a polynomial of degree n evaluated at `x`. The input argument `p` is a vector of length $n + 1$ whose elements are the coefficients in descending powers of the polynomial to be evaluated.
- `rand` generates arrays of random numbers whose elements are uniformly distributed in the interval $(0,1)$.
- `subplot` divides the current figure into rectangular panes that are numbered rowwise. Each pane contains a set of axes. Subsequent plots are outputted to the current pane.
- `sum` If `A` is a vector, `sum(A)` returns the sum of the elements. If `A` is a matrix, `sum(A)` treats the columns of `A` as vectors, returning a row vector of the sums of each column.
- `title('string')` outputs the string at the top and in the center of the current figure.
- `xlabel` and `ylabel` label the `x`-axis and `y`-axis of the current axes.

References

1. J.L. Aron (2001). Mathematical modeling: The dynamics of infection In: Infectious disease epidemiology, theory and practice (Eds: K.E. Nelson, C.M. Williams, N.M.H. Graham), Aspen Pub., Gathersburg, MD, pp. 149–169.
2. R.M. Anderson & R.M. May (1991). Infectious diseases of humans. Oxford University Press, Oxford.
3. D.D. Bainov & P.S. Simeonov (1989). Systems with impulsive effect. Ellis Horwood Ltd, Chichester.
4. D.D. Bainov & P.S. Simeonov (1993). Impulsive differential equations: Periodic solutions and applications. Longman Scientific and Technical, Burnt Mill.
5. B.J. Beaty & W.C. Marquardt (1996). The Biology of disease vectors. University Press of Colorado, Niwot.
6. S.M. Blower, T.C. Porco & G.H. Darby (1998). Predicting and preventing the emergence of antiviral drug resistance in HSV-2. *Nature Medicine* 4:6, 673–678.
7. K.R. Burnham & D.R. Anderson (2002). Model selection and multimodel inference. Springer, New York.
8. K.R. Burnham & D.R. Anderson (2004). Multimodel inference: Understanding AIC and BIC in model selection. *Sociological Methods & Research* 33, 261–304.
9. A.D. Cliff, P. Haggett, J.K. Ord & G.R. Versey (1981). Spatial diffusion: An historical geography of epidemics in an island community. Cambridge University Press, Cambridge.
10. R.F. Costantino, R.A. Desharnais, J.M. Cushing & B. Dennis (1997). Chaotic dynamics in an insect population. *Science* 275:5298, 389–391.
11. O. Diekmann, J.A.P. Heesterbeek & J.A.J. Metz (1990). On the definition and the computation of the basic reproduction ratio R_0 in models for infectious diseases. *J. Math. Biol.* 35, 503–522.
12. O. Diekmann & J.A.P. Heesterbeek (2000). Mathematical epidemiology of infectious diseases: Model building, analysis and interpretation. Wiley, New York.
13. H. Zu Dohna, M. Cecere, R. Gürtler, U. Kitron & J. Cohen (2007). Reestablishment of local populations of vectors of Chagas disease after insecticide spraying. *J. Appl. Ecol.* 44:1, 220–227.
14. C. Dye. (1992). Leishmaniasis epidemiology: The theory catches up. *Parasitology*. 104, Suppl:S7–18.
15. D.J.D. Earn, J. Dushoff & S.A. Levin (2002). Ecology and evolution of the flu. *Trends Ecol. Evol.* 17:7, 334–340.

16. F.R. Giordano & M.D. Weir (1985). A first course in mathematical modeling. Brooks/Cole, Monterey.
17. B.T. Grenfell & A.P. Dobson (1995). Ecology of infectious diseases in natural populations. Cambridge University Press, Cambridge.
18. R. Hagmann, J.D. Charlwood, V. Gil, F. Conceicao, V. do Rosario & T.A. Smith (2003). Malaria and its possible control on the island of Principe. *Malar. J.* 2:1, 15.
19. M.E. Halloran (1998). Concepts of infectious disease epidemiology. In: Modern epidemiology (Eds: K.J. Rothman & S. Greendland), pp. 529–554.
20. HIV/AIDS surveillance report. Centers for Disease Control and Prevention, Division of HIV/AIDS Prevention, National Center for HIV, STD and TB Prevention, Atlanta, 2006.
21. M.E. Halloran, C.J. Struchiner & A. Spielman (1989). Modeling malaria vaccines II: Population effects of stage-specific malaria vaccines dependent on natural boosting. *Math. Biosci.* 94:1, 115–149.
22. J.M. Heffernan, R.J. Smith & L.M. Wahl (2005). Perspectives on the basic reproductive ratio. *J. R. Soc. Interface.* 2:4, 281–293.
23. G.D.H. LaPeare, A.P. Pierce & T.R. Hurd (1998). Let's go modelling!: A student's introduction to modelling and MATLAB programming. McMaster, Hamilton.
24. K. Laurendon, C. Silleo-Zubiri, H. Thompson, F. Shiferaw, S. Thirgood & J. Malcolm (1998). Disease as a threat to endangered species: Ethiopian wolves, domestic dogs and canine pathogens. *Anim. Conserv.* 1, 273–280.
25. R. Levins (1968). Evolution in changing environments: Some theoretical explorations. Princeton University Press, Princeton.
26. M. Lipsitch, T. Cohen, B. Cooper, J.M. Robins, S. Ma, L. James, G. Gopalakrishna, S.K. Chew, C.C. Tan, M.H. Samore, D. Fisman & M. Murray (2003). Transmission dynamics and control of severe acute respiratory syndrome. *Science* 300, 1966–1970.
27. P.M. Luz, C.T. Codeco, E. Massad & C.J. Struchiner (2003). Uncertainties regarding dengue modeling in Rio de Janeiro, Brazil. *Mem. Inst. Oswaldo Cruz.* 98, 871–878.
28. P. Munz, I. Hudea, J. Imad & R.J. Smith? (2009). When zombies attack!: Mathematical modelling of an outbreak of zombie infection. In: Infectious disease modelling research progress (Eds: J.M. Tchuenche & C. Chiyaka), pp133–150.
29. O.G. Pybus, M.A. Charleston, S. Gupta, A. Rambaut, E.C. Holmes & P.H. Harvey (2001). The epidemic behaviour of the Hepatitis C virus. *Science* 292, 2323–2325.
30. E.J. Schwartz, L.A. Szczech, M.J. Ross, M.E. Klotman, J.A. Winston & P.E. Klotman (2005). Highly active antiretroviral therapy and the epidemic of HIV+ end-stage renal disease. *J. Am. Soc. Nephrol.* 16:8, 2412–2420.
31. C. Shan & H. Zhu (2014). Bifurcations and complex dynamics of an SIR model with the impact of the number of hospital beds. *J. Diff. Eq.* 257, 1662–1688.
32. S. Smith? (2023). The Top Ten Diseases of All Time. University of Ottawa Press, Ottawa.
33. R.J. Smith?, P. Cloutier, J. Harrison & A. Desforges (2012). A mathematical model for the eradication of Guinea Worm Disease. In: Understanding the dynamics of emerging and re-emerging infectious diseases using mathematical models (Eds: S. Mushayabasa & C.P. Bhunu), pp157–177.

34. R.J. Smith? & E.J. Schwartz (2008). Predicting the potential impact of a cytotoxic T-lymphocyte HIV vaccine: How often should you vaccinate and how strong should the vaccine be? *Math. Biosci.* 212, 180–187.
35. C.J. Struchiner, M.E. Halloran & A. Spielman (1989). Modeling malaria vaccines I: New uses for old ideas. *Math. Biosci.* 94:1, 87–113.
36. U.S. Renal Data System, USRDS 2006 Annual Data Report: Atlas of End-Stage Renal Disease in the United States, National Institutes of Health, National Institute of Diabetes and Digestive and Kidney Diseases, Bethesda, MD, 2006.
37. S. Ruan, D. Xiao & J.C. Beier (2008). On the Delayed Ross–Macdonald Model for Malaria Transmission. *Bull. Math. Biol.* 70, 1098–1114.
38. P. van den Driessche & J. Watmough (2002). Reproduction numbers and sub-threshold endemic equilibria for compartmental models of disease transmission. *Math. Biosci.* 180, 29–48.
39. M.J. Wonham, T. de-Camino-Beck & M.A. Lewis (2004). An epidemiological model for West Nile virus: Invasion analysis and control applications. *Proc. R. Soc. B* 271, 501–507.

Index

absolute deviation, 18, 19, 24–26
actors, 243
age distribution, 66, 75
ambassador of mathematics, 5
Anderson & May method, 64, 65, 68, 70
army, the, 237
asymptomatic infection, 166, 167, 169, 170
asymptotic contact, 141
asymptotic phase, 191, 192
attractivity, 185
 orbital, 191
 uniform, 185
autonomous systems, 177, 182, 185, 186, 190
behaviour changes, 10, 145, 221, 222, 233
Bernoulli, Daniel, 9
Bible, the, 221
bifurcation, 131–134, 136–138
birth, 46, 53, 69, 78, 82, 87, 88, 141, 142, 144, 145, 154, 155, 162, 163, 222, 232
mosquito, 81, 86
parasite, 222, 229–233
black holes, 4
boundary conditions, 109, 112, 266, 267
braaaiiinssss, 238, 243, 248
brain stem, severance of, 238
Brownian motion, 96
bullets, 238
carrying capacity, 143, 144
Carter, Jimmy, 232
cartoons, 238
chaos, 117, 124, 131, 136, 137, 243
characteristic equation, 83, 85, 167, 168, 171, 190, 240
characteristic exponents, 190
chlorination, 222, 224, 226–228, 230, 233
 fixed, 224, 226, 227, 234
 nonfixed, 224, 226, 228
civilisation, collapse of, 238, 249, 250
cobwebbing, 117, 121, 123, 124, 126–129, 132, 272
conservation, 93
 equation, 96
 law of energy, 96
 law of mass, 93, 95
cross infection, 82
curve fitting, 1, 13, 18–20, 22, 23, 25, 42, 200, 210
cytotoxic T-lymphocytes, 192
Dead Rising, 237
death, 46, 53, 54, 56, 57, 65, 68–70, 73, 76, 78, 87, 88, 141–145, 151, 154–156, 162, 199, 200, 203, 204, 218, 222, 232
zombies (permanent), 238, 240, 244
zombies (temporary), 238
background, 82, 142, 147, 222
disease-specific, 82, 86, 87, 89, 147, 155
mosquito, 81, 86, 217, 218

- parasite, 222, 229, 232
- delay differential equations, 56, 166, 213, 215, 217, 283
- demography, 10, 63, 140, 142, 145–147, 150, 151
- density, 13, 72, 76, 86, 93, 94, 96, 140
- determinant, 69, 73, 83, 259, 263
- difference equations, 66, 117, 118, 121, 128, 129, 131, 137, 138, 154, 174
- differential equations, 1, 48, 54–56, 66, 82, 83, 86, 118, 121, 122, 140, 142, 151, 174, 175, 177–180, 189, 207, 230, 250, 256, 266, 283
- diffusion, 96
 - equation, 97, 109, 110
- dinner parties, well-intentioned, 239
- Dire Straits records, 237
- discrete logistic model, 119, 126
- discrete-time model, 66, 154, 156, 157, 161
- disease
 - black plague, 9, 199
 - brucellosis, 47
 - Chagas' disease, 30
 - chicken pox, 9
 - cocoliztli, 199
 - dengue fever, 72, 75
 - Ebola, 65
 - end-stage renal disease, 199, 200, 203–210, 277, 278
 - foot and mouth, 9
 - Guinea-worm disease, 221, 222, 232, 233
 - hepatitis, 76, 232
 - herpes, 9, 79
 - HIV/AIDS, 9, 20, 47, 64, 65, 141, 192, 199–201, 203, 204, 207, 210
 - human papillomavirus, 9
 - influenza, 13, 17, 18, 23, 25, 47, 65, 118, 124, 140, 166, 167
 - Justinianic plague, 199
 - leprosy, 232
 - malaria, 9, 63, 71, 73, 78, 199, 213, 214, 217, 232
 - measles, 9, 13, 93, 96, 97, 99, 112, 154, 155, 161–164, 199, 232
 - MRSA, 34, 195
 - pandemic, 9, 13–15, 17, 49, 145–147
 - polio, 232
 - rubella, 232
 - SARS, 9, 64, 76
 - schistosomiasis, 47, 77
 - smallpox, 9, 118, 122, 199, 232, 233
 - Spanish flu, 9, 14, 15, 147, 199
 - swine flu, 146
 - TB, 9, 199
 - vector-borne, 63, 74, 78, 81, 90
 - water-borne, 221
 - West Nile virus, 9, 74, 75, 117
 - yaws, 232
 - yellow fever, 81, 87, 90, 232
 - zombies, 4, 5, 141, 237, 238, 240, 241, 243–251
- disease modelling, 1, 4, 5, 10, 11, 45, 63, 85, 142, 166
- displacement, 97, 100–102, 104
- Doctor Who, 137
- doomsday scenario, 250
- drinking water, 221, 222, 224, 229, 233
- ecology, 86, 117, 142
- education, 2, 221, 222, 224, 229, 232, 233, 248
- eigenvalue, 66, 67, 69, 70, 72, 73, 76, 79, 83, 84, 168, 170, 171, 189, 190, 224, 240, 243, 258, 259, 263
- eigenvector, 66, 189, 258
- electric fences, 238
- electricity, 243
- endemic, 49, 53, 63, 64, 68, 77, 131, 232, 260
- epidemic, 45, 47, 49, 50, 56, 59, 71, 76, 77, 200
- equilibrium, 49–53, 65–68, 79, 83, 117, 121–124, 128, 129, 131–134, 137, 155–160, 164, 178, 179, 182, 191, 218, 239, 242, 250, 262, 271
 - disease-free, 65, 67–69, 73, 76, 79, 83, 131, 151, 156, 158, 159, 162–164, 167, 168, 217, 224, 239, 240, 242–247, 250, 262
- doomsday, 239, 240, 242
- endemic, 63, 68, 70, 75, 79, 84, 90, 151, 156–160, 162–164, 217, 224, 225, 233
- nontrivial, 132, 136, 271
- survivors', 247

- eradication, 85, 89, 90, 164, 210, 215, 221, 228, 229, 231–233, 248–250
 - of the human race, 240
- error, 13, 17, 35, 136, 137, 151
 - formulation, 17
 - measurement, 17, 18
 - round-off, 17, 124, 136
 - truncation, 17
- eventual starvation, 237
- ever-increasing force, 248, 250
- explicit solution, 129, 133, 209, 256, 278
- explicit time, 48, 56
- extrapolation, 14
- eyebrows, 4
- farmhouses, isolated, 237
- fashion sense, 248
- filtration, 222, 224, 229, 233
- final conditions, 174, 175
- final-size equation, 76
- Floquet theory, 187, 190, 193, 194
- flow, 47, 57, 93–96
 - geometric series, 176
- Google, disappearance of, 238
- government forces, 248
- grandma, hideously violent death of, 237
- growth, 10, 66, 76, 117, 118, 120, 142, 143
 - constant, 143
 - exponential, 66
 - geometric, 117
 - intrinsic growth rate, 76
 - linear, 118, 119, 144
 - logistic, 143–145, 151
 - maximal, 225
 - of models, 9
 - sub-geometric, 117
- Guinea-worm ceasefire, 232
- Guinness World Record, 4
- guns, 237
- Hamer, 9
- Harry, 97, 112, 113
- herd immunity, 55
- hideous shambling creature of doom, 244
- Higgs Boson, 4
- histogram, 127–129, 138
- hospital beds, 140, 146–148
- human flesh, 237
- human–zombie coexistence (impossibility of), 239
- immigration, 142
- immune system, 10
- immunity, 47, 54, 55, 59, 69, 76, 77, 162, 166, 221, 246, 250
 - permanent, 81, 82, 90, 262
 - temporary, 166, 167
- implicit, 48, 50, 76, 81, 117, 192, 193, 195, 196
- impulsive attacks, 248–250
- impulsive differential equations, 175, 184, 190, 194, 224, 225, 228
 - fixed time, 174, 176, 179, 185, 187, 192
 - nonfixed time, 177, 179, 184, 186
- impulsive effect, 179, 183, 185, 190, 191, 195, 221, 227, 248, 249
- impulsive semidynamical systems, 180–182, 194
- initial conditions, 89, 112, 113, 121, 126, 131, 174–178, 180, 183–186, 207, 208, 216, 243, 247, 256, 266, 277, 283
- integral, 15, 94, 95, 257, 264, 265, 277, 278
- integrating factor, 86, 87, 264, 265, 277
- integration by parts, 278
- interpolation, 14–17
- ivory tower, 5
- Kermack & McKendrick, 9
- kill ratio, 248
- knife-edge case, 49, 144
- larvae, 74, 222, 224, 225, 230, 232
- latent period, 59, 65, 77–79, 241, 242, 244
- least squares, 14, 103, 284
- limit, 95, 118, 176, 180, 276
- line of best fit, 14, 22, 23, 25, 42, 202, 206
- linear regression, 13, 22, 23, 42, 203, 206, 210
- logistic equation, 131, 136
- Lotka–Volterra, 53

Matlab, 1–3, 11, 13, 28, 29, 35, 40, 41, 78, 87, 89, 102, 114, 136, 151, 161, 208, 211, 214, 215, 217, 279
 delay, 216
 for loop, 101, 282
 histfunction, 283
 If/else statement, 59, 100, 282
 line of best fit, 29, 41
 M-File, 57, 59, 280, 283
 matrix, 35, 279, 280, 283
 mesh, 114
 meshgrid, 114
 ode45, 209, 211
 odefunction, 283
 plot, 42, 58, 102–104, 114, 115, 281, 282
 plot surface, 114
 polyfit, 28, 103
 polyval, 103
 random numbers, 100
 stat, 281
 sum, 101, 102
 tspan, 283
 vector, 35, 280, 281
 matrix, 1, 35, 39, 63, 65–67, 69, 73, 81, 83, 100–103, 105, 157, 159, 170, 187–190, 224, 258, 259, 263
 displacement, 100
 fundamental, 187–190
 identity, 187, 258
 inverse, 72, 73
 Jacobian, 63, 65–69, 76, 79, 83, 141, 156–159, 164, 167, 171, 239, 240, 242, 259, 262
 Leslie, 66
 next-generation, 72, 74
 periodic, 188
 random, 100, 101
 row, 102
 time, 102
 upper triangular, 159, 224
 mean distance, 97, 104–106
 mean-value theorem, 273, 274
 media, the, 4, 140, 145, 146
 mid-sized city, doomed, 240
 moats, 238
 modelling cycle, 10, 11, 250
 mosquitos, 9, 10, 72–75, 77, 82, 86–90, 194, 213, 214, 217, 218
 mummies, 221
 mutation, 55, 166
 negative people, 142, 145
 neglected tropical diseases, 221
 next-generation method, 72, 74, 75, 170, 171, 224, 242, 243, 245, 247
 Night of the Living Dead, 237
 Nobel prize, 9
 nonautonomous systems, 185, 187
 null hypothesis, 201, 203
 ODEs, 45–47, 49, 50, 55–57, 59, 63, 65, 71, 74, 81, 88, 93, 109–111, 118, 154, 161, 178–180, 187, 207, 210, 215, 223, 241, 256, 261, 283
 oscillation, 33, 34, 66, 124, 133
 overestimate, 224–226, 230, 233, 234
 parasite, 77, 81, 221, 222, 224–227, 229–234
 partial derivatives, 65, 66, 72, 95, 109, 262
 partial fractions, 255, 257
 PDEs, 13, 55, 93, 109, 110, 115, 118, 266
 periodic orbit, 117, 124, 131–134, 179, 182, 243
 impulsive, 177–179, 182, 183, 187, 191–193, 195, 196
 perturbation, 137, 240
 phase portrait, 50, 51, 53, 56, 58, 65, 87, 161
 photographic disease-recognition cards, 232
 power relationship, 141
 power series, 17
 probability, 45, 71, 81, 97, 140
 survival, 71, 77
 product rule, 264, 278
 pubs, British, 237
 pyrrhic victory, 245
 quarantine, 99, 106, 195, 243–246, 250
 quarks, 4
 R_0 , 63, 64, 66, 68, 70–72, 74–79, 84, 85, 89, 131, 151, 159, 160, 164, 167–172, 213–215, 224, 228, 229, 242, 245, 247, 260, 261

- realism, 241
- reanimation of the dead, 237–240, 244
- recurrence relation, 176
- regression coefficient, 23, 42, 203, 206, 210
- Resident Evil, 237
- respirators, 146
- Ross, Sir Ronald, 9, 213, 214, 218
- Ross–Macdonald malaria model, 213, 214, 218
- Routh–Hurwitz criterion, 85
- saddle point, 114
- SEIR model, 78
- separation of variables, 109
- shopping mall, 237, 241
- SI model, 46, 82, 164
- SIR model, 54, 55, 63, 69, 81, 82, 147, 154, 164
- SIS model, 47, 50, 63, 67, 68
- snipers, 244
- spatial heterogeneity, 17, 47
- spatial variable, 13, 50, 55, 64, 72, 81, 93, 99, 117, 118, 200, 283
- spatially explicit model, 109, 117
- splines, 29, 33, 38
 - interpolation, 33
 - clamped, 38, 41
 - cubic, 28–30, 34, 37, 38, 41, 42
 - interpolation, 14, 41
 - linear, 34–36, 40, 42
 - natural, 38, 39, 41
- stability, 53, 65–67, 70, 83, 84, 117, 121–123, 131–133, 137, 158, 162, 179, 184, 185, 187, 190, 243, 250
- bistability, 247
- orbital, 190–193, 195, 196
- stable, 51–53, 66–68, 84, 123, 124, 132, 134, 137, 157, 159, 160, 164, 217, 224, 240, 243, 274
- uniform, 185
- uniformly asymptotic, 185
- unstable, 52, 53, 66–68, 84, 123, 124, 132, 134, 157, 158, 160, 162, 217, 224, 240, 242, 243, 245, 247, 272
- state space, 48
- stupid idiot, 241
- sum of squares, 97
- survival function, 74
- symptomatic infection, 166, 169
- SZR model, 238
- Taylor’s theorem, 119, 120, 123, 128, 269, 270
- television, 243
- threshold, 63, 64, 66, 68, 70, 71, 75, 76, 84, 85, 168, 169, 177, 213, 227, 229, 234, 247, 260
- time series, 48, 49, 51, 250, 256
- time-to-peak, 174, 175
- trajectory, 52, 89, 98, 99, 104, 122, 124, 127, 128, 180–182, 185
 - impulsive, 181–183, 191
 - parabolic, 99, 106
- transmissibility, 64, 68, 70–72, 74, 88
- transmission
 - density-dependent, 140
 - mass-action, 46, 47, 140, 141, 147, 149–151, 222, 238
 - sexual, 238
 - standard incidence, 140, 141, 145, 147, 149–151
- undead, 237, 238, 246, 249, 250
- units, 56, 83, 87, 88, 93, 140–142, 148, 151, 250
- vaccination, 10, 45, 46, 54, 55
- vector, 72, 258
- vectorial capacity, 77, 78
- vertex, 83
- violence, mindless, 248
- weeping, 238
- zombie apocalypse, 237, 238
- zombies
 - cure for, 247
 - activation of, 241, 243, 247
 - cure for, 246, 247, 250
 - transformation into, 241
- zombies (aargh), 4, 5, 141, 237, 238, 240, 241, 243–251