

Mathematical Biology

University of Cambridge Part II Mathematical Tripos

David Tong

*Department of Applied Mathematics and Theoretical Physics,
Centre for Mathematical Sciences,
Wilberforce Road,
Cambridge, CB3 0BA, UK*

<http://www.damtp.cam.ac.uk/user/tong/mathbio.html>

d.tong@damtp.cam.ac.uk

Recommended Books and Resources

Here are a bunch of books that I've found useful while preparing for this course:

- J.D. Murray, *Mathematical Biology, Volumes 1 and 2*

This is the default reference for practitioner's in the field. It's a remarkably easy read and explains things with clarity without holding your hand. The material that we need for the first part of the course can be found in volume 1, but you'll need to open volume 2 when we get to spatially organised systems.

- Dick Neal, *Introduction to Population Biology*

The mathematics is straightforward and, at times, might come across a little laboured. But the explanations of the underlying biology are well written and the book is peppered with many interesting examples.

- Steven Strogatz, *Nonlinear Dynamics and Chaos*

Not a biology book per se, but a spectacularly good introduction to the dynamical systems that underlie much of this course. It has a number of biologically leaning examples.

In addition, there are many online lecture notes, including ones by past lecturers of this course that I have freely taken from. You can find links to these on the [course webpage](#).

Contents

0	Introduction	1
1	Population Dynamics and Other Stories	3
1.1	First You're Born, Then You Die	4
1.1.1	Exponential Growth	4
1.1.2	The Logistic Equation	6
1.1.3	Fixed Points	10
1.1.4	In Praise of Parsimony	13
1.2	Delay	14
1.2.1	The Linear Delay Differential Equation	16
1.2.2	Blowflies	19
1.2.3	And... Breathe	22
1.3	Age Concern	26
1.3.1	Separable Solutions	27
1.3.2	More General Solutions	28
1.4	Interacting Species	32
1.4.1	Predator-Prey: The Lotka-Volterra Equations	33
1.4.2	Predator-Prey: A Logistic Twist	39
1.4.3	Predator-Prey: I'm Full Now	42
1.4.4	Competition	44
1.4.5	Dengue Fever	49
1.4.6	The Large Diversity Limit	52
1.5	Epidemiology	54
1.5.1	The SIR Model	54
1.5.2	Just When You Thought It Was Safe...	57
1.6	Chemical Reactions	60
1.6.1	The Law of Mass Action	60
1.6.2	Michaelis-Menten Enzyme Kinetics	63
1.6.3	Asymptotic Behaviour	65
1.7	Neuron Excitations	67
1.7.1	FitzHugh-Nagumo Model	68

2	Discrete Time	72
2.1	Linear Examples	72
2.1.1	Hello Poppy	72
2.1.2	Breathe Again	73
2.2	The Logistic Map	75
2.2.1	The Fixed Points	76
2.2.2	Bifurcation	77
2.2.3	And Then. . . Chaos	80
2.2.4	The Logistic Map in Ecology	82
2.3	Universality	83
2.3.1	Zooming in on Bifurcations	85
2.3.2	Renormalisation	87
2.3.3	The Feigenbaum constant α	89
2.3.4	The Feigenbaum Constant δ	91
3	Spatial Variations	96
3.1	Diffusion	99
3.1.1	Diffusion on a Finite Interval	99
3.1.2	How to Cook a Turkey	102
3.1.3	A First Look at Diffusion With Growth	104
3.1.4	Diffusion on the Line	105
3.1.5	Non-Linear Diffusion	109
3.2	Travelling Waves	112
3.2.1	The KPP-Fisher Equation	113
3.2.2	Front Propagation in Bistable Systems	119
3.3	Turing Instability	124
3.3.1	Pattern Formation	128
3.3.2	An Example	129
3.4	Chemotaxis	130
3.4.1	An Example	132
4	Random Variations	135
4.1	Discrete Outcomes	135
4.1.1	Two Outcomes	135
4.1.2	Discrete Population Size	137
4.1.3	Birth and Death Again	140
4.1.4	Extinction	146
4.1.5	Multiple Populations: Wildebeest and Flies	147

4.2	Meet the Fokker-Planck Equation	149
4.2.1	Constant Drift and Diffusion	152
4.2.2	Birth and Death Once More	153
4.2.3	Fokker-Planck With More Variables	154
4.2.4	Wildebeest and Flies Again	156
4.3	An Invitation to Fluctuation and Dissipation	159

Acknowledgements

These lecture notes are far from original. They closely follow the notes of Julia Gog, Ray Goldstein, and Peter Haynes, all of whom gave previous versions of this course. Huge thanks to Julia, Ray, and Peter for putting so much preparation into this course and for helping me with my dumb questions. I'd also like to thank Andrew Gonzalez at McGill University, who knows a thing or two about ecology and was extremely generous in sharing it with me.

0 Introduction

“This process of “model building”, essentially that of discarding all but the essentials and focusing on a model simple enough to do the job but not too hard to see all the way through, is possibly the least understood – and often the most dangerous – of all the functions of a theoretical physicist.”

Philip Anderson

Imagine that you possess the following superpower. When presented with a new object, you pick up a pen and write down a description of it. Maybe you just jot down a few words. Maybe the muse takes hold and you become expansive, filling many paragraphs.

Now comes the superpower. When your description is finished, you stare at the page and have a little think before rearranging the words in a different order, following arcane rules known only to those who possess the superpower. You shuffle an adjective here and a verb there until a feeling of calm comes over you and the words form a pattern that, incredibly, when read, reveals something new about the original object, something that you didn't know before, something that must have been hiding there all along yet only emerges after you play this game. For those imbued with real skill, this superpower can be used to reveal things that no one on the planet previously knew.

This superpower sounds like magic. And yet it's a power that each of you can wield and the purpose of these lectures is merely to hone it. The trick, of course, is that you must describe the original object in the language of mathematics. The existence of such a superpower is the reason why mathematics is special. It is why mathematics is a greater form of expression than, say, poetry: partly because it's more useful, but mostly because it's more magical.

There are many stories about the world written in the language of mathematics. One of the mysteries of this enterprise is that, at the most fundamental level – at the level of quantum field theory and general relativity – the laws of physics are fully described in terms of the most simple equations. But that's not our interest here. Instead, we want to turn to topics in biology: population dynamics and the spread of disease and the interactions of enzymes and many others besides. These topics, like everything in biology, tend to be complicated. And this gives us an immediate headache because, for the kind of systems that we're interested in, no equation with fewer than, say, 10^{23} terms is likely to capture the full complexity of the situation. What to do?




Figure 1. A good model.

One option is to aim for realism. We could continue to add more and more terms to our equations, hoping to match the intricate details of the system, perhaps inspired by results elsewhere in physics where theoretical calculations agree with experiment to many decimal places. This is not the approach we will take here.

Instead, we will strive to create mathematical caricatures of biological systems. We will strip away the complexity and focus on the key principles that lie underneath. The result will be equations that are akin to Picasso's line drawings of animals: they capture the spirit of the beast, but they may not be particularly useful if you're a vet learning how to do surgery. The idea is that, with a few strokes of the pen, we can construct equations that describe the essence of a thing and then solve these equations to learn more about what it does and how it does it.

Throughout these lectures we will solve many such equations. However, one of the things that we will *not* do is to explain how to construct such models in the first place. This is an important skill, but it's typically a skill that first requires a deep knowledge of a particular topic – say, the migratory patterns of monarch butterflies, or the human respiratory system – before you then try to extract from the wealth of information the important features that can be profitably distilled into equations. We won't go through this long and complicated process in what follows, focussing instead on the end product: some simple examples of mathematical models. We will look at many such examples. The hope is that by exploring a wide range of mathematical models, and learning what sorts of behaviour you might expect to emerge, you will be better placed to construct your own models when you finally learn all there is to know about those monarch butterflies.

1 Population Dynamics and Other Stories

People are born. People die. People move on. This, in a nutshell, is population dynamics. Our goal in this section is to drape equations around these words to understand how the population of various species changes over time.

The general approach that we will take will be mirrored throughout these lectures: first we try to isolate the relevant dynamical degree of freedom; then we introduce models of increasing complexity to capture some basic idea about life or death or movement. Each of the complications will usually (but not always) change the form of our equations in some way and a large part of this chapter will be devoted to describing the techniques needed to solve the resulting equations.

The basic question that we have to address when building any mathematical model of the world is: what are the dynamical degrees of freedom? In other words, what are the right variables to use? For our immediate purposes, the answer would seem to be blindingly obvious: we want to understand the size of the population and how it changes with time. We denote this as $N(t)$.

Even here, however, there are subtleties and we may well need to further refine our dynamical variable or add more. Does it matter if we treat $N(t)$ as a continuous variable or should we insist that it's something discrete? Does the age profile of the population matter, in which case we should work with $n(a, t)$, the number of people with a particular age a . Does the way the population is distributed spatially matter, in which case we should work with a population density $n(\mathbf{x}, t)$. Does it matter if the population is interacting with some other species? Are there other things that we've just completely, perhaps even unintentionally, ignored?

The answer to all of these questions is: it depends. And even if you devote your life to the study of some particular population, it still may not be obvious. The art of mathematical modelling is largely in realising what you should include and what you can safely ignore. The way to proceed is to start simple and then introduce each of the potential complications above to see what qualitative and quantitative effect they have on the dynamics. Indeed, it's often only by comparing the data to the mathematical models that we can start to understand which complexities are important and which can be discarded.

As this section progresses, we will be increasingly flexible about what we think of as a "population". Later, we will look at "populations" of viruses, chemicals and electrical signals in the brain. The thing that links these different phenomena is that

they can all be described by coupled, first order differential equations. Indeed, the real purpose of this whole section is to see a few things that can be modelled by these simple dynamical systems. Unlike in other courses, we will not develop the mathematics systematically but, instead, work our way through various examples, most of which will exhibit different mathematical features. As we proceed, we will build up a toolkit of results that allow us to examine these kinds of systems more generally.

1.1 First You're Born, Then You Die

We're going to start by writing down the simplest models for population growth. These describe just the single variable $N(t)$ and we will take both the population N and time t to be continuous. This seems reasonable in the case of time, less so for the population which is, in reality, an integer. But we can justify this approximation if we're dealing with situation in which N is suitably large so that the difference between N and $N + 1$ doesn't make any material difference to the situation.

In what follows, we will ignore both immigration and emigration. This means that population change is determined by birth and death rates alone.

1.1.1 Exponential Growth

For our first attempt at writing down a model, we will assume that the population has some fixed birth and death rates. Each person in the population has some probability of giving birth and some probability of dying. We're not distinguishing population by age and so this probability is the same for everyone.

Suppose that in some time Δt , the number of births is $bN\Delta t$ and the number of deaths is $dN\Delta t$ with b and d constant. We call b and d the birth and death rates, respectively. More precisely, they are the number of births/deaths per individual (or *per capita*) per unit time.

In the time Δt , the change in the population will be

$$N(t + \Delta t) \approx N(t) + (b - d)N(t)\Delta t . \quad (1.1)$$

Dividing by Δt , and taking the limit $\Delta t \rightarrow 0$, we get the first order differential equation

$$\frac{dN}{dt} = (b - d)N . \quad (1.2)$$

This equation only depends on the difference of b and d , not their individual values. (As we will see in Section 4, this conclusion will change when we look at random fluctuations.) The equation is easily solved and we have

$$N = N_0 e^{rt} \quad \text{with} \quad r = b - d . \quad (1.3)$$

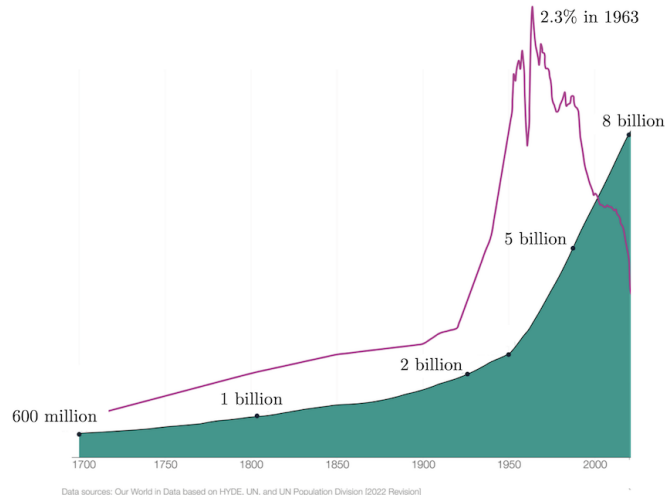


Figure 2. The growth of the world’s population, taken from the [Our World in Data](#) website. The purple line shows the percentage increase each year, so a constant purple line would correspond to exponential growth.

Here N_0 is the initial population at time $t = 0$. We learn that populations will grow exponentially if the birth rate is higher than the death rate, and shrink exponentially if it’s the other way around.

The idea of exponential population growth is often attributed to an essay by Thomas Malthus in 1798. He writes

“This [increase of population] constantly tends to subject the lower classes of the society to distress and to prevent any great permanent amelioration of their condition”

He was not the last person to deduce wide-sweeping, moralising conclusions about the state of society based on flimsy mathematical underpinnings.

So does this exponential growth hold up in practice? Well, like many things in this course, the answer is: to an extent. But often you have to seek out specific examples where it works. An obvious place to look is the global population of humans. The data is shown in Figure 2. For much of the past 300 years, the population growth has been super-exponential, with the exponent r growing over time. (This exponent is roughly what is plotted on the purple line.) But, since 1963, the exponent r has been decreasing. Current projections suggest that r will become negative before the year 2100, with the world’s population peaking at around 10 billion.

1.1.2 The Logistic Equation

In an attempt to be more realistic, we could envisage that the birth and death rates themselves depend on the size of the population. That leads us to the general class of models

$$\frac{dN}{dt} = [b(N) - d(N)]N \quad (1.4)$$

with general functions $b(N)$ and $d(N)$ that encode whatever features of population growth you think are important.

As we're treading slowly, an obvious guess for the right-hand side is to replace the linear function that led to exponential growth with a quadratic function. This suggests that we examine solutions to the equation

$$\frac{dN}{dt} = rN - \lambda N^2 . \quad (1.5)$$

This is a famous differential equation, known as the *logistic equation*. It was first written down in 1838 by the Belgian mathematician Pierre Francois Velhurst. The idea of the extra term is that growth rates become smaller as resources become scarce, perhaps due to lack of food or to overcrowding. The logistic equation comes with a natural maximum population that can be supported. This is most easily seen if we rewrite the equation as

$$\frac{dN}{dt} = r \left(1 - \frac{N}{K} \right) N \quad \text{with} \quad K = \frac{\lambda}{r} . \quad (1.6)$$

Here K is referred to as the *carrying capacity*.

If $r, \lambda > 0$, then small populations with $N < K$ will have an increasing population with $dN/dt > 0$. Meanwhile, larger populations with $N > K$ will have a decreasing population with $dN/dt < 0$. In this way, the logistic equation captures the idea that there is a natural ceiling to the size of a population.

It's straightforward to solve the logistic equation. We have

$$\begin{aligned} \int dN \frac{K}{(K-N)N} &= \int dN \left(\frac{1}{N} + \frac{1}{K-N} \right) \\ &= \log \left| \frac{N}{K-N} \right| - \log \left| \frac{N_0}{K-N_0} \right| = rt . \end{aligned} \quad (1.7)$$

Here we've introduced the integration constant N_0 which is designed so that $N(t = 0) = N_0$. We've taken the modulus sign in the logarithms because we don't know if




Figure 3. Solutions to the logistic equation. On the left, the curves are plotted with (from top to bottom) $N_0 = 2K$, $N_0 = \frac{3}{4}K$, $N_0 = \frac{1}{2}K$ and $N_0 = \frac{1}{2}K$. On the right, the curve is plotted with $N = \frac{1}{100}K$ where it exhibits the kind of characteristic tanh-like squashed s-shape that is sometimes called a *sigmoidal*.

$N > K$ or $N < K$. But if $N(t) > K$ at one time, then it remains so at all times (and similar for $N < K$.) This means that once we combine the two logs we can remove the modulus signs and write

$$\log \left(\frac{N}{K - N} \frac{K - N_0}{N_0} \right) = rt . \quad (1.8)$$

Rearranging, we have

$$N = \frac{N_0 K e^{rt}}{K + N_0 (e^{rt} - 1)} . \quad (1.9)$$

This general form $f(x) = a/(b + e^{-x})$ is known as the *logistic function*. The function is plotted in Figure 3 for various values of N_0 . Regardless of the initial value of the population N_0 , the curves converge on $N \rightarrow K$ at late times. The logistic function is well approximated by our earlier exponential function for times such that $N_0(e^{rt} - 1) \ll K$

There are other ways to write the solutions to the logistic equation that highlight the difference between the two cases $N > K$ and $N < K$. To see this, it's perhaps simplest to return to the integral solution (1.7) and write it as

$$\log \left| \frac{N}{K - N} \right| = r(t - t_0) . \quad (1.10)$$

Now we have the integration constant t_0 and we retain the modulus signs. Rearranging with $N < K$, we find

$$N = \frac{K e^{r(t-t_0)}}{1 + e^{r(t-t_0)}} = \frac{K}{2} \left[1 + \tanh \left(\frac{1}{2} r(t - t_0) \right) \right] . \quad (1.11)$$




Figure 4. On the left: the population of a single celled organism over time. This picture is taken from the classic, and wonderfully named, 1934 book *A Struggle for Existence* by G.F. Gause. (In addition, here is a [more recent paper](#) showing how to fit Gause’s data statistically.) On the right: the percentage of the total corn planted with hybrid seed in different US states, between 1932 and 1956, taken from the 1957 paper [Hybrid Corn: An Exploration in the Economics of Technological Change](#) by Griliches.

Meanwhile, for $N > K$, the same manipulations give

$$N = \frac{K e^{r(t-t_0)}}{e^{r(t-t_0)} - 1} = \frac{K}{2} \left[1 + \coth \left(\frac{1}{2} r (t - t_0) \right) \right]. \quad (1.12)$$

Again, we can ask: how well does the logistic curve do in modelling real world populations? It by no means a universal curve, but it certainly is more ubiquitous than exponential growth and you can find the characteristic s-shape curve appearing in many different places. One example is shown on the left of Figure 4 where the time evolution of the population of single celled organisms known as “paramecium caudatum” is plotted. (Unsurprisingly, when comparing data to equations, things tend to look better when the population in question lives in a bottle in controlled conditions, rather than in the wild.)

The Spread of Beneficial Mutation

There are many other situations where data seems to fit the logistic equation¹, notably when there is a natural ceiling in place. In particular, the equation arises when describing the fraction of a population that has adapted in some advantageous way. One

¹You can find claims that the growth of children’s vocabulary can be well fitted by the logistic function. You can see the data [here](#) and make up your own mind. It’s worth mentioning that the logistic function (1.9) coincides with the Fermi-Dirac distribution in theoretical physics, although I don’t know of a way to view that distribution as a solution to the logistic equation. (You can read more about the way this distribution arises in the lectures on [Statistical Physics](#).)

famous example is the percentage of farmers who adopted a particular breed of hybrid corn, as shown on the right of Figure 4.

For example, the fraction $p(t)$ of a population which carry a gene with a beneficial mutation is described by the logistic equation. To derive this, let $N(t)$ be the total population. A mutation is beneficial if it increases the rate at which individuals have offspring. Suppose that, in the absence of the mutation, the population grows at a rate r but that, with the mutation present, this is increased to $r + s$. At time t ,

$$N_{\text{mutant}}(t) = p(t)N(t) \quad \text{and} \quad N_{\text{normal}}(t) = (1 - p(t))N(t) . \quad (1.13)$$

At time $t + \delta t$, these numbers are

$$\begin{aligned} N_{\text{mutant}}(t + \delta t) &= (1 + (r + s)\delta t) pN \\ N_{\text{normal}}(t + \delta t) &= (1 + r\delta t) (1 - p)N \end{aligned} \quad (1.14)$$

where both p and N on the right-hand side are evaluated at time t and, here and below, we're dropping terms of order $\mathcal{O}(\delta t^2)$. This means that the total population at time $t + \delta t$ is

$$\begin{aligned} N(t + \delta t) &= N_{\text{mutant}}(t + \delta t) + N_{\text{normal}}(t + \delta t) \\ &= (1 + r\delta t + sp(t)\delta t)N(t) . \end{aligned} \quad (1.15)$$

Meanwhile, the fraction of mutants at time $t + \delta t$ is

$$\begin{aligned} p(t + \delta t) &= \frac{N_{\text{mutant}}(t + \delta t)}{N(t + \delta t)} \\ &= \frac{(1 + (r + s)\delta t)}{1 + r\delta t + sp(t)\delta t} p(t) . \end{aligned} \quad (1.16)$$

We now Taylor expand the denominator, again throwing away terms of order $\mathcal{O}(\delta t^2)$, to get

$$\begin{aligned} (t + \delta t) &= (1 + (r + s)\delta t) (1 - r\delta t - sp(t)\delta t) p(t) \\ &= (1 + s(1 - p(t))\delta t) p(t) . \end{aligned} \quad (1.17)$$

We then have

$$\frac{dp}{dt} = \lim_{\delta t \rightarrow 0} \frac{p(t + \delta t) - p(t)}{\delta t} = sp(t)(1 - p(t)) \quad (1.18)$$

which is the promised logistic equation.

Nondimensionalisation

Applied mathematicians have invented one of the ugliest words in the English language and then they go around saying it like it's completely normal. That word is *nondimensionalisation*.

The idea is to work with dimensionless variables. To achieve this, we absorb some of the constants in the equations into the dynamical variables. For example, for the logistic equation (1.6), we define

$$x = \frac{N}{K} \quad \text{and} \quad \tau = rt \tag{1.19}$$

so that the equation becomes

$$\frac{dx}{d\tau} = x(1 - x) . \tag{1.20}$$

Correspondingly, the solution (1.9) is

$$x = \frac{x_0 e^\tau}{1 + x_0(e^\tau - 1)} . \tag{1.21}$$

Nondimensionalisation has the advantage that it allows us to see the wood for the trees, stripping away anything that is inconsequential for the analysis. Of course, it's straightforward to put the constants back in by reverting to the original variables.

You will often see the practice of rescaling certain variables to make them nondimensional, but then retaining their original names. We will also be guilty of this at times, but will flag when we're doing it.

1.1.3 Fixed Points

In general, the class of population models that we're considering take the nondimensionalised form

$$\frac{dx}{dt} = f(x) \tag{1.22}$$

for some function $f(x)$. One particularly simple question that we can ask is: what are the steady state solutions, such that $dx/dt = 0$?

Clearly, the steady state solutions are the roots of the function, meaning those values $x = x_*$ such that $f(x_*) = 0$. Given such a fixed point, we can further ask: is it stable? If the population deviates a little from $x = x_*$, does it move towards x_* or away? The answer to this follows from some simple analysis. We write


$$x(t) = x_* + \epsilon(t) \tag{1.23}$$

with $\epsilon(t) \ll 1$. Taylor expanding, we then have

$$\frac{dx}{dt} = f'(x_*)\epsilon(t) + \mathcal{O}(\epsilon^2) . \tag{1.24}$$

We see that the fixed point $x = x_*$ is stable if $f'(x_*) < 0$ and is unstable if $f'(x_*) > 0$.

In practice, we can just plot the function $f(x)$ and see whether the graph crosses the axis with positive or negative slope. For example, for the logistic equation we have $f(x) = x(1 - x)$ which is plotted on the right. The fixed point at $x_* = 0$ has $f'(0) > 0$ and is unstable, while the other fixed point at $x_* = 1$ has $f'(1) < 0$ and is unstable.



An Example: Adding Predation

Here's a more intricate example. Suppose that we add an extra term to the logistic equation, so that it reads

$$\frac{dN}{dt} = r \left(1 - \frac{N}{K} \right) N - \frac{BN^2}{A^2 + N^2} . \tag{1.25}$$

This additional term represents the effects of predators, which increases monotonically as the population increases but reaches a maximum value as $N \rightarrow \infty$.

There are different ways to nondimensionalise this equation. We could define $x = N/K$ as before. But we will instead make the different choice

$$x = \frac{N}{A} , \quad \tau = \frac{Bt}{A} \quad \text{and} \quad \alpha = \frac{Ar}{B} , \quad \beta = \frac{K}{A} . \tag{1.26}$$

Here α and β are dimensionless parameters. With this rescaling, the equation becomes

$$\frac{dx}{d\tau} = f(x) = \alpha \left(1 - \frac{x}{\beta} \right) x - \frac{x^2}{1 + x^2} . \tag{1.27}$$




Figure 5. On the left: graphic solutions to equation (1.28). the dashed lines are plotted for a fixed value of β and varying α . On the right, the function $f(x)$ defined in (1.27) for values of the parameters where there are four roots.

There's one obvious fixed point at $x = 0$. The others are more complicated as we need to solve

$$\alpha \left(1 - \frac{x}{\beta}\right) = \frac{x^2}{1 + x^2}. \quad (1.28)$$

Multiplying this out gives us an unwieldy cubic. It has either one root or three depending on the parameters. To see this, we can look for solutions graphically. On the left of Figure 5, the function $x^2/(1+x^2)$ is plotted as a solid line, while the linear function $\alpha(1-x/\beta)$ is plotted for a fixed value of β and varying α . We see that for α large and small, there is just a single intersection point, while there is a window of values of α where there are three intersections.

If there is just a single additional root of $f(x)$, other than $x = 0$, then it is necessarily the stable point. If there are three additional roots, then two of them are stable and one is unstable. Indeed, a plot of $f(x)$ for a choice of α and β where there are four roots is shown on the right of Figure 5. You can see that, for any function $f(x)$, the fixed points have to alternate between unstable and stable. The solutions to (1.22) always sit between two fixed points and flow, asymptotically, to the stable one.

The equation (1.27), is thought to be a decent model for a class of tree-eating bugs, with the predation term arising because birds like bugs. These kind of models can inform the strategy that you take to reduce the infection by varying any parameters in the model that may be under your control. For example, you could spray the trees with something unpleasant to reduce the carrying capacity K of the bug. Or you could deliberately increase the population of predators, increasing the value of B .

1.1.4 In Praise of Parsimony


Throughout these lectures, we will typically start with some very simple model to describe a particular phenomenon. It will capture some aspect of what's going on but, when compared to data, we will usually find that things are more complicated. The obvious thing to do is to return to our original equations and improve them by, say, adding an extra term. That's what we did above when including the effect of predation.

It might be tempting to continue to add extra terms and parameters until the mathematical model agrees more closely with the data. We're going to resist this temptation. Mathematical biology doesn't, for the most part, have the precision of more fundamental areas of physics. We're never going to find the biological analog of, say, the blackbody radiation curve that we met in the [Statistical Physics](#) lectures and matches the light from the Big Bang to an accuracy of 10^{-5} . That's because humans and rabbits and bugs and viruses are all more complicated than atoms and photons. And including more and more terms in an attempt to match what you see runs the risk that you aren't adding any kind of insight beyond data fitting.

Instead our goal is more modest in scope. We would like to understand the qualitative features of a biological system, distilling its essence into a few simple equations.

For those who would prefer to see a closer match between predictions and data, it's worth recalling von Neumann's famous words to Fermi (all the more appropriate for a course on mathematical biology)

“With four parameters I can fit an elephant, and with five I can make him wiggle his trunk.”



An almost-realisation of this can be achieved by the parametric plot²

$$\begin{aligned}x(t) &= -60 \cos t + 30 \sin t - 8 \sin 2t + 10 \sin 3t \\y(t) &= 50 \sin t + 18 \sin 2t - 12 \cos 3t + 14 \cos 5t .\end{aligned}$$

²This is taken from the paper “[Drawing an Elephant with Four Complex Parameters](#)” by Jürgen Mayer, Khaled Khairy, and Jonathon Howard.

The result is shown in the figure. If you count only the amplitudes as parameters then you’ve got 8. The authors, somewhat cheekily, paired the parameters together in a random way and declared success with 4 complex numbers. (Presumably they could have equally argued that they did it with just 2 quaternions.)

1.2 Delay

The class of models (1.22) includes an arbitrary function $f(x)$, designed to capture the intricacies of population interactions. But the kind of solutions that we get are always going to be the same: the population will evolve monotonically towards a fixed point. This means that any population that exhibits different behaviour – say oscillations – must have something else going on.

In this section, we explore the mathematics of a different class of models. These have the property that the change in the population depends on its size evaluated at a previous time. This might be argued to be more realistic. For example, it takes time for eggs to hatch. And it takes time for babies to grow into fully functioning, and reproducing, members of society.

We will look at the *delayed logistic equation*, also known as the *Hutchinson-Wright equation*. In nondimensionalised variables, it is

$$\frac{dx(t)}{dt} = \alpha x(t) [1 - x(t - T)] . \quad (1.29)$$

As promised, the change in population depends on the current population both at the current time t and at a previous time $t - T$.

As we will see, the delayed logistic equation (1.29) exhibits interesting behaviour and is used widely to model various biological phenomena. But the equation doesn’t match the words that we’ve used to motivate it: the delay isn’t in the birth term, but instead it is in the carrying capacity term. This means that it is the resources available at the earlier time that affect the growth rate³. For example, you might think of animals

³This is the insight of Hutchinson in his 1948 paper “[Circular Causal Systems in Ecology](#)”. (He has a cute footnote saying that he turned to his friend, Lars Onsager, for help in solving the equation. Onsager has a formidable reputation as one of the greatest, and least comprehensible, physicists of the mid-20th century.) In contrast, Wright’s 1955 paper “[A non-linear difference-differential equation](#)” discusses the equation where the delay is in the birth rate

$$\frac{dx(t)}{dt} = \alpha x(t - T) [1 - x(t)] .$$

We’ll look at an equation in a similar spirit below when we discuss delay models applied to blowfly populations.

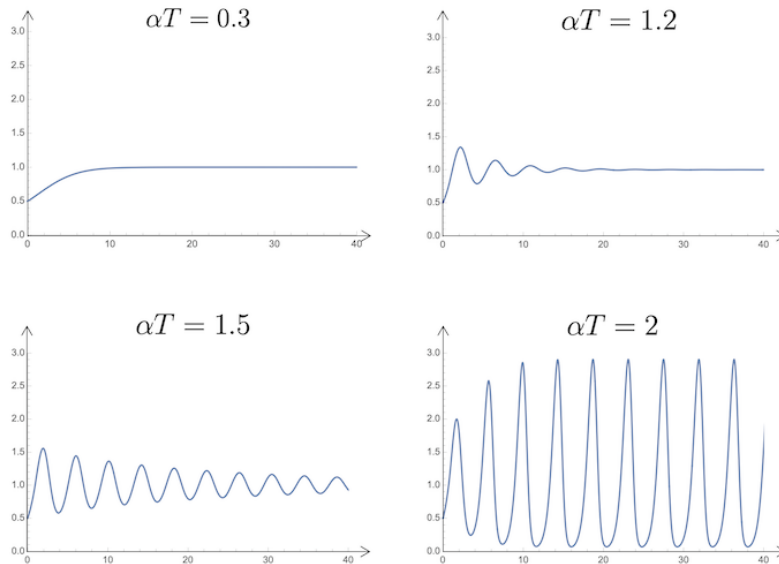


Figure 6. Numerical solutions to the delayed logistic equation, with $x(t)$ plotted vertically and t plotted horizontally. These were produced with $T = 1$ and α given by the value of αT specified in the graph. The initial history was taken to be $x(t) = \frac{1}{2}$ for $-1 \leq t \leq 0$.

storing nutrients, or hoarding resources, like squirrels and nuts, or humans and toilet paper.

The delayed logistic equation is an example of a general class of *delay differential equations*. These kind of equations are rare in theoretical physics, where things tend to happen locally in time, so are likely unfamiliar. Usually a first order differential equation needs just a single initial condition, say $x(t = 0)$. Here we need a functions worth of initial conditions, telling us the history $x(t)$ for $-T \leq t \leq 0$, before we can find a solution.

Even before we try to solve (1.29), we can get a feel for what will happen. In the original logistic equation, the population is capped above at $x = 1$ (in dimensionless variables). But the delay means that the population at time t can exceed $x = 1$ and continue to grow provided that the past population $x(t - T)$ is less than one. The population then stops growing only at a time T after it first hits the would-be ceiling at $x = 1$. Then the population will decrease. This story then repeats, resulting in oscillations around the $x = 1$ fixed point.

Numerical Results

Delay differential equations like (1.29) typically don't have closed form solutions. We will make some analytic progress shortly, but first we can study the solutions numerically. The equation has two parameters: α and T . But, after rescaling time, there is really just a single dimensionless parameter αT . We want to understand how solutions depend on this parameter.

The results are shown in Figure 6, plotted for various values of αT . We do indeed see the observed oscillations when αT is suitably large. But when αT is small, they die away. And by the time that $\alpha T \lesssim 0.5$, the oscillations are pretty much invisible. Our goal is to understand this behaviour.

1.2.1 The Linear Delay Differential Equation

The solution $x = 1$ is still a fixed point of the delayed logistic equation (1.29). We will make progress by looking at small deviations away from this fixed point. To this end, we write

$$x(t) = 1 + \epsilon(t) . \tag{1.30}$$

Substituting this into delayed logistic equation and dropping terms of order ϵ^2 , we find

$$\frac{d\epsilon(t)}{dt} = -\alpha \epsilon(t - T) . \tag{1.31}$$

This is still a delayed differential equation, but it's one of the few that we can solve analytically. These solutions take the form

$$\epsilon(t) = \epsilon_0 e^{st} \tag{1.32}$$

for some constant s . There's a standard trick that we use when solving linear equations of this kind: we take $s \in \mathbb{C}$. At first this looks slightly odd because ϵ is a (rescaled) measure of the population and when you're driving past those road signs that state the population of a town, they very rarely give a complex number. But because (1.31) is linear, if we've got a complex solution then we can simply take the real and imaginary parts to find real solutions.

At heart, the trick of taking $s \in \mathbb{C}$ is simply so that we can write things like e^{it} rather than $\cos t$ and $\sin t$. This means that any imaginary part of s is telling us that the solution oscillates about the fixed point. Meanwhile, a real part of s is telling us that the solution converges to the fixed point (if $\text{Re}(s) < 0$) or moves away from the fixed point (if $\text{Re}(s) > 0$).

The ansatz (1.32) solves our delay differential equation (1.31) provided that


$$se^{sT} = -\alpha . \tag{1.33}$$

If we rescale and define the variable $z = sT$, this becomes

$$ze^z = -\alpha T . \tag{1.34}$$

The solutions to this equation are somewhat subtle. We will proceed slowly.

First, we can look for purely real solutions. This is straightforward. We plot the curve $f(z) = ze^z$ in the figure. The right-hand side of (1.34) is negative so we see immediately that solutions must have $z < 0$. In other words, these are stable. The function has a minimum at $z = -1$ where $f(-1) = -e^{-1}$. This means that there are only purely real solutions provided that the decay time is suitably short,



$$\alpha T \leq \frac{1}{e} \approx 0.37 . \tag{1.35}$$

This agrees with what we saw numerically: for low values of αT , there are no oscillations. If $\alpha T \ll 1$, then we can zoom into the origin of the graph where $f(z) = ze^z \approx z$. We see that the solution for very small αT is just $s = \alpha$, which coincides with the solution to the logistic equation near the fixed point.

When $\alpha T > e^{-1}$, the solutions to (1.34) become complex. The general class of solutions are known as *Lambert W functions*. Once we allow complex values of z , there isn't a unique solution because of the branch cut inherent in taking the log of a complex number. Instead, the different branches of solutions are labelled by $k \in \mathbb{Z}$ and are written as

$$z = W_k(-\alpha T) . \tag{1.36}$$

Here $W_k(x)$ are the Lambert W functions. There are no closed form expressions, but their properties are very well studied. (In particular, there is an inbuilt Mathematica function `LambertW[k, x]` that will do the job for you.)




Figure 7. The real and imaginary parts of the Lambert W function $W_0(-x)$.

The purely real solutions that we have found above are part of the Lambert W function $W_0(x)$. The real and imaginary parts of this solution are plotted in Figure 7. We see that the imaginary part vanishes for $\alpha T < 1/e$, agreeing with what we found above. Meanwhile, the real part has a sharp kink at that point. For our immediate purposes, the important thing is that the real part is negative for $\alpha T < \pi/2$, and then becomes positive.

It's straightforward to reproduce these two key features with a little thought. First note that when we take $z = \pi i/2$, we have

$$z = \frac{\pi i}{2} \implies z e^z = -\frac{\pi}{2} \quad (1.37)$$

So this is indeed a purely imaginary solution when $\alpha T = \pi/2$.

Next, write $z = \sigma + i\omega$. If we rewrite (1.34) as $z = -\alpha T e^{-z}$, we have

$$\sigma + i\omega = -\alpha T e^{-\sigma} (\cos \omega - i \sin \omega) \implies \begin{cases} \sigma = -\alpha T e^{-\sigma} \cos \omega \\ \omega = +\alpha T e^{-\sigma} \sin \omega \end{cases} \quad (1.38)$$

where we have simply decomposed the first equation into real and imaginary parts. Looking at the first equation (for the real part), we see immediately that for $|\omega| < \pi/2$ we must have $\sigma < 0$, while for $|\omega| > \pi/2$, we have $\sigma > 0$.

The upshot of this argument is there are three distinct regimes, in which solutions to the delayed logistic equation have the following behaviour close to the stable fixed point at $x = 1$:

- For $0 < \alpha T < 1/e$, solutions monotonically approach the stable fixed point.

- For $1/e < \alpha T < \pi/2$, solutions oscillate about the stable, with the oscillations decaying exponentially.
- For $\alpha T > \pi/2$, the unstable fixed point is unstable. This isn't sufficient to tell us what actually happens. From the forth numerical solution shown in Figure 6, with $\alpha T = 2$, we see that the solution oscillates about $x = 1$ without decay.

We can, however, use our analysis to estimate the frequency of oscillations. For αT slightly greater than $\pi/2$, we have $\text{Im}(z) \approx \pi/2$ so the oscillations should be close to $e^{i\pi t/2T}$, meaning that they have approximate period $2T/\pi$. That's indeed what is seen numerically.

The discussion above relates only to the first branch of solutions, with $z = W_0(-\alpha T)$. There are also other branches with $z = W_k(-\alpha T)$. For example, it's simple to check that there exists purely imaginary solutions whenever

$$z = (2k + 1)\frac{\pi i}{2} \implies z e^z = -(2k + 1)\frac{\pi}{2} . \quad (1.39)$$

This is where $\text{Re}[W_k(-z)] = 0$ and this particular branch turns from stable to unstable. Importantly, however, higher branches only become unstable at larger values of αT . This ensures that, for $0 < \alpha T < \pi/2$, the fixed point $x = 1$ is stable.

You might then wonder whether it's possible to find stable solutions for larger values of αT , which use only these higher branches with $k \neq 0$. The answer is “yes”, but such solutions are finely tuned. To understand this, recall that we have to specify a whole history $x(t)$ for $-T \leq t \leq 0$ as an initial condition. We can decompose this as

$$x(t) = \sum_{k \in \mathbb{Z}} A_k e^{z_k t/T} \quad \text{for} \quad -T \leq t \leq 0 . \quad (1.40)$$

Here $z_k = W_k(-\alpha T)$ are the exponents from the different branches of the Lambert W functions, and A_k are coefficients that fix the pre-history. As we have seen, the $k = 0$ solution is unstable for $\alpha T > \pi/2$. If we have a pre-history with $A_0 = 0$ then we can discard this solution and extend the regime of stability further, with the $k = 1$ solution the first to turn unstable at $\alpha T = 5\pi/2$.

1.2.2 Blowflies

If you're a sheep farmer, then you really care about blowflies (little bastards). The flies lay their eggs in the fleece, often attracted by open wounds. This doesn't do the sheep much good: mortality rates are around 2 to 3%.




Figure 8. Data of blowfly population. The solid line is the adult population, the dotted line the eggs.

In a famous experiment in the 1950s, the British-Australian zoologist Nicholson kept blowfly as pets in the lab, regulating their supply of food for both the adult and larval population. Nicholson observed wild fluctuations in the blowfly population. His results are reproduced in Figure 8. These oscillations were subsequently explained by the kind of delay differential equation that we’re considering here⁴.

This time we will construct a model in which the birth rate is delayed. We assume that the egg production per capita is given by $P_0 e^{-N/N_0}$. The total egg production is then

$$P(N) = P_0 N e^{-N/N_0} . \quad (1.41)$$

We take a constant death rate μ . Our delay differential equation is then

$$\frac{dN(t)}{dt} = P(t - T) - \mu N(t) . \quad (1.42)$$

We have four constant parameters in this equation: P_0 , N_0 , μ and the delay time T . As a first step, we nondimensionalise our variables and parameters. We write

$$x = \frac{N}{N_0} , \quad \tau = \mu t , \quad \text{and} \quad a = \mu T , \quad b = P_0 T \quad (1.43)$$

This leaves us with the following equation that depends on just two dimensionless parameters, a and b , with a appearing as the delay time,

$$\frac{dx}{d\tau} = \frac{b}{a} x(\tau - a) e^{-x(\tau - a)} - x(\tau) . \quad (1.44)$$




Figure 9. A typical oscillatory solution from the delay differential equation with $b \gg a$.

First, we look at the equilibrium point $x = x_*$, obeying

$$\frac{b}{a}x_*e^{-x_*} = x_* \implies x_* = \log \frac{b}{a}. \quad (1.45)$$

This is a physical equilibrium point with $x_* > 0$ only if $b > a$. We will assume this in what follows.

Next, we look at small perturbations about the fixed point. We write

$$x(\tau) = \log \frac{b}{a} + \epsilon(\tau). \quad (1.46)$$

Substituting into our delay equation (1.45), we have

$$\begin{aligned} \frac{d\epsilon(\tau)}{d\tau} &= \frac{b}{a}(x_* + \epsilon(\tau - a))e^{-x_*}e^{-\epsilon(\tau-a)} - x_* - \epsilon(\tau) \\ &= (x_* + \epsilon(\tau - a))(1 - \epsilon(\tau - a)) - x_* - \epsilon(\tau) \\ &= (1 - x_*)\epsilon(\tau - a) - \epsilon(\tau). \end{aligned} \quad (1.47)$$

Now we follow our previous analysis. We look for solutions of the form

$$\epsilon(\tau) = \epsilon_0 e^{s\tau}. \quad (1.48)$$

These obey (1.47) if

$$s = (1 - x_*)e^{-sa} - 1. \quad (1.49)$$

Again, we can look for (typically) complex solutions to this equation. This time the solutions are harder to come by. But it's rather straightforward to prove the following result:

⁴Nicholson's classical study was [An Outline of the Dynamics of Animal Population](#). The model below was first propose by Gurney, Blythe, and Nisbet in [Nicholson's blowflies revisited](#).




Figure 10. The Hill function plotted for $m = \frac{1}{2}$, for $m = 1$, and for $m = 2$.

Claim: The fixed point $x_* = \log(b/a)$ is stable if $b/a < e^2$.

Proof: To see this, we decompose the equation (1.49) into real and imaginary parts by writing $s = \sigma + i\omega$. We have

$$\sigma = (1 - x_*)e^{-\sigma a} \cos(\omega a) - 1 \quad \text{and} \quad \omega = -(1 - x_*)e^{-\sigma a} \sin(\omega a). \quad (1.50)$$

If $b/a < e^2$ then $|1 - x_*| < 1$. Suppose that the system is unstable, so that $\sigma > 0$. Then the right we have $|(1 - x_*)e^{-\sigma a} \cos(\omega a)| < 1$ and hence the right-hand side of the first equation above is necessarily negative. However the left-hand side is just σ which, by assumption, is positive. Hence we must have $\sigma < 0$ in this regime. \square

Note that we haven't proven that the system is unstable for $b/a > e^2$ and, indeed, this isn't always the case. But it is true that, for suitably large b/a , the system is unstable and again exhibits oscillations. A typical numerical solution with $b \gg a$, showing pretty funky oscillations, is shown in Figure 9.

1.2.3 And... Breathe

While our primary focus in this section is to understand the evolution of population, the techniques that we're introducing are useful in many other contexts. Here we take a slight detour to explain how delay models can be used to model breathing.

There are two reasons that you breathe: one is to bring oxygen into your body, the other is to send carbon dioxide out. Because your body is good at keeping you alive, there is a feedback mechanism at play, and the volume V of the breath that you take depends on the concentration C of CO_2 in your blood. The relation between the two

can be modelled by an equation of the form

$$V(C) = V_{\max} \frac{C^m}{A^m + C^m} \quad (1.51)$$

for some constants A and m . Curves of this kind are sometimes called *Hill functions*. A variety of such curves, for fixed A and varying m , are plotted in Figure 10.

The concentration of CO_2 depends on time, so we have $C(t)$. Suppose that you take in CO_2 at a constant rate p . The amount that you exhale is proportional to both $C(t)$ and to the volume of your breath. But the feedback mechanism isn't entirely efficient, and the breath that you take at time t depends on the concentration at an earlier time $t - T$ for some delay T . The upshot is that we have a delay differential equation,

$$\begin{aligned} \frac{dC(t)}{dt} &= p - bC(t)V(t) \\ &= p - bC(t)V(C(t-T)) \\ &= p - bV_{\max}C(t)\frac{C(t-T)^m}{A^m + C(t-T)^m}. \end{aligned} \quad (1.52)$$

Here b is a constant that specifies how good your breathing is at removing CO_2 .

We can do our usual nondimensionalisation. We define

$$C' = \frac{C}{A}, \quad t' = \frac{pt}{A} \quad \text{and} \quad \beta = \frac{AbV_{\max}}{p}, \quad T' = \frac{pT}{A}. \quad (1.53)$$

This gives us a delay equation for $C'(t')$ but, to keep our equations cleaner, we then just relabel $C' \rightarrow C$ and $t' \rightarrow t$ and also $T' \rightarrow T$. The upshot is that we get the delay differential equation

$$\frac{dC(t)}{dt} = 1 - \beta C(t) \frac{C(t-T)^m}{1 + C(t-T)^m}. \quad (1.54)$$

First, we look for the equilibrium solution C_* , which sits at

$$\frac{C_*^{m+1}}{1 + C_*^m} = \frac{1}{\beta}. \quad (1.55)$$

There is a unique solution C_* to this equation. This corresponds to a regular breathing pattern. That is generally regarded as good. The question we would like to ask is: when is this breathing pattern stable?

The idea here is that something might change to alter your breathing pattern. In particular there is a famous irregularity in breathing known as *Cheyne-Stokes respiration*. If you get it, it's most likely that you have only days to live or you're on the top of a mountain (or possibly both). While it might not be much comfort, it's thought that this can be traced to some complicated physiology which results in a change of the coefficient m in the Hill function (1.51).

We look at small perturbations of form

$$C(t) = C_\star + \epsilon(t) . \quad (1.56)$$

By now, the path should be familiar. We substitute this into (1.54) to derive the linearised delay differential equation. A little algebra gives

$$\begin{aligned} \frac{d\epsilon(t)}{dt} &= -\frac{1}{C_\star}\epsilon(t) - \frac{m}{C_\star} \left(1 - \frac{C_\star^m}{1 + C_\star^m}\right) \epsilon(t - T) \\ &= -p\epsilon(t) - q\epsilon(t - T) . \end{aligned} \quad (1.57)$$

In the second line we've defined the positive constants $p = 1/C_\star$ and $q = m/C_\star(1 + C_\star^m)$. To solve this linear equation, we make the usual ansatz

$$\epsilon(t) = \epsilon_0 e^{st} . \quad (1.58)$$

This leaves us with the algebraic equation

$$s = -p - qe^{-sT} . \quad (1.59)$$

We decompose this into real and imaginary parts by writing $s = \sigma + i\omega$ to get

$$\begin{aligned} \sigma &= -p - qe^{-\sigma T} \cos(\omega T) \\ \omega &= qe^{-\sigma T} \sin(\omega T) . \end{aligned} \quad (1.60)$$

We can immediately read off some results.

First, when $T = 0$ so there is no delay, we have $\sigma = -(p + q) < 0$ so the system is stable.

Second, $\omega = 0$ is always a solution but there will be a second solution with $\omega \neq 0$ whenever $qe^{-\sigma T} > 1$

Third, the system is stable whenever $p > q$. To see this, note that we have $(\sigma + p)^2 + \omega^2 = q^2 e^{-2\sigma T}$, so $qe^{-\sigma T} > |\sigma + p|$. Suppose that the system is unstable, so that $\sigma > 0$. In this case, $q > qe^{-\sigma T} > |\sigma + p| > p$. So instability necessarily implies $q > p$. That isn't to say that all systems with $q > p$ are necessarily unstable; some may be stable. But all those with $q < p$ are definitely stable.





Figure 11. A numerical solution corresponding to an unstable breathing pattern.

This suggests that, if we want to look for when instability occurs, we should start with $p > q$, fix p and then increase q . The point of marginal stability occurs when $\sigma = 0$ and we have

$$\begin{aligned} p &= -\cos(\omega T) \\ \omega &= w \sin(\omega T) \end{aligned} \quad \Longrightarrow \quad \begin{aligned} p \tan(\omega T) &= -\omega \\ p^2 + \omega^2 &= q^2 \end{aligned} .$$



We can always find solutions to the first of these equations in the region $\omega T \in (\pi/2, \pi)$, as shown in the figure on the right where we plot $\tan x$ and $-x$. Call this intersection point $\omega_*(p, T)$. The second equation above then gives a necessary relationship between p , q and T for marginal stability:

$$q^2 = p^2 + \omega_*^2 . \quad (1.61)$$

For q less than this value, breathing is stable. For q greater than this value, you're in trouble.

There's a lot of objects above that we've only defined implicitly – like C_* and ω_* – but, given particular values of the parameters you could easily find these numerically. Nonetheless, the implicit nature of these constants makes it a little tricky to determine, say, the critical value of m in the function (1.51) in terms of T and β (or equivalently C_*). We can make some minor progress if we substitute $p = 1/C_*$ and $q = m/C_*(1+C_*^m)$ into (1.61). Using the fact that $\pi/2T < \omega_* < \pi/T$, we get that the critical value of m lies between

$$1 + \frac{C_*\pi^2}{4T^2} < \frac{m^2}{(1 + C_*^m)^2} < 1 + \frac{C_*^2\pi^2}{T^2} . \quad (1.62)$$

The left-hand inequality immediately tells us that instability only kicks for some value of $m > 1$. A numerical solution to the original delay differential equation (1.52) in the unstable regime is shown in Figure 11.

1.3 Age Concern


Until now, we've treated the population N as just a single number. But, for many questions, we may want to be more discerning. Here we want to take into account one particularly important variable: age.

To this end, we consider the population density $n(a, t)$. Here t is time, as before, while a is the age of individuals. Roughly speaking, $n(a, t)$ counts the number of individuals of age a . More precisely, we should think of $n(a, t) da$ as the number of individual with age between a and $a + da$. The total population at time t is

$$N(t) = \int_0^{\infty} da n(a, t) . \quad (1.63)$$

For mortals, we can reduce the upper limit of this integral to the lifespan of the species: a few weeks for blowfly, somewhat longer for humans.

It's reasonable to expect both the birth rate $b(a)$ and the death rate $\mu(a)$ to depend on the age of the population: the young are fertile, the old vulnerable. We might expect typical functions to look something like those shown on the right.



We now want to think about how the population ages. Suppose that your age is a at time t . It is a sad fact of life that at a short time δt later, you have only two options open to you: either your age is $a + \delta t$, or you're dead. In equations, this translates to

$$n(a + \delta t, t + \delta t) = n(a, t) - \mu(a)n(a, t)\delta t \quad (1.64)$$

where we're dropping terms of order δt^2 . We can Taylor expand the left-hand side and divide through by δt to get the partial differential equation

$$\frac{\partial n}{\partial t} + \frac{\partial n}{\partial a} = -\mu(a) n(a, t) . \quad (1.65)$$

This is sometimes called the *von Foerster equation*.

So far this equation only captures death. To inject a level of joy into the proceedings, we also need to talk about births. It's hard to be born at any age other than $a = 0$ and so the birth rates arise as a boundary condition on the function $n(a, t)$ at $a = 0$:

$$n(0, t) = \int_0^\infty da b(a) n(a, t) . \quad (1.66)$$

In any realistic situation, we don't have any problem with convergence: both the population and birth rate will vanish for some $a > a_{\text{ancient}}$.

A Cute Analogy with Fluid Mechanics

For a one-dimensional fluid with velocity u , we usually define the *material derivative*,

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + u \frac{\partial}{\partial x} . \quad (1.67)$$

This tells us how any quantity changes as we drift along with the fluid. For our population with age structure, we have the corresponding material derivative

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + \frac{\partial}{\partial a} . \quad (1.68)$$

The analog of the fluid velocity in this equation is just $u = 1$, corresponding to the fact that we get older by one year per year. In this case, the material derivative tells us how things change if we're aging with the population (which we are!).

The fluid analogy is nice enough, but it breaks down when it comes to the initial boundary condition (1.66) dictated by the birth rates. This is because, from the perspective of fluid mechanics, this is a “non-local” boundary condition: what happens at $a = 0$ depends on the value of the dynamical field $n(a, t)$ at all values of a .

1.3.1 Separable Solutions

Inspired by the Malthusian growth that we saw earlier, we will look for separable solutions of (1.65) of the form

$$n(a, t) = \tilde{n}(a) e^{rt} \quad (1.69)$$

for some function $\tilde{n}(a)$ and some $r \in \mathbb{R}$. This means that some age profile $\tilde{n}(a)$ remains unchanged, growing (or shrinking) exponentially. We substitute this into (1.65) to get

$$\frac{\partial \tilde{n}}{\partial a} = -[r + \mu(a)] \tilde{n}(a) \implies \tilde{n}(a) = n_0 e^{-ra} \exp\left(-\int_0^a ds \mu(s)\right) . \quad (1.70)$$

We see that the age profile is determined only by the death rate. That exponentiated integral $e^{-\int_0^a ds \mu(s)}$ has the interpretation of the probability of surviving to age a .

We still have to impose the birthing condition (1.66) which leaves us with the following, slightly unusual expression

$$\phi(r) := \int_0^\infty da b(a) e^{-ra} \exp\left(-\int_0^a ds \mu(s)\right) = 1. \quad (1.71)$$

Here that double integral defines the function $\phi(r)$ which, you can see by inspection, is a monotonically decreasing function of r . The birthing condition requires $\phi(r) = 1$ which we view as determining the exponent r .

The population may grow or shrink depending on the sign of r . To get a handle on r , we can look at $\phi(0)$. Because $\phi(r)$ is monotonically decreasing, we can immediately say that

- If $\phi(0) > 1$ then $r > 0$ and the population grows in size.
- If $\phi(0) < 1$ then $r < 0$ and the population shrinks.

In this sense, $\phi(0)$ can be viewed as the average number of offspring per individual.

Note that adding age structure hasn't changed the overall story of exponential growth. In particular, it doesn't open up the possibility for oscillations in the way that delay equations did.

As a sanity check, we can look at what this model gives us in the case that birth and death rates are constant, so $b(a) = b$ and $\mu(a) = d$. Then the function $\phi(r)$ defined in (1.71) becomes

$$\phi(r) = \int_0^\infty da b e^{-(r+d)a} = \frac{b}{r+d}. \quad (1.72)$$

Evaluated at $r = 0$, we have $\phi(0) = b/d$, reproducing our old Malthusian result: if $b > d$ then the population increases exponentially, while if $b < d$ it shrinks exponentially.

1.3.2 More General Solutions

There is a more general class of solutions to the equation (1.65). To motivate this, note that the left-hand-side takes the form of a wave equation. If we could somehow banish death, so that $\mu(a) = 0$, then the solutions would be

$$\frac{\partial n}{\partial a} + \frac{\partial n}{\partial t} = 0 \quad \implies \quad n(a, t) = f(a - t) \quad (1.73)$$

for any function $f(a - t)$. This is telling us that there is some initial population profile $f(a)$ and it just propagates forward in the “age”-direction as people get older. Sadly,


death is unavoidable but, at least in this context, is easily dealt with and (1.65) has the solution

$$n(a, t) = f(a - t) \exp \left(- \int_0^a ds \mu(s) \right) . \quad (1.74)$$

This is that same age profile propagating forward, but now you lose a few as you go.

If we start with some initial age distribution $n(a, 0) = n_0(a)$, then that's enough to tell us what happens for all $a > t$. We have

$$\begin{aligned} n_0(a) &= f(a) \exp \left(- \int_0^a ds \mu(s) \right) \\ \implies n(a, t) &= n_0(a - t) \exp \left(- \int_{a-t}^a ds \mu(s) \right) . \end{aligned}$$



But this form of the solution holds only when $a > t$, which is the shaded region in the diagram. This is the region where it's sufficient to know the initial population profile. For the other region, with $t > a$, we need to take the births into account. For this, we need the boundary condition that specifies the birth rate (1.66). Substituting the wave-like solution (1.74) into this boundary condition gives

$$n(0, t) = f(-t) = \int_0^\infty da b(a) n(a, t) . \quad (1.75)$$

We have the solution for $n(a, t)$ when $a > t$, but not when $a < t$. We can make use of this by splitting the integral above into two pieces and writing

$$\begin{aligned} f(-t) &= \int_0^t da b(a) n(a, t) + \int_t^\infty da b(a) n(a, t) \\ &= \int_0^t da b(a) f(a - t) e^{-\int_0^a ds \mu(s)} + \int_t^\infty da b(a) n_0(a - t) e^{-\int_{a-t}^a ds \mu(s)} . \end{aligned} \quad (1.76)$$

This is a really horrible equation! In principle, it should be used to determine the function $f(-t)$ for $t > 0$ by integrating $f(-t')$ from $t' = 0$ to $t' = t$. We then use this expression for $f(-t)$ in (1.74) to determine the population at times $t > a$. In practice, this is easier said than done.

There is one minor simplification that we can make. The second term in (1.76) gives the contribution from the offspring of the original population $n_0(a)$. The exponential suppression is telling us that must they survive from age $a - t$ to age a . At late times, we

expect that the number of these offspring will tend to zero as the old fail to reproduce, and the resulting population profile will be independent of the initial condition. Indeed, the separable solution didn't give us the opportunity to impose an initial condition. In this case, we have to solve the integral equation

$$f(-t) \approx \int_0^t da b(a) f(a-t) \exp\left(-\int_0^a ds \mu(s)\right). \quad (1.77)$$

It's still not a pleasant equation to solve.

An Example

We can illustrate these ideas with a simple example. We take a constant death rate, $\mu(a) = d$, and a birth rate that is a step function, where you can only give birth if you're younger than some fixed age A ,

$$b(a) = \begin{cases} b & 0 < a < A \\ 0 & a > A \end{cases}. \quad (1.78)$$

To keep things simple, we'll take our initial population to be $n_0(a) = 1$ for all ages a . Admittedly, this isn't particularly realistic (even Methuselah didn't make it past 1000) but, as we saw above, we expect that the initial population will soon be unimportant. We now need to solve for the population in two different regimes: $a > t$ and $a < t$. For $a > t$, we have simply

$$n(a, t) = e^{-dt} \quad \text{for } a > t. \quad (1.79)$$

The other regime $a < t$ is the trickier one. Here we must treat the regimes $t < A$ and $t > A$ separately. For $t < A$, the integral equation (1.76) becomes

$$\begin{aligned} f(-t) &= b \int_0^t da f(a-t) e^{-ad} + b \int_t^A da e^{-dt} \\ &= b \int_0^t d\tau \left(f(-\tau) e^{-d(t-\tau)} \right) + b(A-t) e^{-dt} \quad \text{for } a < t < A \end{aligned} \quad (1.80)$$

where, in the second line, we changed the integration coordinate to $\tau = t - a$. This is actually an equation we can solve. To do this, it's simplest to define $\tilde{n}(t) = n(0, t) = f(-t)$. We have

$$\tilde{n}(t) = b \int_0^t d\tau \left(\tilde{n}(\tau) e^{-d(t-\tau)} \right) + b(A-t) e^{-dt} \quad \text{for } a < t < A. \quad (1.81)$$

We can turn this into a more familiar differential equation simply by differentiating with respect to t . We have

$$\begin{aligned}\frac{d\tilde{n}}{dt} &= b\tilde{n}(t) - bd \int_0^t d\tau \tilde{n}(\tau)e^{-d(t-\tau)} - b(1 + d(A-t))e^{-dt} \\ &= (b-d)\tilde{n}(t) - be^{-dt} \quad \text{for } a < t < A\end{aligned}\tag{1.82}$$

where, to get to the second line, we've substituted in the integral expression in (1.81). But this equation is easily solved. We have

$$\tilde{n}(t) = \hat{N}e^{(b-d)t} + e^{-dt} \quad \text{for } a < t < A.\tag{1.83}$$

Putting this together, we get the following solution

$$n(a, t) = \begin{cases} e^{-dt} & \text{for } t < a \\ e^{-dt} [\hat{N}e^{b(t-a)} + 1] & \text{for } a < t < A \end{cases}.\tag{1.84}$$

A non-vanishing \hat{N} can be viewed as an injection of births and leads to a discontinuity in $n(a, t)$ along the line $a = t$.

We're left with understanding what happens for $A > 0$. Here we can follow the same steps and derive an integral equation for $\tilde{n}(t) = n(0, t) = f(-t)$.

$$\tilde{n}(t) = b \int_{t-A}^t d\tau \tilde{n}(\tau)e^{-d(t-\tau)} \quad \text{for } t > a, A.\tag{1.85}$$

This time, when we differentiate with respect to time, this becomes a delay differential equation

$$\frac{d\tilde{n}(t)}{dt} = (b-d)\tilde{n}(t) - \tilde{n}(t-A)e^{-dA}.\tag{1.86}$$

We saw how to deal with linear delay differential equations of this kind in Section 1.2. We can look for solutions of the form $\tilde{n}(t) \sim e^{st}$ where the exponent s must obey

$$s = (b-d) - be^{-(d+s)A}.\tag{1.87}$$

As usual, the population grows if $\text{Re}(s) > 0$ and shrinks if $\text{Re}(s) = 0$. The population is stable if $s = 0$, which, from (1.87), requires a fine tuning between birth and death rates, now given by

$$b = \frac{d}{1 - e^{-dA}}.\tag{1.88}$$

We see that we must have $b > d$ to maintain a stable population. This is because, in this model, only the young in the population can reproduce while anyone can die.

1.4 Interacting Species

So far we've discussed the evolution of just a single population. At times there was a different species lurking in the background, like when we included terms describing predation, but this other species was very much an NPC in the story.

Things become significantly more interesting when there are two or more populations that interact with each other. This allows for a much richer collection of dynamics.

For much of this section, we will restrict ourselves to just two independent species. But the basic mathematical formalism is just as easy to state regardless of the number. Suppose that we have n different dynamical variables, $u_i(t)$ with $i = 1, \dots, n$. We will study a class of first order dynamical systems that takes the form

$$\frac{du_i}{dt} = f_i(u_1, \dots, u_n) . \quad (1.89)$$

A fixed point of this system is a steady state solution $u_i = u_i^*$ obeying

$$f_i(u_1^*, \dots, u_n^*) = 0 \quad \text{for each } i = 1, \dots, n . \quad (1.90)$$

Given a fixed point, we would like to know whether it is stable or unstable. It's quite possible that it will be stable in some directions and unstable in others. To determine this, we expand

$$u_i(t) = u_i^* + \xi_i(t) \quad (1.91)$$

where $\xi_i(t)$ is taken to be a small perturbation. We then expand (1.89) to linear order in ξ :

$$\frac{d\xi_i}{dt} = f_i(u^*) + \left. \frac{\partial f_i}{\partial u_j} \right|_{u^*} \xi_j + \dots \quad (1.92)$$

The constant term vanishes precisely because we're at a fixed point. We're left with

$$\frac{d\xi_i}{dt} = J_{ij} \xi_j \quad \text{with} \quad J_{ij} = \left. \frac{\partial f_i}{\partial u_j} \right|_{u^*} . \quad (1.93)$$

The matrix J is called the *Jacobian*. The stability of a fixed point is determined by the eigenvalues of this matrix. Suppose that we have an eigenvalue λ with corresponding eigenvector \mathbf{x} ,

$$J\mathbf{x} = \lambda\mathbf{x} . \quad (1.94)$$

In general, λ could be real or complex. If $\text{Re}(\lambda) < 0$, then the fixed point is stable in the \mathbf{x} direction, while if $\text{Re}(\lambda) > 0$ the dynamics will be unstable. If λ includes a complex part, then the trajectories will typically spiral towards or away from the fixed point.

In general, we will also want to plot the trajectories in the space \mathbb{R}^n parameterised by u_i . We will gain plenty of experience in doing this as we go through various examples.

In much of what follows, we will consider situations with $n = 2$ species. In this case J is a 2×2 matrix with eigenvalues λ_1 and λ_2 and it's usually easiest to compute them by considering

$$\text{Tr } J = \lambda_1 + \lambda_2 \quad \text{and} \quad \det J = \lambda_1 \lambda_2 . \quad (1.95)$$

In particular, if $\text{Tr } J < 0$ and $\det J > 0$ then both eigenvalues must be negative and the fixed point is stable.

1.4.1 Predator-Prey: The Lotka-Volterra Equations

The poster child for mathematical biology is the *Lotka-Volterra model*. It's a simple and instructive model for two species interacting. But, as we will see, the results it gives are far from generic.

Here's the setting. There are two species. The first is the prey and has population $N(t)$. These are the cute things that you are most likely rooting for. The other is the predator with population $P(t)$. These are the villains of the piece. Their dynamics is governed by the set of equations

$$\frac{dN}{dt} = N(a - bP) \quad \text{and} \quad \frac{dP}{dt} = P(dN - c) \quad (1.96)$$

with a, b, c and d are all positive numbers. These are the *Lotka-Volterra equations*

The physics behind the various terms and minus signs is straightforward. The $+aN$ term is telling us that, in the absence of predators, the prey prospers (because it eats grass, or because it preys on something even more helpless that doesn't even get a mention in the equations). Meanwhile, the $-cP$ term tells us that, in the absence of prey, the predators die out. Each equation also contains an NP term that captures what happens when the two populations interact. The \pm signs tell us that the result is good for one, less good for the other.

As usual, we can do some rescaling. We write $u = dN/c$ and $v = bP/a$ and $t' = at$ and then, annoyingly, relabel t' as t . The end result is the set of non-dimensionalised equations

$$\frac{du}{dt} = u(1 - v) \quad \text{and} \quad \frac{dv}{dt} = \alpha v(u - 1) . \quad (1.97)$$

These equations depend on the single dimensionless parameter $\alpha = c/a > 0$.




Figure 12. Solutions to the Lotka-Volterra equations plotted in the phase plane on the left, and as a function of time on the right. Both of these were plotted with $\alpha = 0.5$.

Rather unusually, it's straightforward to find exact solutions to these equations. We can think of $P = P(N)$ or, equivalently, $v = v(u)$ to get a direct relationship between the two populations. Dividing the two equations in (1.97) gives

$$\frac{dv}{du} = \frac{\alpha v(u-1)}{u(1-v)}. \quad (1.98)$$

This differential equation is separable and we have

$$\int dv \frac{1-v}{v} = \alpha \int du \frac{u-1}{u} \implies \log v - v + \alpha(\log u - u) = \text{constant}. \quad (1.99)$$

We can plot these orbits in (u, v) -plane which, in this context, is also known as the *phase plane*. The resulting orbits are shown on the left of Figure 12. The trajectories all exhibit the same essential behaviour: they orbit the fixed point $(u, v) = (1, 1)$.

This is telling us that the populations of both species oscillate in time. It's more difficult to extract information about the time dependence, and this typically needs to be done numerically. An example is shown on the right of Figure 12.

We can also compute the average population $\langle u \rangle$ and $\langle v \rangle$. We take, for example, the first equation in (1.97) and write it as $\dot{u}/u = 1 - v$. Integrating the left-hand side over a single period T gives

$$\int_0^T dt \frac{\dot{u}}{u} = \int_{u_{\text{start}}}^{u_{\text{finish}}} du \frac{1}{u} = \left[\log u \right]_{u_{\text{start}}}^{u_{\text{finish}}} = 0. \quad (1.100)$$




Figure 13. Sales of hare and lynx pelts, taken from Murray’s book on Mathematical Biology.

This vanishes because the orbits are closed, so $u_{\text{start}} = u_{\text{finish}}$. Integrating the right-hand side over a single period must similarly vanish: we have

$$\int_0^T dt (1 - v) = T - T\langle v \rangle = 0 \implies \langle v \rangle = 1. \quad (1.101)$$

We see that the average sits at the fixed point which, back in our original variables, is $\langle P \rangle = a/b$. A similar argument shows that $\langle u \rangle = 1$.

From either of the plots in Figure 12, we can extract a story about the underlying ecology. First, the number of prey increases. After a short time, this results in a corresponding increase in the number of predators. This decreases the prey, resulting in less to eat and a decrease in predators which then allows the prey to thrive and so the cycle of life repeats.

An obvious question is: does this match what is seen in the wild? The answer is: sometimes but not often. As we will soon see, the mathematical structure of the Lotka-Volterra equations is rather special and doesn’t survive most perturbations. For that reason, we should be suspicious about its applicability. Nonetheless, there is a famous and rather wonderful story that seems to give support to this model. For many decades, the Hudson Bay Trading Company kept records of the numbers of pelts they sold. These include pelts of hares which are natural prey, and lynx which are natural predators. The results are shown in Figure 13 and are closely resemble the Lotka-Volterra oscillations. Clearly there are many other factors at play here, and so it’s unsurprising that the data is rather messy. One puzzle is that, around 1885, it looks like the lynx population rises before the hare population, which prompted many theories, the best of which can be [viewed here](#).




Figure 14. The population of prey plankton (in green) and predator algae (in red), both normalised to one, together with the trajectories in the phase plane. This data is taken from the paper [Long-term cyclic persistence in an experimental predator-prey system](#) by Blasius et al.

Cleaner data can be found for populations in bottles, rather than in the wild. The population of rotifer, which is a kind of plankton, and their algae prey are shown in [Figure 14](#), clearly exhibiting predator-prey cycles.

The Hamiltonian Structure of Lotka-Volterra

We see from [\(1.99\)](#) that there is a conserved quantity in the Lotka-Volterra equations,

$$H = v - \log v + \alpha(u - \log u) . \quad (1.102)$$

The value of H determines the chosen orbit. It takes its minimum value at the fixed point, where $H_{\min} = (1 + \alpha)$. Orbits that are further out have larger values of H .

The existence of a conserved quantity is reminiscent of the conserved energy in classical mechanics and you might wonder if there's a deeper connection. It turns out that the Lotka-Volterra system is an example of a Hamiltonian system of the kind that we met in the course on [Classical Dynamics](#). Specifically, we can view the coordinates $q = \log v$ and $p = \log u$ as position and momenta, and the Lotka-Volterra equations can then be written a

$$\frac{dq}{dt} = \frac{\partial H}{\partial p} \quad \text{and} \quad \frac{dp}{dt} = -\frac{\partial H}{\partial q} \quad \text{with} \quad H = e^q - q + \alpha(e^p - p) . \quad (1.103)$$

This Hamiltonian structure makes the Lotka-Volterra equations rather special and, to some extent, unrealistic. In physics, the fundamental laws are all Hamiltonian, a fact that can ultimately be traced to the quantum nature of reality. In mathematical biology, there is no such reason that the underlying laws should be pretty. Indeed, we will soon see more realistic generalisations of the Lotka-Volterra equations that do not preserve this Hamiltonian structure.

Stability Analysis

While we can understand the solutions to Lotka-Volterra equations by direct integration, this won't be true of the generalisations that we look at. To prepare ourselves, it will be useful to look again at the Lotka-Volterra equations from other perspectives which, although they are less powerful, will easily generalise to the more complicated situations.

The first method is to do a stability analysis of the fixed point. In fact, there are two fixed points at $(u, v) = (0, 0)$ and $(u, v) = (1, 1)$. The Jacobian matrix (1.93) takes the general form

$$J = \begin{pmatrix} 1 - v & -u \\ \alpha v & \alpha(u - 1) \end{pmatrix}. \quad (1.104)$$

To understand the stability, we need to compute the eigenvalues λ of this matrix at each of the fixed points. For the trivial fixed point we have

$$(u, v) = (0, 0) \implies \lambda = 1, -\alpha. \quad (1.105)$$

This means that the origin is a saddle point. This is fortunate, but unsurprising. It's fortunate because it means that our two species will not naturally be driven to extinction. It's unsurprising because we set things up such that the prey flourish in the absence of predators, while the predators suffer in the absence of prey.

The second fixed point is more interesting. We have

$$(u, v) = (1, 1) \implies \lambda = \pm i\sqrt{\alpha}. \quad (1.106)$$

So in this case, the eigenvalues are pure imaginary. Taken at face value, this tells us that the trajectories in the phase plane orbit the fixed point. To see this, note that the eigenvectors of the Jacobian are $\mathbf{x}_{\pm} = (\pm i, \sqrt{\alpha})$. The most general linearised solution is then

$$(u, v) = (1, 1) + A_+ \mathbf{x}_+ e^{i\sqrt{\alpha}t} + A_- \mathbf{x}_- e^{-i\sqrt{\alpha}t} \quad (1.107)$$

where A_+ and A_- can be viewed as initial conditions and should be chosen so that (u, v) is real. For example, we could take $A_+ = A_- = \frac{1}{2}$, in which case we have the trajectory

$$(u, v) = (1, 1) + (-\sin(\sqrt{\alpha}t), \sqrt{\alpha} \cos(\sqrt{\alpha}t)). \quad (1.108)$$

As we have seen, the trajectories in the Lotka-Volterra model do, in fact, orbit the fixed point but the linear analysis shown here isn't enough to demonstrate it. That's because, in general, a fixed point with purely imaginary eigenvalues could be rendered stable or unstable from higher order effects.

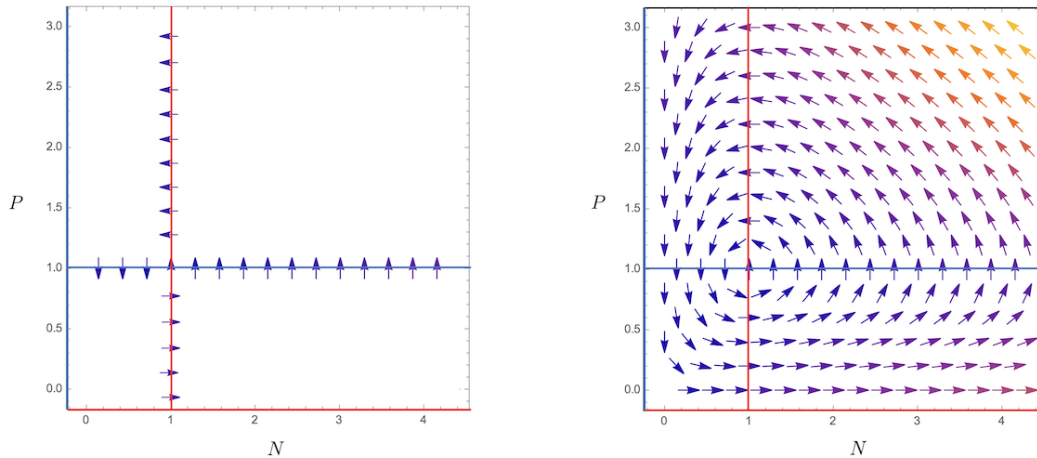


Figure 15. On the left, the nullclines divide the phase plane into four quadrants, where the flows are up/down and left/right. The fixed point sits at the point where the different nullclines meet. On the right, the full flow in the phase plane.

Nullclines

The second method to analyse dynamical systems of this type is to search for the *nullclines*. These are simply the curves

$$\begin{aligned} \dot{u} = 0 &\implies u = 0 \text{ or } v = 1 \\ \dot{v} = 0 &\implies v = 0 \text{ or } u = 1 . \end{aligned} \tag{1.109}$$

These are plotted on the left of Figure 15. On the red nullclines, we have $\dot{u} = 0$, and hence the flows are either left or right. On the blue nullclines we have $\dot{v} = 0$ and hence the flows are up or down. We necessarily have a fixed point whenever different nullclines meet because both $\dot{u} = \dot{v} = 0$. We see that the nullclines split the phase plane into quadrants where flows point north-east, north-west, south-west, or south-east. From this, we can surmise the general topology of the flows. The full flow structure is shown on the right of Figure 15.

Gone Fishing

We will see several generalisations of the Lotka-Volterra equations shortly. But there is one that is very straightforward and yet comes with a counterintuitive punchline.

Suppose that there is an additional effect at play, where the one or both of the species is hunted. This is usually phrased as “fishing” in the literature, as a nod to Volterra’s

initial motivation for writing down this system of equations. We could, for example, add an additional term to capture the effect of hunting the prey. In this case, the original equations (1.96) could be changed to

$$\frac{dN}{dt} = N(a - bP) - hN \quad \text{and} \quad \frac{dP}{dt} = P(dN - c). \quad (1.110)$$

The additional hN term captures the effect of hunting. This is straightforward because it doesn't change the structure of the equations at all, at least if $h < a$. Everything that we said above still holds, but the fixed point shifts to

$$(N, P) = \left(\frac{c}{d}, \frac{a - h}{b} \right). \quad (1.111)$$

The surprise is that hunting the prey hasn't changed the average prey population at all. Instead, counterintuitively, it has reduced the average predator population!

This cute effect is the reason why biologists first studied the Lotka-Volterra equations. (Lotka was the first to introduce the equation but that was in the context of chemical reactions.) The Italian biologist Umberto D'Ancona noted that the proportion of sharks and skates and other slightly dangerous predatory things dramatically increased in the Adriatic during the First World War when fishing was largely curtailed. He spent a long time puzzling over this until finally doing the sensible thing and asking a mathematician for help. Fortunately, his father-in-law was Vito Volterra, one of Italy's great 20th century mathematicians. This was the motivation for Volterra to write these equations and show that, as observed, in the absence of hunting, the predator population should increase.

1.4.2 Predator-Prey: A Logistic Twist

No set of equations is ever the last word in mathematical biology. There are always ways to finesse the model, to include some extra factor that may, or may not, change the qualitative behaviour. In this section and the next, we look at two such generalisations of the Lotka-Volterra equations.

We already introduced the logistic equation back in Section 1.1: it includes an additional term for a single species that limits the ability to reproduce as the population grows. We can easily incorporate such terms into our predator-prey model, so that they Lotka-Volterra equations (1.96) become

$$\frac{dN}{dt} = N \left(a - bP - \frac{N}{K_1} \right) \quad \text{and} \quad \frac{dP}{dt} = P \left(dN - c - \frac{P}{K_2} \right). \quad (1.112)$$

Here we've introduced two carrying capacities, K_1 and K_2 , for the two species. After rescaling, the equations can be written as

$$\frac{du}{dt} = u(1 - v - \mu_1 u) \quad \text{and} \quad \frac{dv}{dt} = \alpha v(u - 1 - \mu_2 v) \quad (1.113)$$

where the (inverse) carrying capacities are to be found in the positive constants μ_1 and μ_2 . We will take $\mu_i < 1$.

These equations have three fixed points, provided that $\mu_1 < 1$. There is the trivial fixed point $(u, v) = (0, 0)$ and a new fixed point in which the predators are extinct while the prey reach their logistic equilibrium: $(u, v) = (1/\mu_1, 0)$. The coexistence fixed point is now

$$(u^*, v^*) = \frac{1}{1 + \mu_1 \mu_2} (1 + \mu_2, 1 - \mu_1) . \quad (1.114)$$

Note that we have $v^* > 0$ when $\mu_1 < 1$. Note also that $u^* > 1$ and $v^* < 1$, so the additional terms have the effect of reducing the predator population while increasing the prey population at the fixed point.

We can perform the usual stability analysis. About the fixed point (1.114), the Jacobian is

$$J = \begin{pmatrix} 1 - v^* - 2\mu_1 u^* & -u^* \\ \alpha v^* & \alpha(u^* - 1 - 2\mu_2 v^*) \end{pmatrix} = \begin{pmatrix} -\mu_1 u^* & -u^* \\ \alpha v^* & -\alpha\mu_2 v^* \end{pmatrix} \quad (1.115)$$

where, to get to the second expression, we've used the fact that, for example, $1 - v^* - \mu_1 u^* = 0$. At this point, we don't need to use the explicit expressions in (1.114) for the fixed point: it's sufficient to stare at the signs of the terms. Recall that the determinant of a 2×2 matrix is the product of eigenvalues while the trace is the sum. We have $\det J > 0$ and $\text{Tr} J < 0$ which means that both eigenvalues must be negative. We learn that the fixed point is now stable.

The resulting flows in phase plane are shown in Figure 16, together with the nullclines which now sit at an angle. Note that there is an additional fixed point on the u -axis, where two different nullclines meet. On the right of Figure 16, the time-dependent solutions are plotted. This makes the dynamics clear: after a number of mild oscillations, the two populations settle down to an equilibrium with more (non-dimensionalised) prey than predators.




Figure 16. On the left: the flow in phase plane. The nullclines are shown in red and blue and determine where the flow is horizontal or vertical respectively. On the right, the time dependence of $u(t)$ and $v(t)$. Both plots were made with $\alpha = 0.5$, $\mu_1 = 0.3$ and $\mu_2 = 0.5$.

A Lyapunov Function

There is another way of seeing that all trajectories spiral into the fixed point. This comes from looking at the “Hamiltonian” that was constant on Lotka-Volterra orbits

$$H = v - \log v + \alpha(u - \log u) . \quad (1.116)$$

For our new model, it will prove useful to introduce the slightly different function

$$\tilde{H} = v - v^* \log v + \alpha(u - u^* \log u) . \quad (1.117)$$

This coincides with the Hamiltonian H when $u^* = v^* = 1$ for the Lotka-Volterra model. The importance of this new function can be seen by taking its time derivative

$$\begin{aligned} \frac{d\tilde{H}}{dt} &= \left(1 - \frac{v^*}{v}\right) \dot{v} + \alpha \left(1 - \frac{u^*}{u}\right) \dot{u} \\ &= \alpha(v - v^*)(u - 1 - \mu_2 v) + \alpha(u - u^*)(1 - v - \mu_1 u) . \end{aligned} \quad (1.118)$$

Now recall that $1 = v^* + \mu_1 u^*$ and $1 = u^* - \mu_2 v^*$. We use this to replace the 1’s in the expression above to get

$$\begin{aligned} \frac{d\tilde{H}}{dt} &= \alpha[(v - v^*)(u - u^* - \mu_2(v - v^*)) - (u - u^*)(v - v^* + \mu_1(u - u^*))] \\ &= -\alpha[\mu_2(v - v^*)^2 + \mu_1(u - u^*)^2] \\ &\leq 0 . \end{aligned} \quad (1.119)$$

That's rather nice. We've managed to construct a monotonically decreasing function $\tilde{H}(u, v)$ along the flow. Indeed, it stops changing only when we hit the fixed point (u^*, v^*) . This, again shows that all flows necessarily end up at the fixed point. This is an example of a *Lyapunov function*.

1.4.3 Predator-Prey: I'm Full Now

In both the original Lotka-Volterra model, and the logistic generalisation above, there is no end to the predator's appetite: they continue to voraciously work their way through the prey, like pac-man munching those dots.

A more realistic model might be to allow the predators to become satiated at some point, so that the amount of predation saturates as $u \rightarrow \infty$, rather than continuously increasing. Here is an example of a model that has this property,

$$\frac{du}{dt} = u \left(1 - u - \frac{bv}{u+d} \right) \quad \text{and} \quad \frac{dv}{dt} = \alpha v \left(1 - \frac{v}{u} \right) . \quad (1.120)$$

We see that there is novelty in both equations. For the prey population $u(t)$, the final predation term has the promised effect that, while proportional to the predators, the coefficient plateaus as $u \rightarrow \infty$. Meanwhile, the evolution of predators is different from what we had before. Now the reproduction term is $\alpha > 0$, but their carrying capacity is equal to the prey population.

We can use the techniques that we developed above to see the behaviour of this system. The non-trivial fixed point has $v^* = u^*$ and

$$1 - u^* = \frac{bu^*}{u^* + d} \quad \implies \quad u^* = \frac{1}{2} \left[1 - b - d + \sqrt{(1 - b - d)^2 + 4d} \right] \quad (1.121)$$

where we've picked the root that is positive. We can now check the stability. Here there's a short calculation to do in evaluating the Jacobian J at the fixed point. It's straightforward to show that

$$\det J = \alpha u^* \left[1 + \frac{db}{(u^* + d)^2} \right] > 0 . \quad (1.122)$$

This tells us that the eigenvalues λ are either both positive, both negative or, come in complex conjugate pairs. Stability is determined by whether $\text{Re } \lambda$ is positive or negative and we can see this by looking at the trace,

$$\text{Tr } J = u^* \left[\frac{bu^*}{(u^* + d)^2} - 1 \right] - \alpha . \quad (1.123)$$

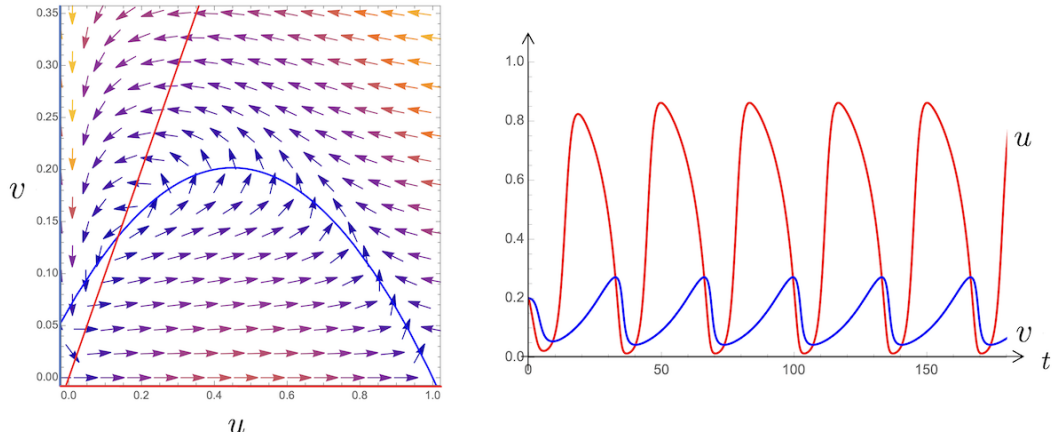



Figure 17. On the left: flows in the phase plane with an unstable fixed point, together with the nullclines. On the right, a numerical solution to the equations of motion with initial conditions that start close to the fixed point. Both of these plots were made with $b = 1.5$ and $d = 0.1$ which gives $\alpha_{\text{crit}} \approx 0.36$. We have then chosen $\alpha = 0.1$ to be in the unstable regime.

One way to read this is that for fixed b and d (and hence fixed u^*), there is a critical value of α given by

$$\alpha_{\text{crit}} = u^* \left[\frac{bu^*}{(u^* + d)^2} - 1 \right]. \quad (1.124)$$

For $\alpha > \alpha_{\text{crit}}$ the fixed point is stable and for $\alpha < \alpha_{\text{crit}}$ the fixed point is unstable.

The space of parameters is somewhat more interesting because for certain values of b and d , we have $\alpha_{\text{crit}} < 0$ so the system is stable for all $\alpha > 0$. An example of the resulting flows, together with the nullclines, is shown in the figure to the right (plotted with $b = 2$ and $d = 1$ and $\alpha = 0.5$).



Things are more interesting when the fixed point is unstable. If we don't end up at the fixed point, then where do we end up? The answer is that the dynamics converges towards a *limit cycle*, meaning a closed trajectory in the phase plane that attracts nearby trajectories. The flows in the phase plane, together with a numerical solution to the equations of motion are shown in Figure 17. We see that this model again gives rise to the oscillations in populations that was characteristic of the original Lotka-

Volterra model, but with one important difference: all initial conditions converge to the same cyclic behaviour.

The Poincaré-Bendixson Theorem


It's natural to ask: is there a way to demonstrate the existence of a limit cycle without resorting to numerics? There is a standard technique that can be applied to do this which we sketch here.

The general idea is to find a closed, bounded region $S \subset \mathbb{R}^2$ for which all flows go *into* S , and none come out. More precisely, we pick an outward pointing normal \mathbf{n} on the boundary ∂S of S and require that

$$\mathbf{n} \cdot \begin{pmatrix} \dot{u} \\ \dot{v} \end{pmatrix} < 0 \quad \text{everywhere on } \partial S. \quad (1.125)$$

This is telling us that, once inside the region S , you're trapped there forever. The *Poincaré-Bendixson theorem* then states (roughly) that if there are no fixed points in S then there will necessarily be a limit cycle.

Given this theorem, we can prove the existence of a limit cycle in our model if we can find such a region S . For our particular model, a sketch of the region S is shown shaded in the figure. It has two boundaries. An inner boundary consists of a small circle that excludes the fixed point and the flows are necessarily into S because the fixed point is unstable. The outer boundary is constructed so that (1.125) holds. This is straight forward to show for three of the four sides because the nullclines dictate the direction of the flow. The slightly tricky one is the left-hand boundary and one has to work a little harder to show that this too can be made to obey (1.125).



1.4.4 Competition

For our next example, we turn to a slightly different scenario, albeit one that is modelled by the same kind of equations. Rather than considering a populations of predators and prey, we will try to even the playing field a little. Instead, we consider two species competing for the same resources.

We call the populations N_1 and N_2 and describe their dynamics by the coupled logisticesque equations

$$\begin{aligned}\frac{dN_1}{dt} &= r_1 N_1 \left(1 - \frac{N_1}{K_1} - b_1 \frac{N_2}{K_2} \right) \\ \frac{dN_2}{dt} &= r_2 N_2 \left(1 - \frac{N_2}{K_2} - b_2 \frac{N_1}{K_1} \right) .\end{aligned}\tag{1.126}$$

We take $r_1, r_2, b_1, b_2 > 0$.

There are various rescalings that we can do to simplify this equation. An obvious one is to remove the carrying capacities by defining $N'_i = N_i/K_i$. We do this, then rename $N'_i \rightarrow N_i$. There are further rescalings that we can do by absorbing one of the reproduction rates r_i into time, but they destroy the symmetry of the equations so we choose not to. We then have

$$\begin{aligned}\frac{dN_1}{dt} &= r_1 N_1 (1 - N_1 - b_1 N_2) \\ \frac{dN_2}{dt} &= r_2 N_2 (1 - N_2 - b_2 N_1) .\end{aligned}\tag{1.127}$$

If we set either $N_1 = 0$ or $N_2 = 0$, then the dynamics of the other is described by the familiar logistic equation that we discussed in Section 1.1. Here, we would like to understand how the two species interact with each other.

As we'll see, this model is rather straightforward, certainly compared to some of the predator-prey generalisations that we discussed above. To start, we can look at the fixed points. There are four. One is $(N_1, N_2) = 0$ and is devoid of life. Two others have one of the species extinct,

$$(N_1, N_2) = (1, 0) , \quad (N_1, N_2) = (0, 1)\tag{1.128}$$

while the third equilibrium point has the two species coexisting:

$$(N_1, N_2) = \frac{1}{1 - b_1 b_2} (1 - b_1, 1 - b_2) .\tag{1.129}$$

Populations have to be positive, so the last of these is a viable fixed point if either both $b_1, b_2 < 1$ or if both $b_1, b_2 > 1$. But if one of these coefficients is greater than one and the other less than one, then there is no coexistence.

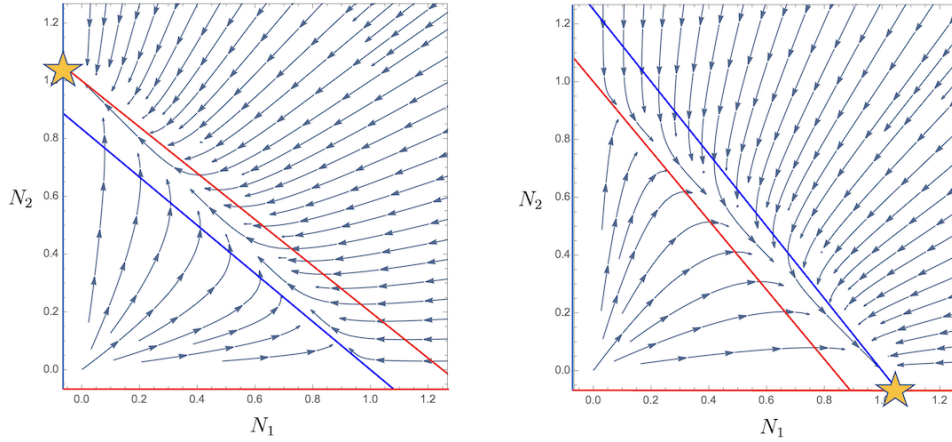


Figure 18. The phase plane flows for two cases $b_2 < 1 < b_1$ (on the left) and $b_1 < 1 < b_2$ (on the right) with the nullclines also shown. The fixed points are where a red and blue nullcline meet. The gold star denotes the stable fixed point.

What about stability? The Jacobian matrix takes the form

$$J = \begin{pmatrix} r_1(1 - 2N_1 - b_1N_2) & -r_1b_1N_1 \\ -r_2b_2N_2 & r_2(1 - 2N_2 - b_2N_1) \end{pmatrix}. \quad (1.130)$$

We need to evaluate this on each of the four fixed points (assuming that all four exist) and compute the eigenvalues λ . This is a straightforward exercise. We have:

$$(N_1, N_2) = (0, 0) \implies \lambda = r_1, r_2. \quad (1.131)$$

So this fixed point is unstable. The next two fixed points we get

$$\begin{aligned} (N_1, N_2) = (1, 0) &\implies \lambda = -r_1, r_2(1 - b_2) \\ (N_1, N_2) = (0, 1) &\implies \lambda = -r_2, r_1(1 - b_1). \end{aligned} \quad (1.132)$$

Here we see a more interesting pattern: the first fixed point is stable if $b_2 > 1$ and the second is stable if $b_1 > 1$. If either of these coefficients is less than one, then the respective fixed point becomes a saddle. The flow in phase plane for two the cases $b_1 < 1 < b_2$ and $b_2 < 1 < b_1$ are shown in Figure 18.

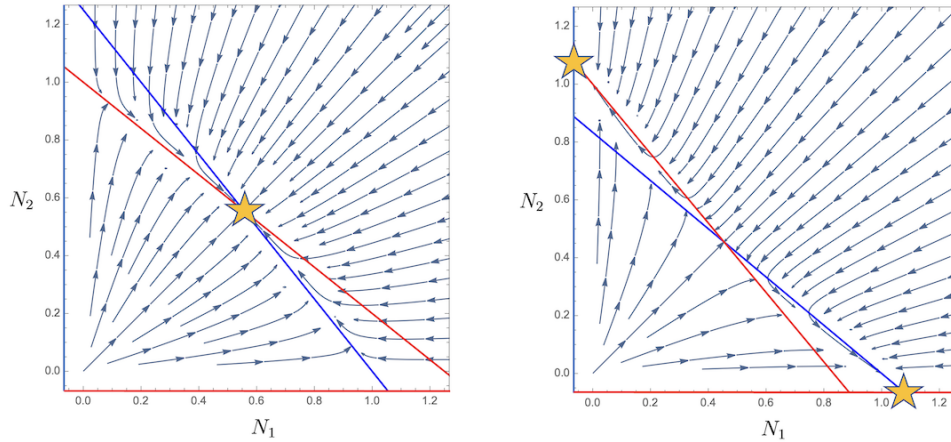


Figure 19. The phase plane flows for two cases $b_1, b_2 < 1$ (on the left) and $b_1, b_2 > 1$ (on the right) with the nullclines also shown. The fixed points are where a red and blue nullcline meet and the stable fixed point(s) denoted by a gold star.

Finally, for the fixed point corresponding to coexistence, the eigenvalues are more complicated. A little algebra gives

$$\begin{aligned}
 (N_1, N_2) &= \frac{1}{1 - b_1 b_2} (1 - b_1, 1 - b_2) \\
 \implies \lambda &= \frac{1}{2(1 - b_1 b_2)} \left[r_1(b_1 - 1) + r_2(b_2 - 1) \right. \\
 &\quad \left. \pm \sqrt{(r_1(b_1 - 1) + r_2(b_2 - 1))^2 + 4(b_1 - 1)(b_2 - 1)(b_1 b_2 - 1)} \right].
 \end{aligned} \tag{1.133}$$

You can check that the number under the square root is always positive. Moreover, both eigenvalues are negative provided that $b_1, b_2 < 1$, while one is negative and the other positive if $b_1, b_2 > 1$.

The flows in the phase plane are plotted for these two cases in Figure 19. We can now piece the full story together. The coefficient b_1 is a measure of the disruption that the second species has on the first. Similarly, b_2 measures the disruption of the first species on the second. If $b_1, b_2 < 1$, then there is a stable coexistence fixed point and the two species share the resources, with the population of each lower due to the presence of the other species.

Meanwhile if, say, $b_1 > 1$ and $b_2 < 1$ then this means that the second species is the more disruptive and it's game over for the first: the fixed point has $N_1 = 0$.

Finally, if both $b_1 > 1$ and $b_2 > 1$ then there are stable fixed points with one or the other species extinct. Which fixed point you end up in depends on the initial conditions. Increasing, say, b_1 increases the basin of attraction for the second species to become the winner.

There is a well known dictum in ecology known as the *principle of competitive exclusion*. It says that two species which compete for the same limited resources cannot coexist as one will have an advantage and will ultimately win out. This is indeed what the simple model above predicts whenever $b_1 > 1$ or $b_2 > 1$. But, as we've seen, the two species can certainly live in happy coexistence in this model when $b_1, b_2 < 1$. I'm not an ecologist but it seems strange to me to take something that is *not* predicted by equations and then elevate it to a "principle".


Be Kind

It may be that our two species get along in some symbiotic way. In this case, we can continue to describe their interactions through the equations (1.127), but we take $b_1, b_2 < 0$.

As before, there are four fixed points: the trivial one, the two with one species extinct (1.128), and the coexisting fixed point (1.129),

$$(N_1, N_2) = \frac{1}{1 - b_1 b_2} (1 - b_1, 1 - b_2) .$$

We see that this fixed point only exists if $b_1 b_2 < 1$. This fixed point is stable, and the resulting flows are shown in the figure. This time, the presence of each species enhances the population of the other.



As $b_1 b_2 \rightarrow 1$ the two species are *too* beneficial, and the populations run away to infinity! Indeed, for any $b_1 b_2 > 1$, the population runs off to infinity.

1.4.5 Dengue Fever

In this section we will develop a slightly different competition model, this one designed to describe attempts to mitigate a particularly unpleasant disease.

Dengue fever is transmitted by mosquitos in the tropics. It's pretty unpleasant and you don't want to catch it. Happily, help is at hand in the form of a bacteria known as *Wolbachia*. When mosquitos are infected with this bacteria, it blocks transmission of the dengue virus. This brings hope that by introducing *Wolbachia*-infected mosquitos into the wild, they may help reduce or eradicate the dengue virus.

We will assume that all mosquitos carry dengue, but those infected with *Wolbachia* are harmless. At this point there's something of a story to tell. Mosquitos can only pass *Wolbachia* onto their offspring. (Just kissing is not enough.) And we have the following complications:

- If a female is infected, all her eggs will be infected.
- An infected female lays fewer eggs than the uninfected.
- An infected female can mate with an infected male. But if an uninfected female mates with an infected male, there will be no viable eggs.
- Infected mosquitos don't live as long.

Our task is to translate these facts into equations. Indeed, the art of mathematical modelling is constructing equations from words. We introduce the following variables

$$\begin{aligned}x &= \text{number of uninfected females} \\y &= \text{number of infected females} .\end{aligned}\tag{1.134}$$

We assume that the number of (un)infected males simply tracks the female population. Then we model the facts above by the following equations,

$$\begin{aligned}\frac{dx}{dt} &= x \left(r \frac{x}{x+y} - d - \epsilon(x+y) \right) \\ \frac{dy}{dt} &= y(\lambda r - \mu d - \epsilon(x+y)) .\end{aligned}\tag{1.135}$$

If we set either $x = 0$ or $y = 0$ then the other equation reduces to the logistic equation. The various terms in the equations have the following interpretation:

- The proportion of uninfected males is $x/(x+y)$. An uninfected female mating can only mate with an uninfected male and this increases the uninfected proportional to r .
- An infected female can mate with any male. This increases the rate of infected mosquitos by λr , with $\lambda < 1$ because infected females lay fewer eggs.
- Uninfected mosquitos die off at a rate proportional to d . Infected mosquitos die off at the quicker rate of μd with $\mu > 1$.
- There is a logistic-like competition for resources, giving rise to the $\epsilon(x+y)$ terms. This is the same for infected and uninfected.

With these equations in hand, we can now rescale to remove unnecessary constants, with $t \rightarrow rt$ and $x \rightarrow \epsilon x/r$ and $y \rightarrow \epsilon y/r$. We're left with the pair of equations,

$$\frac{dx}{dt} = x \left(\frac{x}{x+y} - \frac{d}{r} - (x+y) \right) \quad \text{and} \quad \frac{dy}{dt} = y \left(\lambda - \frac{\mu d}{r} - (x+y) \right) . \quad (1.136)$$

It's useful to collect the remaining constants together so we write

$$\frac{dx}{dt} = x \left(x_0 - \frac{y}{x+y} - (x+y) \right) \quad \text{and} \quad \frac{dy}{dt} = y (y_0 - (x+y)) . \quad (1.137)$$

Here

$$x_0 = 1 - \frac{d}{r} \quad \text{and} \quad y_0 = \lambda - \frac{\mu d}{r} . \quad (1.138)$$

(Note that the extra 1 in x_0 is what turns the $x/(x+y)$ in (1.136) into $y/(x+y)$ in (1.137)). The uninfected mosquitos breed more than they die (sadly), so $r > d$ and, correspondingly, $0 < x_0 < 1$. There is no a priori bound on y_0 but we will assume that $y_0 > 0$ so that populations of infected mosquitos also grow otherwise our goal of eradicating dengue is hopeless. This means that we have $0 < y_0 < x_0 < 1$.

Now we're in business. There are four fixed points: the trivial one $(x, y) = (0, 0)$, two where one of the populations is extinct, $(x, y) = (x_0, 0)$ and $(x, y) = (0, y_0)$ and one where both infected and uninfected coexist (x^*, y^*) , with

$$x^* = y_0(1 - x_0 + y_0) \quad \text{and} \quad y^* = y_0(x_0 - y_0) . \quad (1.139)$$

You can check that this is a saddle.





Figure 20. Flows in the phase plane have two stable fixed points shown with gold stars. These correspond to all mosquitos infected (on the y -axis, or all mosquitos uninfected (on the x -axis). Where you end up depends on the initial conditions. The basis of attraction to have uninfected mosquitos is shown shaded. This is plotted with $x_0 = 0.5$ and $y_0 = 0.2$.

A plot of the flows in the phase plane, together with the nullclines, are shown in Figure 20. The two stable fixed points have either all mosquitos infected, or all uninfected. If you want to end up solely with infected mosquitos, then you need to introduce a sufficient number so that you sit in the basin of attraction of the fixed point $(x, y) = (0, y_0)$. This is the unshaded region in Figure 20.

Suppose that your goal is to eradicate dengue fever. You might wonder if it's possible to change the parameters in our equations to work to our benefit. We see that the end result ultimately depends on just two parameters: x_0 and y_0 . There's not much that we can do about x_0 , which depends on the birth and death rates of the uninfected mosquitos. In contrast, as shown in (1.138), y_0 depends on λ , which is the drop in egg production for infected mosquitos, and on μ which captures the reduced lifespan of infected mosquitos, and you might optimistically hope that these can be changed by altering the strain of Wolbachia. We could either try to increase λ , or to decrease μ .



From (1.138), we see that μ multiplies $d/r \ll 1$, suggesting that it might be more profitable to attempt to increase λ so that the number of eggs produced by infected mosquitos is closer to those produced by the uninfected. Either way, if we increase y_0 then the nullcline moves. An example is shown in the right, where the phase plane flow is plotted for $x_0 = 0.5$ and $y_0 = 0.4$ (as opposed to $y_0 = 0.2$ in Figure 20). Again the basin of attraction for uninfected mosquitos is shaded and is visibly smaller than that shown in Figure 20.

1.4.6 The Large Diversity Limit

Many systems in nature involve more than two interacting species. At the extreme end, the human gut contains somewhere close to 1000 different kinds of microbes. Some marine ecosystems contain a similar number of species. Can we fruitfully model such complicated systems?

To start, we can proceed as before. If we have M different species, each with population $N_i(t)$ with $i = 1, \dots, M$, then we can describe the dynamics by the set of equations

$$\frac{dN_i}{dt} = f_i(N) . \tag{1.140}$$

In general, this dynamics is likely to be complicated. But there's one question that has a rather nice answer: is it possible for multiple species to coexist?

Rephrasing this, is it possible to find a stable fixed point $f(N^*) = 0$ when the number of interacting species is very large, so $M \gg 1$? Expanding about the fixed point, the dynamics is


$$\frac{dN_i}{dt} = J_{ij}(N_j - N_j^*) \quad \text{with} \quad J_{ij} = \frac{\partial f_i(N^*)}{\partial N_j} . \tag{1.141}$$

This fixed point is stable if all the eigenvalues of J_{ij} have negative real part.

Now, it is of course always *possible* to find functions f_i such that all the eigenvalues of J_{ij} all have negative real part. But is it likely?

There is an important branch of mathematics, known as *random matrix theory*, that allows us to answer questions of this kind. Suppose that you have a matrix where each element is chosen randomly, and independently, from some probability distribution. We then ask: what is the probability distribution of the eigenvalues? Rather wonderfully, it turns out that many features of the eigenvalue distribution are independent of the original choice of probability distribution that you choose for the individual elements.

We won't derive the key properties of random matrices here, but instead just state the key result. Suppose that you sample the individual elements of the matrix from a Gaussian distribution with zero mean and variance σ^2 . Then, in the limit $M \rightarrow \infty$, the eigenvalues λ are uniformly distributed on the complex plane in a disc of radius $r = \sqrt{M\sigma^2}$, as shown in the figure.



How can we use this for our question of population stability? Here's a simplified model. We will take the Jacobian matrix in (1.141) to have the form

$$J = -\mathbf{1} + J_{\text{random}} . \tag{1.142}$$

That is: in the absence of interactions, the fixed point is stable and all eigenvalues have been scaled to $\lambda = -1$. But we then add to this random interactions with other species, captured by the additional term J_{random} . Invoking the result from random matrix theory, we see that, in the $M \rightarrow \infty$ limit, the eigenvalue with the largest real part has

$$\text{Re}(\lambda_{\text{max}}) = -1 + \sqrt{M\sigma^2} . \tag{1.143}$$

In other words, a stable ecosystem where the interactions between different species have variance σ^2 can support at most

$$M_{\text{max}} = \frac{1}{\sigma^2} \tag{1.144}$$

different species. This is known as *May's stability criterion*.

The model above involves many simplifying assumptions but, nonetheless, the idea that there is a trade-off between diversity and stability has been influential among ecologists.

1.5 Epidemiology

2020 was a weird year. For many of us, it's a blur of Zoom calls and government mandated daily walks and lots of talk about the number R_0 and why it's too big. The purpose of this section is to re-live this experience, this time with equations.

1.5.1 The SIR Model

The classic epidemic model is named after its three variables,

S = number of people who are susceptible to the disease

I = number of people who are infected

R = number who have recovered or are dead.

One of the lessons that came out of the 2020 pandemic is that there is a tension between the goals of scientists and those of politicians. This model highlights the tension pretty clearly. If you're a mathematical modeller then, at least at this basic level, it doesn't matter if someone recovers or dies from the disease: either way, they no longer contribute to its spread. But I've been told that there are some politicians who appreciate the distinction between these two outcomes. If you want to phrase things more delicately, you could say R = "removed".

The equations that model the spread of the disease are:

$$\frac{dS}{dt} = -\beta IS, \quad \frac{dI}{dt} = +\beta IS - \nu I, \quad \frac{dR}{dt} = \nu I \quad (1.145)$$

with positive constants $\beta, \nu > 0$.

These equations have an intuitive underpinning. The susceptible turn into infected at a rate that is proportional to IS . The idea is that this product captures the interaction between the two groups, in way that is analogous to the NP terms that we met in the predator-prey equations. Meanwhile, the infected either recover or die at a constant rate. Whichever path they take, they are removed from the process and no longer contribute to the dynamics.

The first thing to note is that the total number of people $N = S + I + R$ is a constant. That's tautologically true if you still count people who die! Moreover, we don't really care about the recovered/dead at all as they don't feed back into the other two. That means that we can focus just on the first two equations in (1.145) and then reconstruct $R(t) = N - S(t) - I(t)$.

The first question to ask is: are we in trouble? Is the number of infected people going to increase? From the second equation in (1.145), we see that the answer doesn't depend on how many are already infected, as long as $I \neq 0$: it depends only on how many susceptible there are:

$$\dot{I}(0) > 0 \iff \beta S(0) > \nu . \quad (1.146)$$

Here ν is the rate at which individuals recover, and β is the rate at which each susceptible is infected. The epidemic starts if the rate at which total population is infected is faster than the rate at which they recover.

Typically, a disease will start with just a few infected among a large population and we can take $S(0) \approx N$. This motivates us to define the *reproductive ratio*, pronounced “R-naught” or “R-zero”,

$$R_0 = \frac{\beta N}{\nu} . \quad (1.147)$$

Here $1/\nu$ is the transmission period; β is the transmission rate per contact; and N is the initial number of contacts. The epidemic grows if $R_0 > 1$. For the original strain of COVID-19, R_0 was somewhere between 2 and 5. For polio, $R_0 \approx 4 - 6$; for mumps $R_0 \approx 10 - 12$. For measles, $R_0 \approx 16 - 18$.

Suppose that we are well prepared and vaccinate a fraction p of the population before the disease hits. Then the number of susceptibles is reduced to $(1 - p)N$, lowering the effective reproductive ratio to $(1 - p)R_0$. Or, said differently, if the unvaccinated population has a given R_0 , then we need to vaccinate a fraction $p > (R_0 - 1)/R_0$ to stop the disease spreading. Crucially, we don't need to vaccinate everyone, just enough to reach herd immunity which protects the whole community.

This is a good point to pause and comment on a more philosophical aspect of mathematical modelling. All the mathematical models that we write down in this course are, to put it bluntly, wrong. They are all, at best, caricatures of the underlying reality. Most likely they omit many important details. All of which begs the question: why should we trust them? This is particularly important when it comes to putting in place mechanisms to counter disease. One answer to this question is that you can search for things in the models that are robust: things that don't change when you add extra layers to the model and include further details. Crucially, it turns out that in epidemiology, one such robust quantity is the threshold for vaccination that we have computed above. And, indeed, it works well in practice.




Figure 21. On the left, flows of the SIR model in the phase plane. On the right, a typical time evolution of the susceptible and infected population.

We can solve for the trajectories in the phase plane in much the same way as the Lotka-Volterra model. We view $I = I(S)$ and, dividing the first two equations in (1.145), we're left with

$$\frac{dI}{dS} = \frac{\nu}{\beta S} - 1 = \frac{N}{R_0 S} - 1. \quad (1.148)$$

This is easily integrated to give

$$I(S) = \frac{N}{R_0} \log S - S + c \quad (1.149)$$

with c the integration constant. Curves for different choices of c are plotted on the left of Figure 21, adorned with arrows which show that the number of susceptibles are always decreasing. The curves are plotted in the triangular region to reflect the fact that $S + R < N$, the total population. An epidemic that starts with just a few infected is described by the blue curve which intersects the corner of the triangle.

To get the time data, you need to solve the equations numerically. A typical example is shown on the right of Figure 21. The number of infected grows and then shrinks to zero, while the number of susceptibles falls to a constant value S_∞ .

From both graphs, we see that the number of infected reaches a maximum value. From (1.148), we can easily see that the maximum number of infected occurs when $S = N/R_0$. An intuitive way to think about this comes from defining the *effective reproductive ratio* $R_{\text{eff}} = \beta S/\nu$. The epidemic starts to recede when $R_{\text{eff}} = 1$.




Figure 22. The lucky ones, as a function of R_0 .

There's one piece of good news hiding in this analysis: we're not all going to die. Or even become infected. The epidemic burns itself out and some number S_∞ never catch the disease. This occurs when $I(S) = 0$. From the left plot of Figure 21, we see that there are two values of S when this happens: we call these S_0 and S_∞ . Setting $I(S) = 0$ in (1.149), we have

$$\frac{N}{R_0} \log S_0 - S_0 = \frac{N}{R_0} \log S_\infty - S_\infty . \quad (1.150)$$

It's natural to assume that we start with very few infected and lots of susceptibles, so $I_0 \ll N$ and $S_0 \approx N$. We will denote the lucky fraction of the population as σ , so that

$$S_\infty = \sigma N . \quad (1.151)$$

From (1.150), this lucky fraction satisfies

$$\log \sigma = R_0(\sigma - 1) . \quad (1.152)$$

The result is plotted in Figure 22. For $R_0 < 1$ there's no need to panic: essentially 100% of the population escapes unscathed. But we can see that things get more worrisome as R_0 increases. In particular, for R_0 large we have $\sigma \approx e^{-R_0}$. So by the time you get to $R_0 \approx 10$, that's more or less everyone that's going to get infected.

1.5.2 Just When You Thought It Was Safe...

I caught covid last month. So it's certainly not gone away. Of course, what happened is that my immunity wore off and I moved from the recovered R group back into the susceptible S group.

It's straightforward to incorporate this, or other variations, into the SIR model. Here we consider a different feature: people die (of something other than the disease in question); and people are born. And all of those new born babies sit straight in the susceptible camp. To capture this, we add various terms to our original SIR equations (1.145), which now becomes

$$\begin{aligned}\frac{dS}{dt} &= -\beta IS + bN - \mu S \\ \frac{dI}{dt} &= \beta IS - \nu I - \mu I \\ \frac{dR}{dt} &= \nu I - \mu R.\end{aligned}\tag{1.153}$$

Here $N = S + I + R$ as before; b is the birth rate, and μ the death rate. We've assumed that this disease doesn't kill anyone, an assumption that manifests itself in the fact that members of all groups contribute to the births. Moreover, we assume that the death rate is constant regardless of whether you've had the disease or not. We see that $dN/dt = (b - \mu)N$ like the Malthus model of Section 1.1.

To make things simple, let's assume (perhaps unrealistically) that the population is stable, with $b = \mu$ so that N is constant. Our reproductive ratio again comes from looking at the \dot{I} equation when $S = N$; it is

$$R_0 = \frac{\beta N}{\nu + \mu}.\tag{1.154}$$

At first glance, it looks like that's good news: dying reduces R_0 ! But some thought suggests that the effect is minimal: for any disease in which you recover over a period of time that's much shorter than your lifetime, we necessarily have $\nu \gg \mu$, and so $R_0 \approx \beta N/\nu$ as before.

Since N is constant, the first two equations again decouple from the third. This time, the influx of births means that there is a fixed point (S^*, I^*) with

$$S^* = \frac{\nu + \mu}{\beta} = \frac{N}{R_0} \quad \text{and} \quad I^* = \mu \frac{N - S^*}{\beta S^*} = \frac{\mu}{\beta}(R_0 - 1).\tag{1.155}$$

By now, we know what we're doing. We determine the stability by computing the Jacobian at the fixed point

$$J = \begin{pmatrix} -\mu R_0 & -(\nu + \mu) \\ \mu(R_0 - 1) & 0 \end{pmatrix}.\tag{1.156}$$




Figure 23. On the left, flows of the modified SIR model in the phase plane. On the right, a typical time evolution of the susceptible and infected population.

The eigenvalues are complex and given by

$$\lambda = -\frac{\mu R_0}{2} \pm \frac{1}{2} \sqrt{\mu^2 (R_0 - 2)^2 - 4\mu\nu(R_0 - 1)}. \quad (1.157)$$

If we're in the realistic situation where recovery is not measured in a timescale less than decades then $\mu \ll \nu$ and we have

$$\lambda \approx -\frac{\mu R_0}{2} \pm i\omega \quad \text{with} \quad \omega = \sqrt{\mu\nu(R_0 - 1)}. \quad (1.158)$$

These eigenvalues tell us that we have a stable fixed point, which trajectories spiral towards with a period. The resulting dynamics is plotted, both in the phase plane, and as a function of time, in Figure 23

This suggests that we will observe transient oscillations on a time period that is the geometric mean of the lifetime and infection recovery time.

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{\sqrt{\mu\nu(R_0 - 1)}} \approx \frac{1}{\sqrt{\mu\nu}}. \quad (1.159)$$

For example, measles has $R_0 \approx 20$ with a recovery rate of about 12 days. So putting $1/\nu \approx 12$ days and setting a human lifetime at $1/\mu \approx 70$ years, we expect to see oscillations with a period of $T \approx 2.2$ years. In fact, there's an extra complication here, namely schools. There is a delta-function injection of a susceptible population every September when kids start school. This is sufficient to keep the measles outbreaks occurring, but with an oscillation that is roughly 2 years rather than 1 year. Data of measles outbreaks in the UK over the past decade is shown in Figure 24.




Figure 24. Data from the UK government on measles outbreaks over the past decade.

1.6 Chemical Reactions

The kinds of dynamical system that we’ve been developing so far have many applications that removed from population dynamics. Here we give an example in which we treat the population of chemicals in reactions.

1.6.1 The Law of Mass Action

The simplest chemical reaction is that two chemicals, A and B , form a product C ,



Here k is the rate constant of the reaction. It’s straightforward to model this in terms of equations. We let A , B and C denote the concentration of the chemical and write the equations

$$\frac{dA}{dt} = -kAB , \quad \frac{dB}{dt} = -kAB , \quad \frac{dC}{dt} = +kAB . \quad (1.161)$$

The fact that the time derivatives are proportional to the product of the concentrations $A \times B$ is known as the *law of mass action* in chemistry. It’s the same conceptual idea that we’ve invoked in both the predator-prey models and the epidemiology models but without giving it a name. However, the law of mass action is on much firmer footing in the context of chemical reactions where it can be derived using the kind of technology that we introduced in the lectures on [Statistical Physics](#). In contrast, in population dynamics and epidemiology, the idea that we get the right phenomenology by multiplying together the interacting variables is an assumption.

It's straightforward to solve these equations. First we note that $A + C = A_0$ is constant, as is $B + C = B_0$. We can then use this to write the last equation in (1.161) purely in terms of C ,

$$\frac{dC}{dt} = k(A_0 - C)(B_0 - C) . \quad (1.162)$$

This has the solution

$$C(t) = A_0 B_0 \frac{1 - e^{(A_0 - B_0)kt}}{B_0 - A_0 e^{(A_0 - B_0)kt}} \quad (1.163)$$

where we've chosen the integration constant so that $C(0) = 0$. This expression has the nice property that $\lim_{t \rightarrow \infty} C(t) = \min(A_0, B_0)$, reflecting the fact that the reaction stops whenever one of the chemicals is exhausted.

What is done can sometimes be undone. It may be that the reaction can also go in the opposite direction, albeit at a different rate:



Chemists often use the notation k_{-1} for the inverse reaction rate k_2 . In this case, the third equation in (1.161) becomes

$$\frac{dC}{dt} = k_1 AB - k_2 C \quad (1.165)$$

with k_1 and k_2 the two reactions rates. We still have $A + C = A_0$ constant and $B + C = B_0$ constant, so this equation becomes

$$\frac{dC}{dt} = k_1(A_0 - C)(B_0 - C) - k_2 C \quad (1.166)$$

Now there is a fixed point, when the right-hand side vanishes. We define

$$\kappa = \frac{k_2}{k_1} . \quad (1.167)$$


From (1.165), we see that, despite the similarity in their names, k_2 and k_1 actually have different dimensions, so that κ has the dimensions of concentration. Indeed, the fixed point obeys the quadratic

$$C^2 - (A_0 + B_0 + \kappa)C + A_0 B_0 = 0 \quad (1.168)$$

together with the additional requirement that $0 < C < \min(A_0, B_0)$.

Lotka-Volterra Revisited

Here's an interesting class of reactions: we have four different chemicals, A , B , C and D which react as



We take the reaction rates to be k_1 , k_2 , and k_3 respectively. The dynamics is then given by

$$\begin{aligned} \frac{dA}{dt} &= -k_1 AB \\ \frac{dB}{dt} &= +k_1 AB - k_2 BC \end{aligned} \quad (1.170)$$

$$\begin{aligned} \frac{dC}{dt} &= +k_2 BC - k_3 CD \\ \frac{dD}{dt} &= +k_3 CD . \end{aligned} \quad (1.171)$$

Suppose that we now intervene and ensure that there is a constant concentration of the chemicals A and D . Then we can replace these variables with these constant values in (1.170) and (1.171). These two equations then reduce to the Lotka-Volterra equations that we studied in detail in Section 1.4 where we saw that the concentration of the chemicals B and C will oscillate over time. Indeed, this was the context in which Alfred Lotka first wrote down these equations.

Stoichiometry

Consider a reaction in which m molecules of type A combine with n molecules of type B to produce a molecule of type C . If the reaction is reversible, then we write



This is modelled by the equations

$$\begin{aligned} \frac{dA}{dt} &= -mk_1 A^m B^n + mk_2 C \\ \frac{dB}{dt} &= -nk_1 A^m B^n + nk_2 C \\ \frac{dC}{dt} &= k_1 A^m B^n - k_2 C . \end{aligned} \quad (1.173)$$

Here the powers A^m and B^n are a consequence of the law of mass action. Meanwhile, the pre-factors of m and n are designed to ensure that $A + mC$ and $B + nC$ are both constant, as they should be.

1.6.2 Michaelis-Menten Enzyme Kinetics

Some biochemical reactions can be yawningly slow. They get a helping hand from enzymes. These are proteins which give the chemicals a small hug, increasing the reaction rate.

The original set of chemicals are called the *substrate* S . The end result is called the *product* P . The reaction is helped on its way by a population of enzymes E . When the enzymes bind with the substrate, they form a combination known as a *complex* C . The reaction takes place through the *Michaelis-Menten reaction*:



Note that the first reaction, with the enzymes binding to the substrate, can go both ways. But the second, where the enzyme releases the final product, goes only in one direction. The equations governing this reaction are

$$\begin{aligned} \frac{dS}{dt} &= -k_1SE + k_2C \\ \frac{dE}{dt} &= -k_1SE + (k_2 + k_3)C \\ \frac{dC}{dt} &= +k_1SE - (k_2 + k_3)C \\ \frac{dP}{dt} &= +k_3C . \end{aligned} \quad (1.175)$$

We will assume that $S(0) = S_0$ and $E(0) = E_0$, while $C(0) = P(0) = 0$ as initial conditions. We usually assume that $S_0 \gg E_0$, meaning that the limit on the reaction rate is set by the number of enzymes rather than the lack of substrate.

There are two conservation laws within these equations. These are

$$E + C = E_0 \quad \text{and} \quad S + C + P = S_0 . \quad (1.176)$$

The first reflects the fact that the enzymes aren't used up: they are either empty and denoted as E or full and denoted as C . The second tells us that the end result is to change substrate into product without losing any.

We can use the first of these to eliminate the enzyme concentration E , and focus on equations just for S and C . (Note that the product P is just a dumping ground and doesn't affect the other variables.) We have

$$\begin{aligned} \frac{dS}{dt} &= -k_1E_0S + (k_1S + k_2)C \\ \frac{dC}{dt} &= k_1E_0S - (k_1S + k_2 + k_3)C . \end{aligned} \quad (1.177)$$




Figure 25. The flows in the (s, c) phase plane, plotted with $\epsilon = 0.1$. We have also taken $\mu = 1.5$ and $\lambda = 0.5$.

We can clean these up a little by rescaling variables. We write $s = S/S_0$ and $c = C/E_0$. This is the proportion of substrate and the proportion of occupied enzymes respectively, with $s(0) = 1$ and $c(0) = 0$. If we also rescale time by defining $\tau = k_1 E_0 t$ then the equations become

$$\begin{aligned} \frac{ds}{d\tau} &= -s + (s + \mu - \lambda)c \\ \frac{dc}{d\tau} &= \frac{1}{\epsilon} (s - (s + \mu)c) \end{aligned} \quad (1.178)$$

where we've now got just three constants

$$\lambda = \frac{k_3}{k_1 S_0} \quad \text{and} \quad \mu = \frac{k_2 + k_3}{k_1 S_0} \quad \text{and} \quad \epsilon = \frac{E_0}{S_0}. \quad (1.179)$$

It's straightforward to solve these equations numerically and the result is shown in Figure 25 where, crucially, we have taken $\epsilon \ll 1$. The result is rather striking: all trajectories head immediately to a common curve, which they then follow down to the origin.

The key to understanding this feature is to appreciate that, with $\epsilon \ll 1$, the enzyme dynamics $c(t)$ happens on a much faster time scale than the substrate dynamics $s(t)$. The system therefore relaxes quickly to the $\dot{c} = 0$ curve, given by

$$c = \frac{s}{s + \mu}. \quad (1.180)$$

That's the curve shown in red in Figure 25. There is then a much slower precession along this curve which is captured by substituting (1.180) into the equation for \dot{s} ,

$$\frac{ds}{d\tau} = -\frac{\lambda s}{s + \mu} . \quad (1.181)$$

This tells us how the substrate is depleted towards $s = 0$. As the amount of substrate changes, so too does the concentration of complexes, quickly adapting to remain in the equilibrium given by (1.180).

Biochemists are often interested in the reaction velocity R , given by

$$R = \frac{dP}{dt} = k_2 C = k_1 k_2 E_0 \frac{S}{k_1 S + k_2 + k_3} . \quad (1.182)$$

This is the *Michaelis-Menten equation*, often written in terms of the maximum reaction velocity $V = k_2 E_0$ and the so called *Michaelis-Menten constant* $\kappa = (k_2 + k_3)/k_1$. You can extract these constants by plotting the linear graph of R^{-1} against S^{-1} .

1.6.3 Asymptotic Behaviour

It's possible to make analytic progress by solving the equation perturbatively in the small parameter ϵ . However, naively setting $\epsilon = 0$ changes the character of the differential equation so we need to tread a little carefully.

Let's start by being naive. If we set $\epsilon = 0$ in (1.178) then we get the differential equation (1.181). We'll call the function that solves this $\tilde{s}(t)$: it is easy to check that this is given implicitly by

$$\tilde{s} + \mu \log \tilde{s} = -\lambda\tau + \text{constant} . \quad (1.183)$$

We can fix the constant by requiring the initial condition $\tilde{s}(\tau = 0) = 1$, which tells us that

$$\tilde{s} + \mu \log \tilde{s} = -\lambda\tau + 1 . \quad (1.184)$$

This should be viewed as an approximation to the late time behaviour of the system: it holds after the fast dynamics has happened and we can impose the constraint (1.180).

We can also get an approximation for the short time behaviour. To do this, we rescale the time coordinate and write

$$\tau = \epsilon T . \quad (1.185)$$

The two equations of motion (1.178) become

$$\frac{ds}{dT} = \epsilon(-s + (s + \mu - \lambda)c) \quad \text{and} \quad \frac{dc}{dT} = s - (s + \mu)c. \quad (1.186)$$

Now the ϵ is sitting on the right-hand side of the equation, rather than the left-hand side, and we can do a standard perturbative expansion, writing

$$s(T) = s_0(T) + \epsilon s_1(T) + \dots \quad \text{and} \quad c(T) = c_0(T) + \epsilon c_1(T) + \dots. \quad (1.187)$$

This is a short time expansion. To leading order in ϵ , we have

$$\frac{ds_0}{dT} = 0 \quad \text{and} \quad \frac{dc_0}{dT} = s_0(T) - (s_0(T) + \mu)c_0(T) \quad (1.188)$$

subject to the initial conditions $s(0) = 1$ and $c(0) = 0$. The first of these equations tells us that $s_0(T) = 1$, while the second is then solved by

$$c_0(T) = \frac{1}{1 + \mu} [1 - e^{-(1+\mu)T}] \quad \implies \quad c_0(\tau) = \frac{1}{1 + \mu} [1 - e^{-(1+\mu)\tau/\epsilon}]. \quad (1.189)$$

We see that $c_0(\tau)$ has the characteristic $e^{-\tau/\epsilon}$ behaviour at short times. Famously, this kind of function vanishes faster than any polynomial as $\epsilon \rightarrow 0$.

For both $s(\tau)$ and $c(\tau)$, we now have two approximations to the solutions. The short time solutions are $s_0(\tau)$ (which, admittedly, is trivially constant) and $c_0(\tau)$. The long time solution, after the fast dynamics is exhausted, is $\tilde{s}(\tau)$ given implicitly by (1.184) and, from (1.180), $\tilde{c}(\tau) = \tilde{s}(\tau)/(\tilde{s}(\tau) + \mu)$. In the world of matched asymptotic expansions, $s_0(\tau)$ and $c_0(\tau)$ are called *inner solutions* while $\tilde{s}(\tau)$ and $\tilde{c}(\tau)$ are called *outer solutions*. Our next task is to patch these together.

This is straightforward for $s(\tau)$ because the short time solution is just a constant, $s_0(\tau) = 1$ and, indeed, this matches the short time behaviour $\lim_{\tau \rightarrow 0} \tilde{s}(\tau) = 1$. You can check that $\tilde{s}(\tau)$ is a good approximation to the true solution when $\epsilon \ll 1$.

For $c(\tau)$, we have to be a little more careful. We have

$$\lim_{\tau \rightarrow \infty} c_0(\tau) = \lim_{\tau \rightarrow 0} \tilde{c}(\tau) = \frac{1}{1 + \mu}. \quad (1.190)$$

This is telling us that they will patch nicely together. We do this by adding the two solutions and subtracting their common piece, a process known as the *uniform approximation*,

$$c_{\text{uniform}}(\tau) = c_0(\tau) + \tilde{c}(\tau) - \frac{1}{1 + \mu}. \quad (1.191)$$

This function is still known only implicitly because $\tilde{s}(\tau)$ is given by the relation (1.184). Nonetheless, you can check that $c_{\text{uniform}}(\tau)$ gives good agreement to the numerical solution for $c(\tau)$ when $\epsilon \ll 1$.

1.7 Neuron Excitations

One of the great advances of mathematical biophysics was a system of equations, first written down in 1952 by Hodgkin and Huxley, to describe the way that neurons fire in the brain. Neurons have a long sticky-out bit called an axon, which conducts pulses of electricity known as *action potentials* and these mediate signalling from one neuron to the next. These action potentials form from some external chemical prompt, but in a way such that a reasonably small prompt is amplified to an unambiguous pulse of electricity. The question is: how does this happen?

Rather than looking at the propagation of the signal along the axon, Huxley and Hodgkin instead considered the easier situation in which the signal is constant in space and focussed on how the signal fires in time. They wrote down a rather complicated set of equations, in which there are three kinds of electric current, carried by potassium ions, sodium ions, and other stuff that they called the “leakage current”. These currents have to pass through gates and the probability that these gates are open or closed depends in some complicated way on the voltage. We won’t go into the detail of the Huxley-Hodgkin model, but the end result is four equations in four variables which they were able to solve numerically. The solutions matched well the experimental data from studies of squid giant axons. (As opposed to giant squid axons; the squid in question was rather small by squid standards, the axon rather large by axon standards.) The giant axon is the part of the squid’s nervous system that is responsible for escaping danger by squirting out water like a jet propulsion system; the large size of the axon means that it has lower resistivity and so signals propagate faster, allowing quicker reaction times. You can read about the Huxley-Hodgkin model in many places, including a fairly decent summary on [Wikipedia](#).

As an aside: Hodgkin and Huxley were successive masters of Trinity College Cambridge in the 70s and 89s. Hodgkin was long before my time, but I did meet Andrew Huxley a couple of years before he died. I sat down next to him at lunch, embarrassingly ignorant of who he was even though his portrait was hanging on the wall in front of me. He asked me what I did and, when I told him I was a physicist, there was a glorious pause before he said: “I met J. J. Thomson once”.

The key feature in the Hodgkin-Huxley model is that a small input gives rise to a big, spiked response. There are many other biological situations where similar behaviour is seen, including heart muscle cells and the blooms of plankton. The general class of models are called *excitable systems* and in this section we explore the phenomenology of the simplest.

1.7.1 FitzHugh-Nagumo Model

The FitzHugh-Nagumo model is the following pair of coupled differential equations,

$$\begin{aligned}\frac{du}{dt} &= \frac{1}{\epsilon} \left(u + v - \frac{1}{3}u^3 - z(t) \right) \\ \frac{dv}{dt} &= -(u - a + bv) .\end{aligned}\tag{1.192}$$

In the context of signal propagation by neurons, $u(t)$ represents the voltage difference across a membrane, while $v(t)$ represents the ease with which sodium and potassium ions can cross the membrane.

The model depends on three positive constants, a , b and ϵ . The first pair sit in the regions


$$0 < b \leq 1 , \quad 1 - \frac{2b}{3} < a < 1 .\tag{1.193}$$

The remaining constant ϵ is arbitrary but, as the name suggests, we will soon take $\epsilon \ll 1$ which ensures that the dynamics of $u(t)$ is fast, while that of $v(t)$ is slow. In addition, there is an *input function* $z(t)$ which we get to specify and it acts as a forcing term in the first equation.

We start by setting $z(t) = 0$ and look at the dynamics of the equations. There is a single fixed point (u^*, v^*) which is the intersection of the two nullclines

$$v = u \left(\frac{1}{3}u^2 - 1 \right) = \frac{a - u}{b} .\tag{1.194}$$

We have to solve a cubic which is a little awkward, but we can extract the essential features if we plot the two graphs together, as shown in the figure. First note that the slope of the cubic at the origin is -1 , while the slope of the line is $-1/b$ which, by the first inequality in (1.193), is necessarily steeper. This tells us that there is only one intersection point.



Second, the cubic has roots at $u = 0$ and $u = \pm\sqrt{3}$. The minimum is at $u = 1$ where we have $v = -2/3$. The slightly weird inequality in (1.193) ensures that the linear graph sits above the cubic at its minimum, which means that they must meet later. In other words, we have

$$u^* > 1 .\tag{1.195}$$

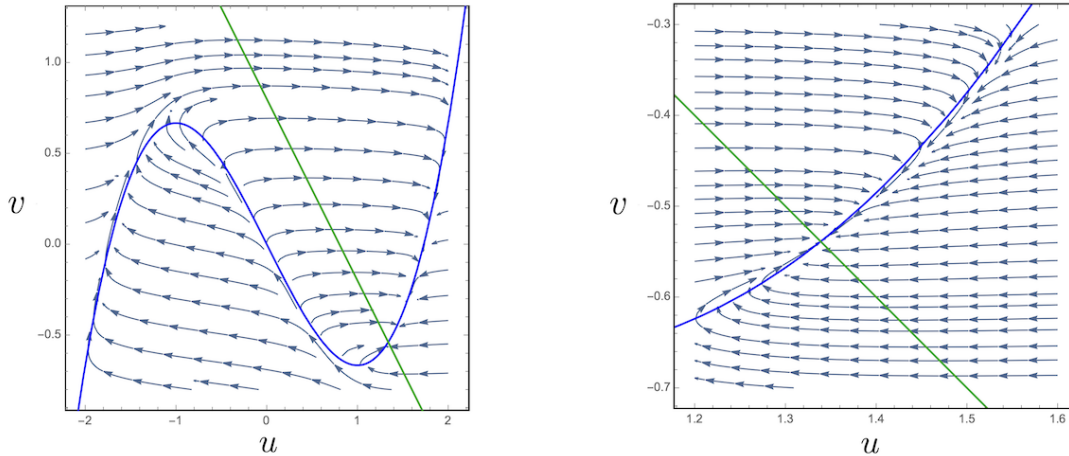


Figure 26. The phase plane dynamics of the FitzHugh-Nagumo model, plotted with $a = 0.8$, $b = 1$ and $\epsilon = 0.2$. Because ϵ is reasonably small, the dynamics separates into a fast mode, which heads towards the cubic, and a much slower mode which traces the curve of the cubic towards the fixed point where the nullclines meet. A close up of the flow near the fixed point is shown on the right.

Happily, this is all the information we need to determine the stability of the fixed point. The Jacobian is

$$J = \begin{pmatrix} (1 - u^2)/\epsilon & 1/\epsilon \\ -1 & -b \end{pmatrix}. \quad (1.196)$$

We have $\text{Tr } J < 0$ and $\det J > 0$, ensuring that the real part of both eigenvalues is negative and the fixed point is stable.

Now we assume that $\epsilon \ll 1$ so that the dynamics separates into a fast piece and a slow piece. The $u(t)$ variable is the fast piece, quickly tending almost horizontally towards the cubic $v = u^3/3 - u$. (In the context of neurons, this happens on a time scale of about 1 ms.) Then the $v(t)$ dynamics is slower. The result is plotted numerically in Figure 26. Note that the flow is unusual. If the flow hits the cubic to the right of the local minimum at $u = 1$, then the dynamics simply follows the cubic down towards the fixed point. If, in contrast, it hits the cubic to the left of that local minimum then it takes the long way round to the fixed point, moving up, before sweeping right as shown in Figure 26. The fact that some paths take these long deviations will be important in what follows.




Figure 27. On the left: the phase plane motion, shown in red, after the perturbation. The system starts at the original fixed point, where the lower cubic curve intersects the straight line. It then goes on a long detour in phase space before ending up at the new fixed point, where the upper cubic curve intersects the straight line. On the right: the resulting time dependence of $u(t)$ and $v(t)$. These plots are made with the same values as Figure 26, together with $V_0 = 0.3$.

Response to an External Input

Now we explore the role of the forcing term $z(t)$ in the equations (1.192). Suppose that by time $t = 0$, the system has happily settled down at the fixed point. Now we perturb the system by turning on $z(t)$, with the simple profile

$$z(t) = \begin{cases} 0 & \text{for } t < 0 \\ V_0 & \text{for } t > 0 \end{cases} \quad (1.197)$$

with V_0 constant.

The effect of the perturbation $V_0 > 0$ is to shift the cubic nullcline upwards, while leaving the linear nullcline alone. This means that the fixed point also shifts up and to the left. What happens next depends crucially on whether the new fixed point is stable or unstable. We saw previously that the fixed point is stable provided that $u^* > 1$ which ensures that it sits to the right of the local minimum in the cubic. It's simple to check that this persists for our perturbed system provided that $V_0 < V_{\text{crit}}$ with

$$V_{\text{crit}} = \frac{2}{3} + \frac{a-1}{b} . \quad (1.198)$$

We start by looking at the situation with $V_0 < V_{\text{crit}}$. If the perturbation V_0 is very small, so that the old fixed point sits above the minimum of the new cubic, then nothing dramatic happens: the system simply flows directly to the new fixed point.




Figure 28. On the left: the phase plane motion, shown in red, after the perturbation. Now the new fixed point is unstable and the system performs indefinite loops. On the right: the resulting time dependence of $u(t)$ and $v(t)$. These plots are made with the same values as Figure 26, together with $V_0 = 0.5$.

But a slightly larger V_0 , still with $V_0 < V_{\text{crit}}$, has a much more dramatic effect. If the old fixed point sits below the minimum of the new cubic then the system takes a large diversion to reach the new fixed point. This is shown on the left of Figure 27 which depicts both the original and the new nullclines, together with the trajectory in the phase space shown in red. This trajectory starts at the original fixed point, where the lower cubic curve intersects the line, and ends at the new fixed point at the intersection of the upper cubic curve. Although the two fixed points are close to each in phase space, the motion takes a huge detour to get from one to the other. The result is that a small perturbation of the original system can lead to a large pulse in the variables $u(t)$ and $v(t)$, before they settle down to equilibrium values that are close to their original values. This is shown on the right of Figure 27. In the context of neurons, this pulse in the voltage is what causes the neuron to fire.

The situation for $V_0 > V_{\text{crit}}$ is shown in Figure 28. Once again, the system takes a long detour to get from the old fixed point to the new. But now, once it arrives at the new fixed point, it finds that it's not particularly welcome as the fixed point is unstable and it gets pushed away. It is now condemned to loop forever on the limit cycle. The resulting dynamics for $u(t)$ and $v(t)$ are shown on the right of Figure 28 and exhibit an infinite series of pulses.

For both $V_0 < V_{\text{crit}}$ and for $V_0 > V_{\text{crit}}$, we see how a small change in the equation can lead to a large change in the resulting dynamics, either to initiate a single, transient pulse or to send the system into a new limit cycle.