Mathematical Modelling in Biology
Lecture Notes

Robin Thompson

Trinity Term 2026

Contents

Chapter 0: Acknowledgements

Thanks to the many lecturers who have contributed to the content of this course. Most of all, thanks to Prof. Ruth Baker for compiling the previous version of these notes.

If you spot any typos or other errors in these notes, please let me know!

My email address is robin.thompson@maths.ox.ac.uk

Chapter 1: Why study mathematical biology?

The field of mathematical biology involves using mathematics to understand real-world biological systems. The key goal is to further our understanding of the mechanisms underlying biological processes by building and analysing models that capture the main features of the system being studied. Models can be tested and refined based on real-world observations or experimental data, and can be used to make predictions.

Sub-topics of mathematical biology are wide ranging, and include population dynamics, infectious disease modelling, modelling cancer tumour dynamics and treatment, population genetics, bio-fluid dynamics and many other things. Mathematical biology was at the centre of the news during the COVID-19 pandemic, when infectious disease models were used to guide public health measures in countries around the world.

While mathematics and biology have been linked for centuries (e.g., Fibonacci's mathematical description of the hypothetical growth of rabbit populations in the 13th century), the field of modern mathematical biology is relatively new compared to other branches of applied mathematics. For example, the first infectious disease models of the type commonly used today were developed early in the 20th century. The UK, and Oxford in particular, has been at the forefront of the development of the subject in the last few decades, with the Centre for Mathematical Biology (now the Wolfson Centre for Mathematical Biology) established in the Mathematical Institute in Oxford in the 1980s.

The goal of this course is to provide an introduction to some of the methods used by mathematical biologists. Specifically, we will consider population dynamics models (including single species models and models of interacting populations), infectious disease models, models of chemical reactions in the body and models of neuron signalling. Through this course, you will learn some of the wide range of mathematical techniques used to study these systems.

If you enjoy this lecture course, a natural extension in Part B is the Further Mathematical Biology course (B5.5). Additional opportunities exist for you to use mathematics to study biological systems, including dissertations (e.g., a Part B Structured Project or Part C Dissertation on a Mathematical Topic) and other lecture courses (e.g., Part B Stochastic Modelling of Biological Processes or Part C Mathematical Physiology). I hope that this introductory course inspires you to want to learn more mathematical biology going forwards!

Chapter 2: Discrete-time models for a single species

In this chapter, we will introduce models describing the population dynamics of a single species. We will focus here on discrete-time models, in which the number of individuals in timestep $t+1$, N_{t+1}, depends on the number of individuals in timestep t, N_t, according to

$$N_{t+1} = f(N_t). \tag{2.1}$$

Usually, $f(N_t)$ is of the form $f(N_t) = N_t g(N_t)$, since we expect the population size to remain at zero if $N_t = 0$.

While the use of discrete time is a modelling simplification (time is clearly a continuous quantity!), it can be convenient in scenarios in which model simulations are compared with data that are collected at discrete time intervals. For example, the size of a population might be measured each week, and a discrete-time model may be used to represent these measurements.

We will consider techniques for analysing these models, specifically: i) finding steady states and assessing their linear stability; ii) bifurcation diagrams; and iii) cobwebbing.

2.1 Examples

Some common examples of discrete-time single species population growth models are given below. The model to use, and the values of its parameters (e.g., the value of r in the exponential model) can be chosen so that the model output is as similar as possible to available real-world data describing the population size at each timestep (e.g., each day or each week).

- Exponential model (or Malthusian model)

If, in a single timestep, the average number of births by each existing individual is b, and the probability that each individual dies is d, then an appropriate model might be $N_{t+1} = (1+b-d)N_t$. Under this model, the number of individuals in timestep $t+1$ depends linearly on the number of individuals in timestep t. This motivates the discrete-time exponential growth model, modelled by the so-called Malthusian equation,

$$N_{t+1} = rN_t, \tag{2.2}$$

in which $r > 0$. The parameter r is the exponential growth rate, or the *per capita* growth rate ("per capita" means that, in a single timestep, the change in the population size is *relative to the current population size*).

This model can be solved exactly to give

$$N_t = r^t N_0 \to \begin{cases} \infty & \text{if } r > 1, \\ N_0 & \text{if } r = 1, \\ 0 & \text{if } r < 1, \end{cases} \quad \text{as } t \to \infty\,. \tag{2.3}$$

- Logistic model

If $r > 1$, an unrealistic feature of the exponential model is that the population size will grow indefinitely. In reality, populations often exhibit a "carrying capacity" - a population size at which the population will no longer grow. This can be due, for example, to the food supply not being sufficient to enable larger population sizes.

A simple model that includes a carrying capacity is the logistic model, given by

$$N_{t+1} = N_t + rN_t\left(1 - \frac{N_t}{K}\right), \tag{2.4}$$

in which $r > 0$ and $K > 0$.

Under this model, if $N_t < K$ then the population size will increase in the next timestep, and if $N_t > K$ then the population size will decrease in the next timestep. If $N_t = K$, then the population size remains constant (a steady state; see later in this chapter).

- Ricker model

An unrealistic feature of the discrete-time logistic model is that, if N_t is large, then N_{t+1} will be negative; clearly, real-world population sizes cannot be negative. An alternative model that also includes a carrying capacity, but does not exhibit this unrealistic behaviour, is the Ricker model, given by

$$N_{t+1} = N_t \exp\left[r\left(1 - \frac{N_t}{K}\right)\right], \tag{2.5}$$

in which again $r > 0$ and $K > 0$.

As with the logistic model, if $N_t < K$ then the population size will increase in the next timestep, and if $N_t > K$ then the population size will decrease in the next timestep. We will return to this model later in this chapter. The Ricker model was originally formulated to study fish populations, and has been used in the context of managing fisheries (i.e., understanding how fish populations respond to different fishing levels).

In non-dimensionalised form (letting $N_t = Ku_t$), the Ricker model is given by

$$u_{t+1} = u_t \exp\left[r\left(1 - u_t\right)\right]. \tag{2.6}$$

- Hassell model

Another discrete-time model that can be used to model the temporal evolution of the population size of a single species is the Hassell model, given by

$$N_{t+1} = \frac{rN_t}{(1 + aN_t)^b}, \tag{2.7}$$

in which $r > 0$, $a > 0$ and $b > 0$. As the population size increases, the per capita growth rate decreases. As with all of the models described here, the parameters (here r, a and b) can be chosen so that the model output matches real-world data as closely as possible.

2.2 Steady states and linear stability

Definition. A *steady state* of the discrete-time single species population model $N_{t+1} = f(N_t)$ is a value, N^*, satisfying

$$N^* = f(N^*). \tag{2.8}$$

Given a steady state N^*, we can investigate its stability by considering an (initially) small perturbation around it, setting

$$N_t = N^* + n_t. \tag{2.9}$$

Then, noting that $N_{t+1} = N^* + n_{t+1}$ and that $N_{t+1} = f(N_t) = f(N^* + n_t)$, gives

$$N^* + n_{t+1} = f(N^* + n_t). \tag{2.10}$$

The right-hand-side can be expressed as a Taylor series about N^*, leading to

$$N^* + n_{t+1} = f(N^*) + n_t f'(N^*) + \mathcal{O}(n_t^2). \tag{2.11}$$

Since $N^* = f(N^*)$, we have

$$n_{t+1} \approx f'(N^*) n_t, \tag{2.12}$$

where $f'(N^*)$ is a constant, and thus

$$n_t \approx \left[f'(N^*)\right]^t n_0. \tag{2.13}$$

This means that N^* is linearly stable if $|f'(N^*)| < 1$ (as the perturbation n_t will then tend to zero as $t \to \infty$) and linearly unstable if $|f'(N^*)| > 1$.

2.3 Bifurcation diagrams

Definition. A *bifurcation point* is a point in parameter space at which the number of steady states changes or the stability of the steady states changes (or both of these things).

A *bifurcation diagram* is a plot of the values of the steady states (i.e., all N^* satisfying $N^* = f(N^*)$) as a function of a model parameter value. By convention, stable steady states are plotted using a solid line, and unstable steady states are plotted using a dashed line.

While this may at first seem complicated, we consider an example below to illustrate how steady states can be found and classified, how bifurcation diagrams can be plotted and how bifurcation points can be identified.

2.3.1 Example. Steady states and bifurcation diagrams

As an example, we consider analysing the model

$$N_{t+1} = rN_t\left(1 - N_t\right), \tag{2.14}$$

in which $r > 0$.

First, we find the steady states, which satisfy $N^* = f(N^*)$, where $f(N^*) = rN^*\left(1 - N^*\right)$. Hence, the steady states are $N^* = 0$ and $N^* = \frac{r-1}{r}$.

To classify these steady states, we calculate

$$f'(N^*) = r - 2rN^*. \tag{2.15}$$

For $0 < r < 1$, we have:

- $N^* = 0$ is a stable steady state, since $|f'(0)| = |r| < 1$;

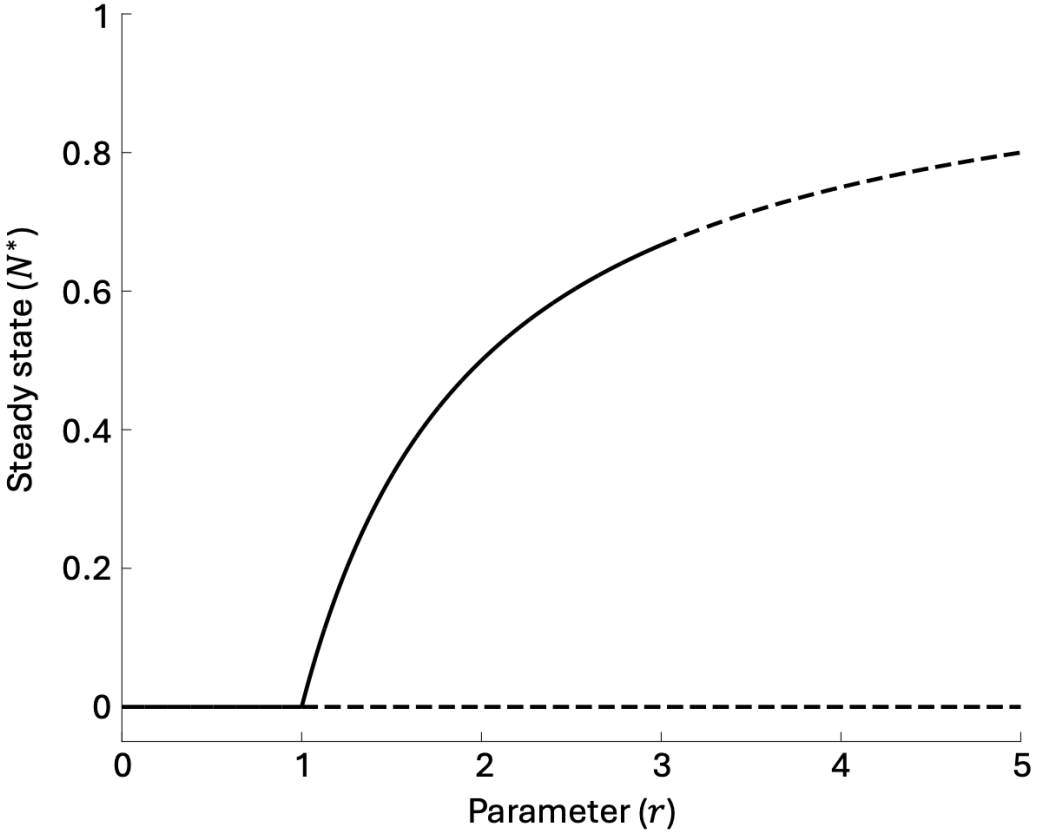

Figure 2.1: Bifurcation diagram for the model given in equation (2.14).

- $N^* = \frac{r-1}{r}$ is an unstable steady state, since $|f'(\frac{r-1}{r})| = |2 - r| > 1$. It is also unrealistic (it arises at a negative population size), and thus irrelevant for the purpose of modelling real-world populations.

For $1 < r < 3$, we have:

- $N^* = 0$ is an unstable steady state, since $|f'(0)| = |r| > 1$;
- $N^* = \frac{r-1}{r}$ is a stable steady state, since $|f'(\frac{r-1}{r})| = |2 - r| < 1$.

For $r > 3$, we have:

- $N^* = 0$ is an unstable steady state, since $|f'(0)| = |r| > 1$;
- $N^* = \frac{r-1}{r}$ is an unstable steady state, since $|f'(\frac{r-1}{r})| = |2 - r| > 1$.

The corresponding bifurcation diagram is shown in Figure 2.1. This figure is plotted only for positive values of N^*, since negative values correspond to negative population sizes (which are unrealistic). As noted above, stable steady states are indicated by solid lines and unstable steady states by dashed lines. At the values $r = 1$ and $r = 3$, there are changes in the stability properties of the steady states, and so $r = 1$ and $r = 3$ are bifurcation points.

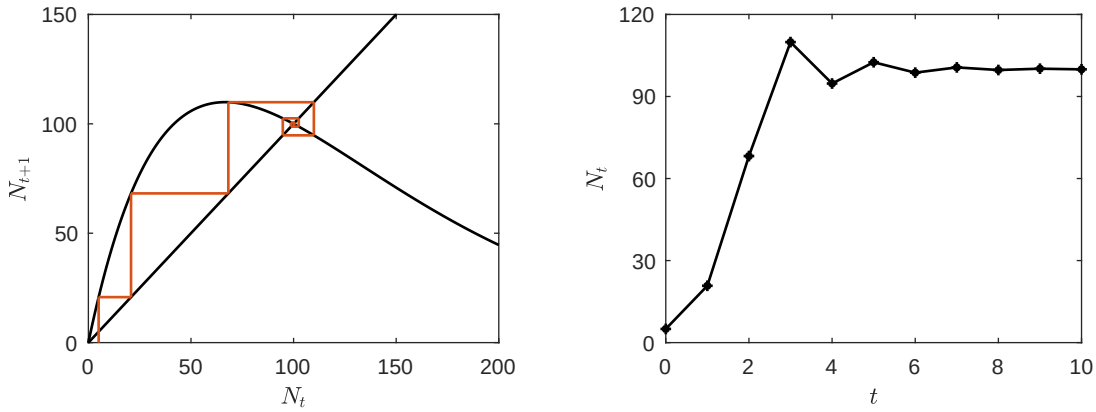

Figure 2.2: Dynamics of the Ricker model. The left panel shows a plot of $N_{t+1} = N_t \exp\left[r\left(1 - N_t/K\right)\right]$ alongside $N_{t+1} = N_t$ with the cobwebbing technique shown (starting at the beginning of the red line on the "x" axis). The right panel shows N_t at successive timesteps, $t = 1, 2, \ldots, 10$, as determined in the left panel. Parameter values are $N_0 = 5$, $r = 1.5$ and $K = 100$.

2.4 Cobwebbing

So far, we have seen how the steady states of a discrete-time single species population model can be found and analysed. In addition, we may wish to analyse the temporal evolution of the system.

To do this, we can construct a cobweb map using a technique that is known informally as *cobwebbing*.

To generate a cobweb map, we first plot N_{t+1} as a function of N_t (i.e., $N_{t+1} = f(N_t)$). We also plot the line $N_{t+1} = N_t$. We then perform the following steps:

- Step 1. Start on the "x-axis" at the initial value N_0.
- Step 2. Plot vertically to the line $N_{t+1} = f(N_t)$.
- Step 3. Plot horizontally across to the line $N_{t+1} = N_t$.
- Repeat steps 2-3 indefinitely!

Through this process, the temporal evolution of N_t can be visualised. Specifically, the "y" co-ordinate following step 2 in each iteration is N_1, N_2, N_3 and so on.

2.4.1 Example. Ricker model

In the left panel of Figure 2.2, we illustrate the construction of a cobweb map for the Ricker model (equation (2.5)), with $N_0 = 5$, $r = 1.5$ and $K = 100$. Successive values of N_t are plotted in the right panel.

2.4.2 Consistency with linear stability analysis

For the Ricker model with the parameters specified above, it can be seen from Figure 2.2 that the system proceeds to the steady state $N^* = K$. This is consistent with the results of a linear

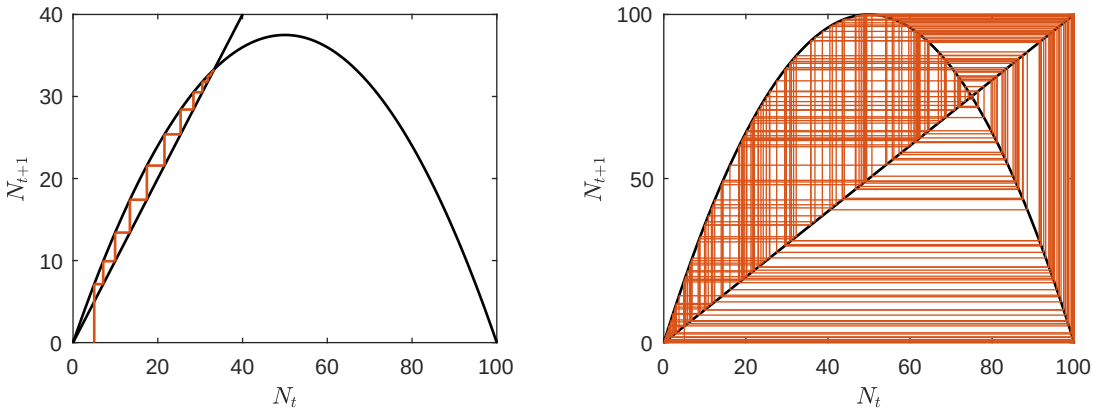

Figure 2.3: Dynamics of the model given by equation (2.16). The left panel shows results for $r = 1.5$ and the right panel shows results for $r = 4$. Other parameter values are $N_0 = 5$ and $K = 100$.

stability analysis performed using the method described in section 2.2; in particular, for this model, $|f'(K)| = |1-r|$, which is less than one when $r = 1.5$ (so the steady state is linearly stable).

More generally, the construction of cobweb maps also supports the idea that the stability of a steady state, N^*, depends on the value of $f'(N^*)$. For example:

- $-1 < f'(N^*) < 0$

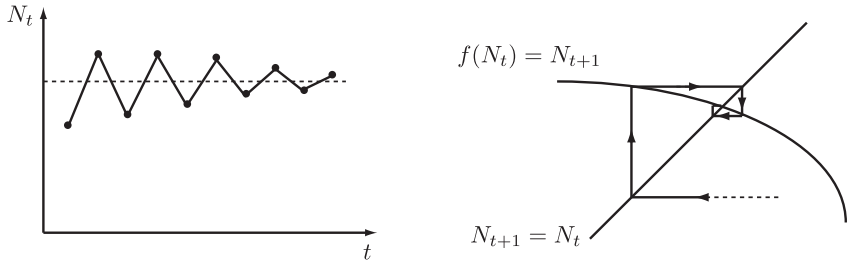

- $f'(N^*) = -1$

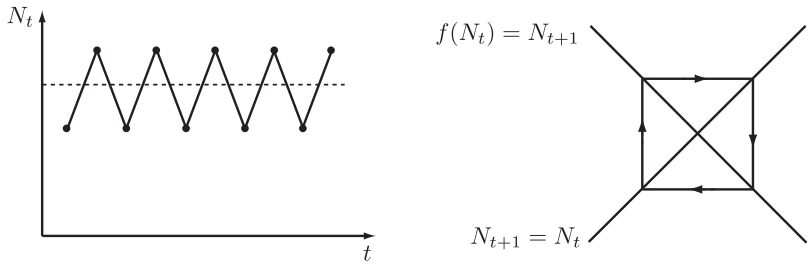

- $f'(N^*) < -1$

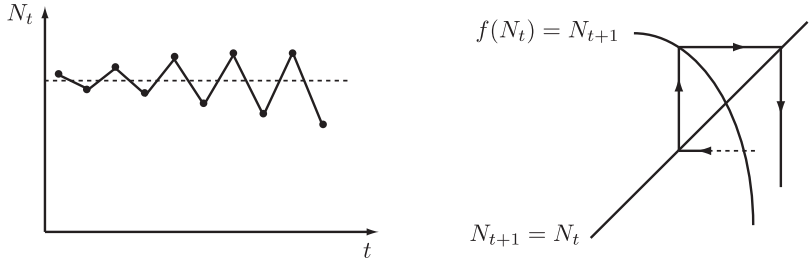

2.4.3 Warning!

It should be noted that cobwebbing does not always provide a straightforward depiction of a system's behaviour, even for simple models.

Consider, for example, the model

$$N_{t+1} = rN_t\left(1 - \frac{N_t}{K}\right), \quad r > 0, \quad K > 0. \tag{2.16}$$

This model is sometimes referred to as the logistic model in the mathematical biology literature. However, we note that it is not the same as our formulation of the logistic model in equation (2.4) (for example, in our formulation, the carrying capacity $N^* = K$ is a steady state, whereas in equation (2.16) it is not).

In the left panel of Figure 2.3, for $r = 1.5$ and $K = 100$, cobwebbing indicates that the model given by equation (2.16) progresses straightforwardly to the positive-valued steady state. However, in the right panel of Figure 2.3, the equivalent cobwebbing diagram is shown for $r = 4$; the dynamics are not so straightforward!

2.5 Summary

Now that we have reached the end of this chapter, given a discrete-time single species population model, you should be able to:

- Find the steady states.
- Undertake a linear stability analysis to classify their stability.
- Plot bifurcation diagrams and identify bifurcation points.
- Construct cobweb maps.

Chapter 3: Continuous-time models for a single species

In this chapter, we turn our attention from discrete-time models to continuous-time models describing the population dynamics of a single species. Such models are of the form

$$\frac{\mathrm{d}N}{\mathrm{d}t} = f(N), \tag{3.1}$$

in which $N(t)$ describes the number of individuals in the population at time t.

We will consider techniques that can be helpful for analysing these models, specifically: i) non-dimensionalisation; ii) finding steady states and assessing their linear stability; iii) bifurcation diagrams. In the context of bifurcation diagrams, we will introduce the concept of hysteresis.

Finally, as a case study, we will consider models that can be used to inform an ecologically acceptable strategy for harvesting a natural population (e.g., plants or animals). In that context, we will show how a model can be used to maximise the harvest yield and introduce the concept of the recovery time of the population.

3.1 Examples

Some examples of continuous-time single species population growth models are given below. As with discrete time models, the values of the parameters can be chosen so that the model output is similar to observed real-world data.

- Exponential model (or Malthusian model)

If, on average, each individual in the population generates new individuals at birth rate b, and each individual dies at rate d, then

$$\frac{\mathrm{d}N}{\mathrm{d}t} = (b-d)N. \tag{3.2}$$

This motivates the exponential model, in which the population grows (or declines) at an overall rate that is proportional to the current population size,

$$\frac{\mathrm{d}N}{\mathrm{d}t} = rN. \tag{3.3}$$

This model can be solved via separation of variables to give

$$N(t) = N_0 e^{rt}, \tag{3.4}$$

where N_0 is the initial population size (at $t = 0$).

- Logistic model (or Verhulst model)

As with the discrete-time exponential model, an unrealistic feature of the continuous-time exponential model is that population sizes can grow indefinitely (in the continuous-time case, this occurs whenever $r > 0$).

An alternative model is the logistic model, given by

$$\frac{\mathrm{d}N}{\mathrm{d}t} = rN\left(1 - \frac{N}{K}\right), \tag{3.5}$$

in which $r > 0$ and $K > 0$. The term on the right-hand-side is sometimes included in more complex models (see the spruce budworm model below); this term is referred to as "logistic growth".

The parameter K is the carrying capacity. When the population size N is less than K, then the population will grow. If instead N is greater than K, the population will decline. If $N = K$, the population size remains constant (a steady state; see later in this chapter).

Like the exponential model, the logistic model can be solved analytically. To do this, we separate variables to obtain

$$\int \frac{1}{rN\left(1 - \frac{N}{K}\right)}\mathrm{d}N = \int \mathrm{d}t. \tag{3.6}$$

Applying the method of partial fractions on the left-hand side gives

$$\frac{1}{r}\int \frac{1}{N} + \frac{1}{K - N}\mathrm{d}N = \int \mathrm{d}t, \tag{3.7}$$

which can be integrated to give

$$\frac{1}{r}\left(\ln(N) - \ln(K - N)\right) = t + C, \tag{3.8}$$

in which C is a constant of integration. Applying the initial condition $N(0) = N_0$ and rearranging then gives the solution

$$N = \frac{N_0}{\frac{N_0}{K} + \left(1 - \frac{N_0}{K}\right)\mathrm{e}^{-rt}}. \tag{3.9}$$

As $t \to \infty$, solutions of the logistic model tend monotonically to K; this can be seen in the plots in Figure 3.1.

- Spruce budworm predation model

The spruce budworm is a pest that feeds on balsam fir trees, eventually killing the trees. In 1978, Ludwig *et al.* introduced an equation for modelling the population dynamics of the spruce budworm, given by

$$\frac{\mathrm{d}N}{\mathrm{d}t} = r_B N\left(1 - \frac{N}{K_B}\right) - \frac{BN^2}{A^2 + N^2}, \tag{3.10}$$

in which A and B are positive constants. The first term on the right-hand side assumes logistic growth of the population, with carrying capacity K_B and growth rate r_B. The second term represents the effect of predation by birds upon the budworm population.

3.2 Non-dimensionalisation

Before analysing a model, a useful technique is non-dimensionalisation. The purpose of non-dimensionalisation is to reduce the number of model parameters by identifying dimensionless groupings that determine the dynamics.

For the purpose of this course, there are two scenarios in which you will be expected to be able to non-dimensionalise a model: i) when you are given the non-dimensional scalings; ii) when you are given the non-dimensional equation(s). We illustrate each of these scenarios in the context of the spruce budworm predation model.

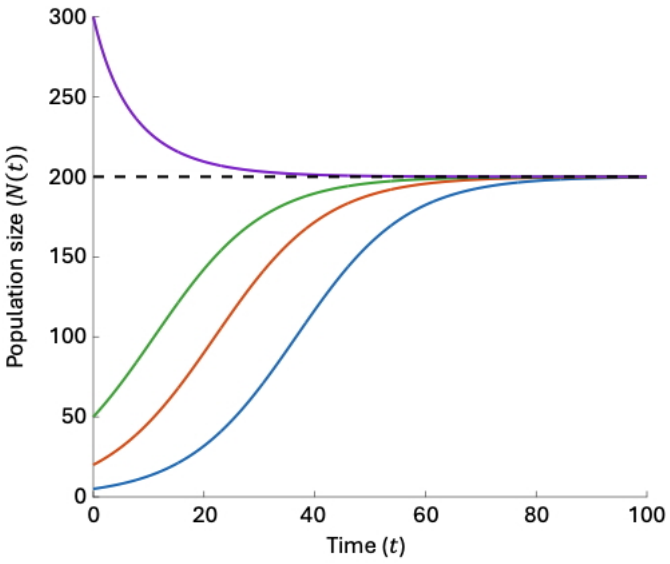

Figure 3.1: Logistic growth for $N_0 < K$ (blue, red and green) and $N_0 > K$ (purple). In this figure, $K = 200$ (as marked by the black dashed line), $r = 0.1$, and N_0 takes the values 5 (blue), 20 (red), 50 (green) or 300 (purple).

3.2.1 Non-dimensionalisation given scalings

Example. Non-dimensionalise the spruce budworm predation model by introducing the non-dimensional variables: $u = \frac{N}{A}$, $r = \frac{r_B A}{B}$, $q = \frac{K_B}{A}$, $\tau = \frac{Bt}{A}$.

To do this, we first note that $\frac{\mathrm{d}}{\mathrm{d}t} = \frac{\mathrm{d}\tau}{\mathrm{d}t}\frac{\mathrm{d}}{\mathrm{d}\tau} = \frac{B}{A}\frac{\mathrm{d}}{\mathrm{d}\tau}$.

Then, substituting the non-dimensional variables into the spruce budworm predation model gives (please check this!):

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = ru\left(1 - \frac{u}{q}\right) - \frac{u^2}{1+u^2}. \tag{3.11}$$

In this example, we have reduced the number of parameters in the model from four to two, thereby simplifying the model.

3.2.2 Non-dimensionalisation given the final model

Example. Non-dimensionalise the spruce budworm predation model to obtain

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = ru\left(1 - \frac{u}{q}\right) - \frac{u^2}{1+u^2}. \tag{3.12}$$

To do this, we need to find the scalings that transform the original model into the non-dimensional model that has been provided (assuming that we do not already know the scalings from the previous example!) We therefore introduce general scalings for the dependent and independent variables, $N = N^* u$ and $t = t^* \tau$, where N^* and t^* are constants that we must find. We must also find expressions for the parameters r and q of the non-dimensional model (in terms of the parameters of the original model).

Substituting these variables into the spruce budworm predation model, noting that $\frac{\mathrm{d}}{\mathrm{d}t} = \frac{\mathrm{d}\tau}{\mathrm{d}t}\frac{\mathrm{d}}{\mathrm{d}\tau} =$

$\frac{1}{t^*}\frac{\mathrm{d}}{\mathrm{d}\tau}$, gives

$$\frac{N^*}{t^*}\frac{\mathrm{d}u}{\mathrm{d}\tau} = r_B N^* u\left(1-\frac{N^* u}{K_B}\right) - \frac{B(N^*)^2u^2}{A^2+(N^*)^2u^2}, \tag{3.13}$$

$$\implies \frac{\mathrm{d}u}{\mathrm{d}\tau} = r_B t^* u\left(1-\frac{N^* u}{K_B}\right) - \frac{Bt^*N^*u^2}{A^2+(N^*)^2u^2}. \tag{3.14}$$

$$\implies \frac{\mathrm{d}u}{\mathrm{d}\tau} = r_B t^* u\left(1-\frac{u}{K_B/N^*}\right) - \frac{u^2}{(A^2/(Bt^*N^*))+((N^*)^2/(Bt^*N^*))u^2}. \tag{3.15}$$

We then compare equation (3.15) with the non-dimensional model that we are aiming for (equation (3.12)). This indicates that we would like to choose

$$r_B t^* = r, \qquad \frac{K_B}{N^*} = q, \qquad \frac{A^2}{Bt^*N^*} = 1, \qquad \frac{(N^*)^2}{Bt^*N^*} = 1. \tag{3.16}$$

These equations can then be solved to find

$$N^* = A, \qquad t^* = \frac{A}{B}, \qquad r = \frac{r_B A}{B}, \qquad q = \frac{K_B}{A}. \tag{3.17}$$

Hence, we have found scalings ($N = Au$ and $t = \frac{A}{B}\tau$) and parameter choices ($r = \frac{r_B A}{B}$ and $q = \frac{K_B}{A}$) that transform the spruce budworm predation model into

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = ru\left(1-\frac{u}{q}\right) - \frac{u^2}{1+u^2}, \tag{3.18}$$

as required.

As an aside, we note that if a model contains functions such as logarithms, trigonometric functions or exponentials, then the arguments of those functions must be non-dimensional. This is because these functions are defined to take purely numerical values. While this does not apply to the spruce budworm predation model, you might find this useful in the context of question 4 of problem sheet 1. ☺

We note that, while here we have considered a single species population model, the same approaches for non-dimensionalisation can be used in the context of other models, including models of interacting species (as studied in the next chapter).

3.3 Steady states and linear stability

Definition. A *steady state* for the continuous-time single species population model $\mathrm{d}N/\mathrm{d}t = f(N)$ is a value, N_s, satisfying

$$f(N_s) = 0. \tag{3.19}$$

Given a steady state N_s, we can investigate its stability by considering an (initially) small perturbation around it. We consider two approaches for assessing the stability of a steady state: i) graphical method; ii) linear stability analysis.

Graphical method

The graphical method involves plotting $f(N)$ (i.e. $\mathrm{d}N/\mathrm{d}t$) as a function of N. Steady states then correspond to points where the graph crosses the N axis.

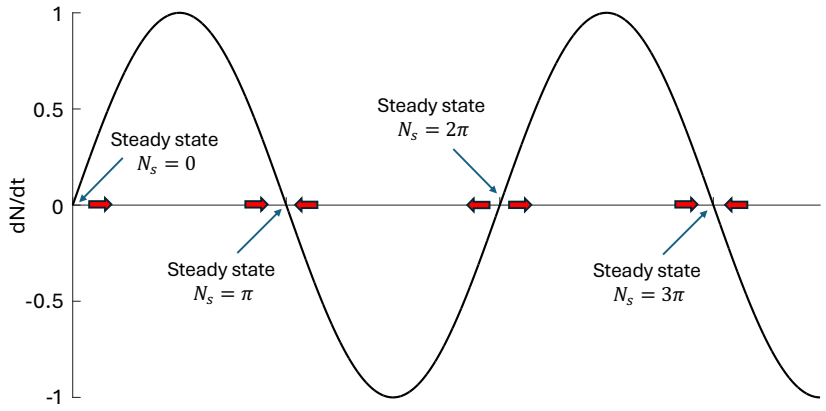

Figure 3.2: Assessing the stability of the steady states of the model $\mathrm{d}N/\mathrm{d}t = \sin(N)$ using the graphical method. As an example, for the steady state $N_s = \pi$, a small positive perturbation away from the steady state leads to values of N at which $\mathrm{d}N/\mathrm{d}t$ is negative, so that N decreases back to the steady state. Similarly, a small negative perturbation away from the steady state $N_s = \pi$ leads to values of N at which $\mathrm{d}N/\mathrm{d}t$ is positive, so that N increases back to the steady state. Thus, the steady state $N_s = \pi$ is stable, as indicated by the red arrows. By similar considerations, the steady state $N_s = 2\pi$ is unstable, the steady state $N_s = 3\pi$ is stable, and so on.

To assess their stability, we simply consider small perturbations around the steady state; if a small positive perturbation leads to a positive value of $f(N)$, and a small negative perturbation leads to a negative value of $f(N)$, then the steady state is unstable. If instead a small positive perturbation leads to a negative value of $f(N)$, and a small negative perturbation leads to a positive value of $f(N)$, then the steady state is stable.

We illustrate this procedure in Figure 3.2 for the model $\mathrm{d}N/\mathrm{d}t = \sin(N)$.

Linear stability analysis

We can assess the stability of a steady state, N_s, of the model $\mathrm{d}N/\mathrm{d}t = f(N)$ more formally by undertaking a linear stability analysis. To do this, we make a small perturbation, $n(t)$, about N_s, setting

$$N(t) = N_s + n(t). \tag{3.20}$$

Then

$$f\left(N(t)\right) = f\left(N_s + n(t)\right) = f(N_s) + n(t)f'(N_s) + \frac{1}{2}n(t)^2 f''(N_s) + \ldots, \tag{3.21}$$

where we denote $' = \mathrm{d}/\mathrm{d}N$. The second equality above results from a Taylor expansion. Hence

$$\frac{\mathrm{d}n}{\mathrm{d}t} = \frac{\mathrm{d}N}{\mathrm{d}t} = f\left(N(t)\right) = f(N_s) + n(t)f'(N_s) + \frac{1}{2}n(t)^2 f''(N_s) + \ldots. \tag{3.22}$$

We note that $f(N_s) = 0$ since N_s is a steady state, and then linearise by neglecting $\mathcal{O}(n^2)$ terms (since these terms are small), giving

$$\frac{\mathrm{d}n}{\mathrm{d}t} \approx f'(N_s)n(t) \qquad \Longrightarrow \qquad n(t) \approx n(0)\exp\left[\frac{\mathrm{d}f}{\mathrm{d}N}(N_s)t\right]. \tag{3.23}$$

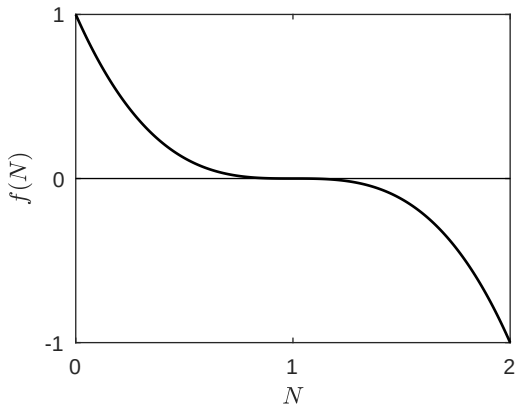

Figure 3.3: The function $f(N) = (1-N)^3$.

The steady state N_s is *linearly stable* if $n(t) \to 0$ as $t \to \infty$. In other words, N_s is linearly stable if

$$\frac{\mathrm{d}f}{\mathrm{d}N}(N_s) < 0. \tag{3.24}$$

Warning!

Note that we can find functions $f(N)$ such that $\mathrm{d}N/\mathrm{d}t = f(N)$ has a steady state which is stable and *not* linearly stable. For example, the equation

$$\frac{\mathrm{d}N}{\mathrm{d}t} = (1-N)^3, \tag{3.25}$$

has a steady state at $N_s = 1$ that is stable (see Figure 3.3, and consider applying the graphical method). However, if $f(N) = (1-N)^3$, then $f'(1) = 0$, so the steady state is not linearly stable.

Example. Steady states of the non-dimensional spruce budworm predation model

Here, we find and classify the steady states of the non-dimensional spruce budworm predation model. We recall that, following non-dimensionalisation, the model is given by

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = ru\left(1-\frac{u}{q}\right) - \frac{u^2}{1+u^2} := f(u). \tag{3.26}$$

The steady states are then given by the solutions of

$$ru\left(1-\frac{u}{q}\right) - \frac{u^2}{1+u^2} = 0. \tag{3.27}$$

We proceed graphically, and note that the steady states are given by the intersection of the graphs

$$f_1(u) = ru\left(1-\frac{u}{q}\right) \quad \text{and} \quad f_2(u) = \frac{u^2}{1+u^2}. \tag{3.28}$$

The left panel of Figure 3.4 shows plots of $f_1(u)$ and $f_2(u)$ for different values of q. We see that, depending on the value of q, we have either one or three non-zero steady states.

Noting that

$$\left.\frac{\mathrm{d}f}{\mathrm{d}u}\right|_{u=0} = r > 0, \tag{3.29}$$

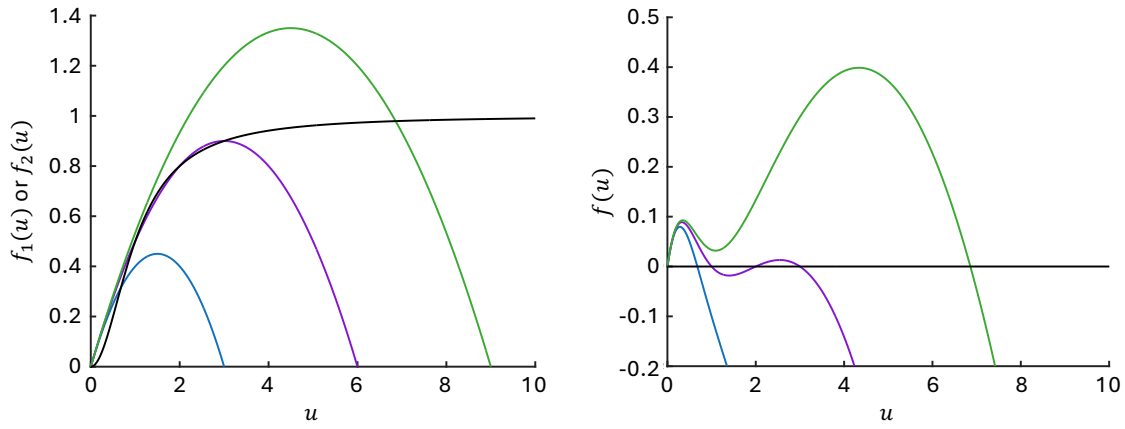

Figure 3.4: Characterising the dynamics of the non-dimensional spruce budworm predation model. Left: plots of the functions $f_1(u)$ (coloured lines; for $q = 3$ (blue), $q = 6$ (purple) and $q = 9$ (green)) and $f_2(u)$ (black line). Right: plot of $f(u)$ for $q = 3$ (blue), $q = 6$ (purple) and $q = 9$ (green). In both panels, $r = 0.6$.

typical plots of $\mathrm{d}u/\mathrm{d}\tau = f_1(u) - f_2(u)$ as a function of u are shown in the right panel of Figure 3.4 for a range of values of q. These plots can be used to classify the steady states using the graphical method. For example, when $r = 0.6$ and $q = 9$ (corresponding to the green line in the right panel of Figure 3.4), the steady state at $u = 0$ is unstable and the steady state that arises at a positive value of u is stable.

Bifurcation diagrams and hysteresis

As we saw when considering discrete-time single population models, a *bifurcation point* is a point in parameter space at which the number of steady states changes or the stability of the steady states changes (or both of these things).

A *bifurcation diagram* is a plot of the values of the steady states (in the context of models of the form $\frac{\mathrm{d}N}{\mathrm{d}t} = f(N)$, this is all N_s satisfying $f(N_s) = 0$) as a function of a model parameter value. Again, by convention, stable steady states are plotted using a solid line, and unstable steady states are plotted using a dashed line.

Bifurcation diagrams are useful to identify systems that exhibit *hysteresis*.

Definition. A system displaying *hysteresis* exhibits a response to the increase (or decrease) of a parameter that is not precisely reversed as the parameter is decreased (increased) to its original value.

A useful analogy to consider is walking off the edge of a cliff. If you walk forwards, you will fall off the cliff. If you then walk backwards, you will not return to the top of the cliff!

Example. Hysteresis in the non-dimensional spruce budworm predation model

We now demonstrate the occurrence of hysteresis in the non-dimensional spruce budworm predation model.

Suppose that we fix $r = 0.6$ in the model, and consider the bifurcation diagram as the parameter q varies. The top left panel of Figure 3.5 depicts $\frac{\mathrm{d}u}{\mathrm{d}\tau}$ as a function of u, for a range of values of

q. When q is small, there is only one non-zero steady state (S_1, say) which is stable. As q is increased past a critical value, $q = q_1$ say, then three non-zero steady states exist, (S_1, S_2 and S_3). As q is increased further, past a new critical value, $q = q_2$ say, the upper steady state S_3 is the only non-zero steady state that remains. In the top left panel of Figure 3.5, the red line is a plot for $q = q_1$ and the blue line is a plot for $q = q_2$. The corresponding bifurcation diagram is shown in the top right panel of Figure 3.5.

We now consider what happens to the system under the assumptions that: i) if the system is near a stable steady state, then it remains there; ii) if the system is near an unstable steady state, then it transitions away to near a stable steady state.

Suppose that q is small, and we are near the stable steady state S_1 (red dot in the bottom left panel of Figure 3.5). Then, as q increases past q_1, the new steady states (S_2 and S_3) appear but we remain near S_1 (following the red horizontal arrows in the bottom left panel of Figure 3.5). When q is eventually increased further past q_2, the steady state S_1 no longer exists and the system progresses to near the steady state S_3 (red vertical arrows in the bottom left panel of Figure 3.5). However, if q is then decreased back below q_2, since the stable steady state S_3 remains then the system would stay near S_3 rather than returning to near S_1 (blue horizontal arrows in the bottom left panel of Figure 3.5). Consequently, the change that arose as q is increased past q_2 (i.e. the system progressing from S_1 to S_3) is not reversed when q is then reduced below q_2 again (in fact, this change is only reversed later when q is lowered below q_1; see the blue vertical arrows in the bottom left panel of Figure 3.5).

We could ask "what is the biological interpretation of the presence of hysteresis in this model?" The answer is that, if the carrying capacity, q, is accidentally manipulated such that the spruce budworm population increases suddenly ($S_1 \to S_3$), then reversing this change in q is not sufficient to reduce the population substantially again.

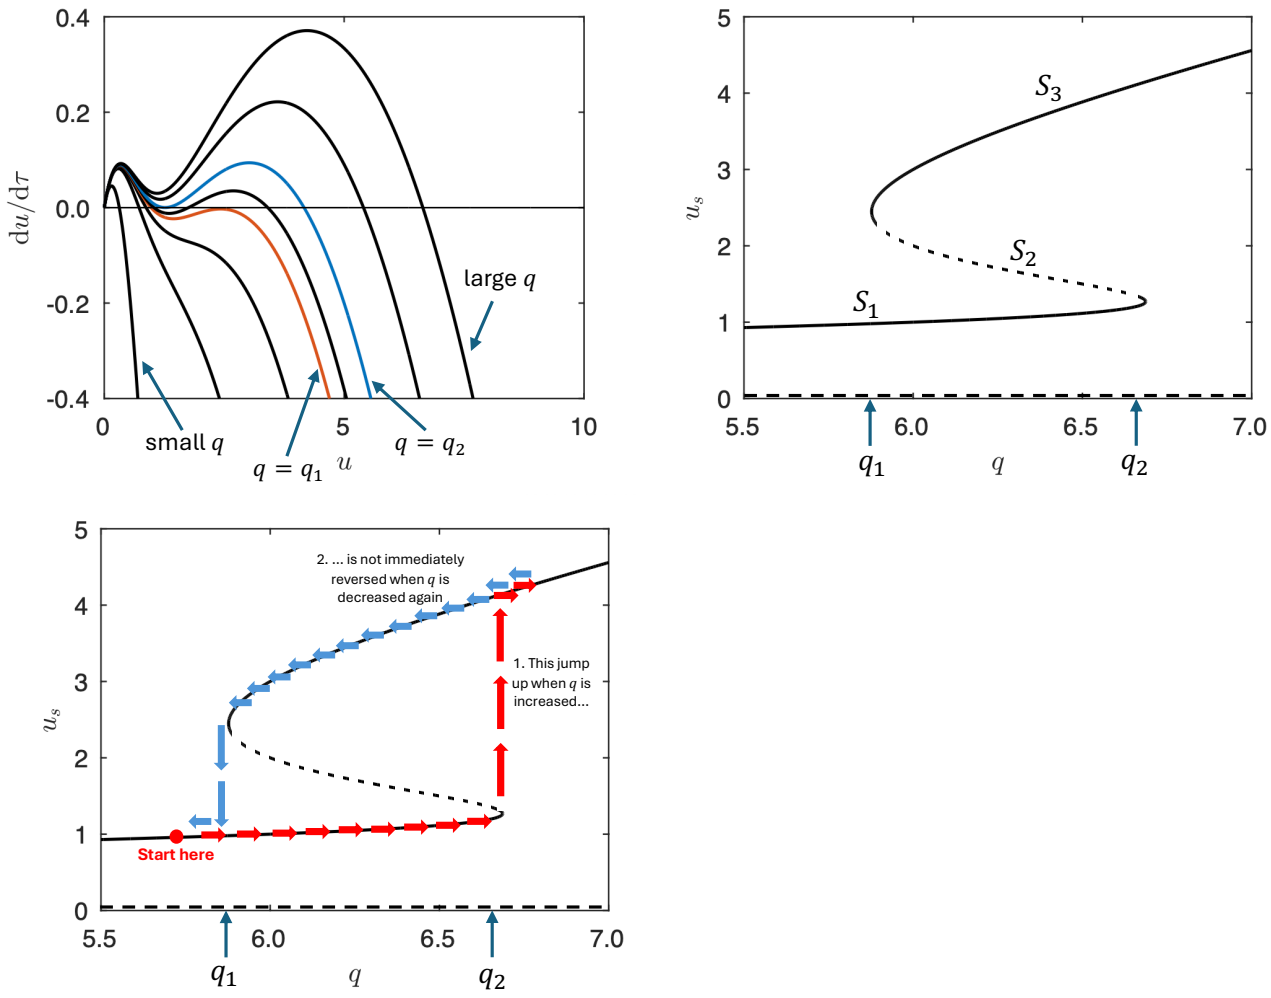

Figure 3.5: Top left: $\mathrm{d}u/\mathrm{d}\tau$ for the non-dimensional spruce budworm predation model as a function of u, plotted for different values of the parameter q. For small q, there is one, small, positive steady state. For $q \in (q_1, q_2)$, there are three positive steady states. For large q, there is one, large, positive steady state. Top right: The corresponding bifurcation diagram (the steady states plotted as a function of the parameter q). Bottom left: Arrows superimposed on the bifurcation diagram indicating the presence of hysteresis. Starting near the stable steady state at a low value of q (red dot), the system remains near the stable steady state as q is increased until q exceeds q_2. If q is then decreased below q_2 again, the system does not immediately revert to near the original stable steady state. The trajectory mapped out by the red and blue arrows is known as a "hysteresis loop" (although we note that the presence of a full loop is not a necessary condition for hysteresis).

3.4 Models of harvesting

To end this chapter, we consider two simple models in which a population is being harvested.

In the absence of harvesting, we assume that the population grows logistically. We then consider two harvesting models: the constant yield model and the constant effort model.

3.4.1 Constant yield

Under the constant yield model, the population size N is governed by

$$\frac{\mathrm{d}N}{\mathrm{d}t} = rN\left(1 - \frac{N}{K}\right) - Y_0. \tag{3.30}$$

Under this model, a constant harvest, Y_0, is collected from the population each unit of time.

As in the graphical method for finding and analysing steady states, we can plot $\frac{\mathrm{d}N}{\mathrm{d}t}$ as a function of N - see the left panel of Figure 3.6.

We consider starting from a large initial population size (i.e., $N(0)$ is assumed to be large).

If small yields are collected (i.e. Y_0 is small), the population size will approach a positive stable steady state.

If, however, larger yields are collected (i.e. Y_0 is larger) then $\frac{\mathrm{d}N}{\mathrm{d}t}$ is negative for all values of N and so the population size will decrease indefinitely. We note that, in this second scenario, $N(t)$ will exhibit unrealistic behaviour - the population size will become negative.

The threshold value of Y_0 above which $N(t)$ exhibits unrealistic behaviour starting from a large initial value of $N(0)$ can be found. Specifically, this occurs when the number of steady states changes from two to zero.

The steady states are given by the solutions, N_s, of

$$-\frac{r{N_s}^2}{K} + rN_s - Y_0 = 0 \quad \Rightarrow \quad N_s = \frac{r \pm \sqrt{r^2 - 4rY_0/K}}{2r/K}. \tag{3.31}$$

Therefore unrealistic behaviour will occur whenever

$$Y_0 > \frac{rK}{4}. \tag{3.32}$$

Due to the potential for this model to generate impossible behaviour (negative population sizes!), we focus instead on an alternative model: the constant effort model.

3.4.2 Constant effort

Under the constant effort model, the population size N is governed by

$$\frac{\mathrm{d}N}{\mathrm{d}t} = rN\left(1 - \frac{N}{K}\right) - EN. \tag{3.33}$$

Under this model, the population is harvested at a rate that is proportional to the current size of the population, $Y = EN$.

If the value of E is too large, then too much of the population is harvested and thus little harvest is available in the long-term. On the other hand, if the value of E is too small, then we do not collect much harvest.

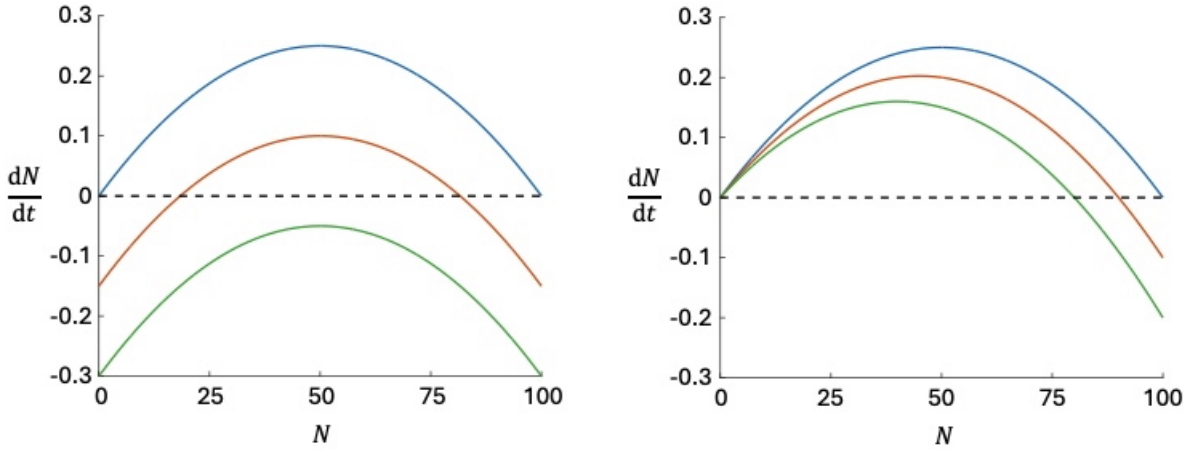

Figure 3.6: $\mathrm{d}N/\mathrm{d}t$ as a function of N for the constant yield harvesting model (left) and constant effort harvesting model (right). For the constant yield model, results are shown for $Y_0 = 0$ (blue), $Y_0 = 0.15$ (red) and $Y_0 = 0.3$ (green); as Y_0 is increased beyond a critical value, the steady states disappear and N decreases indefinitely. For the constant effort model, results are shown for $E = 0$ (blue), $E = 0.001$ (red) and $E = 0.002$ (green); unless the system starts from $N = 0$, it evolves in the long-term to the positive steady state (which can be seen from this figure to be stable). For both models, values of $K = 100$ and $r = 0.01$ are used.

This leads us to consider: how much effort should we put into harvesting if we want to maximise the long-term yield?

Again, we can plot $\frac{\mathrm{d}N}{\mathrm{d}t}$ as a function of N - see the right panel of Figure 3.6. We assume that $E < r$, otherwise the population will fade out and no harvest will be obtained in the long-term. Thus, the system approaches the positive (stable) steady state, which is given by

$$N_s = (r - E)\frac{K}{r}. \tag{3.34}$$

The long-term yield per unit time is then

$$Y = EN_s = (r - E)\frac{EK}{r}. \tag{3.35}$$

This is a quadratic equation with roots $E = 0$ and $E = r$. The effort that maximises this yield is therefore $E = \frac{r}{2}$, which obtains long-term yield $Y = \frac{rK}{4}$.

3.4.3 Recovery times

Finally, we present an approach for analysing whether or not a harvesting strategy is robust to perturbations (caused, for example, by differences between years or different weather events). To do this, we consider the recovery time of the population.

Definition. The *recovery time* of the population (in a harvesting model, or in a continuous time model of the population size of a single species) is defined to be the time for a small perturbation from a stable steady state to decrease by a factor of e.

We illustrate the notion of the recovery time by considering the logistic model and the constant effort harvesting model.

Recovery time in the logistic model

In the logistic model,

$$\frac{\mathrm{d}N}{\mathrm{d}t} = rN\left(1 - \frac{N}{K}\right), \tag{3.36}$$

and there is a positive stable steady state $N_s = K$.

We consider a small perturbation, $n(t)$, and set $N = K + n$. Substituting this into the logistic model and linearising (neglecting the quadratic term in n) gives

$$\frac{\mathrm{d}n}{\mathrm{d}t} \approx -rn, \tag{3.37}$$

so that

$$n(t) \approx n(0)e^{-rt}. \tag{3.38}$$

Hence, the recovery time for the logistic model is $\frac{1}{r}$.

Recovery time in the constant effort model

In the constant effort model,

$$\frac{\mathrm{d}N}{\mathrm{d}t} = rN\left(1 - \frac{N}{K}\right) - EN. \tag{3.39}$$

The positive stable steady state is $N_s = (r - E)\,\frac{K}{r}$.

We consider a small perturbation away from this steady state, $n(t)$, and set $N = N_s + n$. Substituting this into the logistic model and linearising (neglecting the quadratic term in n) gives

$$\frac{\mathrm{d}n}{\mathrm{d}t} \approx \left[rN_s\left(1 - \frac{N_s}{K}\right) - EN_s\right] + (r - E)n - 2rN_s\frac{n}{K}. \tag{3.40}$$

Noting that the terms in square brackets sum to zero (since, from the model, they are equal to $\frac{\mathrm{d}N}{\mathrm{d}t}\big|_{N=N_s}$, which must be zero), and then substituting in $N_s = (r - E)\,\frac{K}{r}$, gives

$$\frac{\mathrm{d}n}{\mathrm{d}t} \approx n(E - r), \tag{3.41}$$

so that

$$n(t) \approx n(0)e^{(E-r)t}. \tag{3.42}$$

Hence, the time for the perturbation to decrease by factor e (for example, from $n(0)$ to $n(0)\frac{1}{e}$) is $\frac{1}{r-E}$.

3.5 Summary

Now that we have reached the end of this chapter, given a continuous-time single species population model, you should be able to:

- Demonstrate how the model can be non-dimensionalised (either given scalings or given the non-dimensionalised model).
- Find the steady states.
- Classify their stability both graphically and via a linear stability analysis.
- Plot bifurcation diagrams.

- Identify bifurcation points and the presence of hysteresis.
- Obtain real-world conclusions (e.g., calculate the harvesting effort required to maximise the long-term yield in the constant effort harvesting model).
- Calculate the recovery time.

Chapter 4: Continuous-time models for interacting species

4.1 Background

In this chapter, we consider models of two species, with population sizes u and v, the dynamics of which can be described using the system of coupled ordinary differential equations

$$\frac{\mathrm{d}u}{\mathrm{d}t} = f(u,v), \tag{4.1}$$

$$\frac{\mathrm{d}v}{\mathrm{d}t} = g(u,v), \tag{4.2}$$

where f and g are prescribed functions that model the interactions between the species.

To analyse these models, we find and classify their steady states, and sketch the phase plane (via plotting the nullclines and considering the directions of trajectories in each region of the plane). After describing how these steps can be undertaken, we will consider three types of interaction between two species (predator-prey, competition and mutualism) and analyse models of these interactions.

4.1.1 Finding and classifying steady states

The *steady states*, (u_s, v_s), satisfy both $f(u_s, v_s) = 0$ and $g(u_s, v_s) = 0$.

To classify the steady states, we recall the approach presented in the Differential Equations I course (A1). In particular, we consider the Jacobian

$$\boldsymbol{J} = \begin{pmatrix} \frac{\partial f}{\partial u} & \frac{\partial f}{\partial v} \\ \frac{\partial g}{\partial u} & \frac{\partial g}{\partial v} \end{pmatrix}. \tag{4.3}$$

We calculate the eigenvalues of this matrix at the steady state (u_s, v_s); the eigenvalues determine the type of steady state. Denoting the eigenvalues by λ_1 and λ_2, then the possibilities are:

- $\lambda_1 < 0$, $\lambda_2 < 0$ (both real): Stable node.
- $\lambda_1 > 0$, $\lambda_2 > 0$ (both real): Unstable node.
- $\lambda_2 < 0 < \lambda_1$ (both real): Saddle point.
- $\lambda_1 = \mu + \nu i$, $\lambda_2 = \mu - \nu i$, $\mu < 0$: Stable spiral.
- $\lambda_1 = \mu + \nu i$, $\lambda_2 = \mu - \nu i$, $\mu > 0$: Unstable spiral.
- $\lambda_1 = \nu i$, $\lambda_2 = -\nu i$, $\mu > 0$: Centre.

Note. As an aside, in situations where we want to demonstrate that a steady state is stable, rather than to determine its precise type, a useful trick can sometimes be applied. Specifically, when calculating the eigenvalues of the Jacobian matrix, if the characteristic equation is of the form

$$\lambda^2 + b\lambda + c = 0, \tag{4.4}$$

with $b > 0$ and $c > 0$, then the steady state must be stable. This is because the eigenvalues are then of the form

$$\lambda_{1,2} = \frac{-b \pm \sqrt{b^2 - 4c}}{2}, \tag{4.5}$$

where either $b^2 - 4c < 0$ (in which case the steady state is a stable spiral) or $b^2 - 4c > 0$ (in which case the steady state is a stable node). This trick can sometimes be useful when demonstrating the stability of a steady state in a scenario in which finding precise expressions for the eigenvalues requires a large amount of calculation.

4.1.2 Sketching the phase plane

To sketch the phase plane, we undertake following steps.

- Sketch the nullclines; these are the lines on which either $\mathrm{d}u/\mathrm{d}t = 0$ (a "u nullcline") or $\mathrm{d}v/\mathrm{d}t = 0$ (a "v nullcline"). Note that, when plotting the phase plane (in (u, v) space), trajectories must cross each u nullcline exactly vertically and each v nullcline exactly horizontally.
- Identify the steady states. These are the points at which a u nullcline intersects with a v nullcline.
- The nullclines divide the phase diagram into different regions. Consider the direction of trajectories in each region of the plane (by considering $\mathrm{d}u/\mathrm{d}t$ and $\mathrm{d}v/\mathrm{d}t$ in each region).
- Draw trajectories on the phase plane, indicating their directions.

We note that this procedure should confirm the locations and types of steady states found using the approach described in section 4.1.1.

We illustrate how steady states can be found and classified, and phase planes sketched, for common ecological examples that we present in the following subsections.

4.2 Predator-prey models

4.2.1 Lotka-Volterra model

A widely used predator-prey model is the Lotka-Volterra model, which was devised early in the 20th century. Denoting the number of prey by N and the number of predators by P, the Lotka-Volterra model is given by

$$\frac{\mathrm{d}N}{\mathrm{d}t} = aN - bNP, \tag{4.6}$$

$$\frac{\mathrm{d}P}{\mathrm{d}t} = cNP - dP, \tag{4.7}$$

in which a, b, c and d are positive parameters and $c < b$.

Non-dimensionalisation

To reduce the number of model parameters, we first non-dimensionalise the model by setting $u = (c/d)N$, $v = (b/a)P$, $\tau = at$ and $\alpha = d/a$. This leads to the non-dimensional system

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = u(1-v) \coloneqq f(u,v), \tag{4.8}$$

$$\frac{\mathrm{d}v}{\mathrm{d}\tau} = \alpha v(u-1) \coloneqq g(u,v). \tag{4.9}$$

Steady states

Working with the non-dimensional model, we find and classify the steady states and sketch the phase plane. We also use the phase plane to sketch $u(\tau)$ and $v(\tau)$.

Steady states arise when both $\frac{\mathrm{d}u}{\mathrm{d}\tau} = 0$ and $\frac{\mathrm{d}v}{\mathrm{d}\tau} = 0$. Thus, the steady states are $(u_s, v_s) = (0,0)$ and $(u_s, v_s) = (1,1)$.

The Jacobian matrix is

$$\boldsymbol{J} = \begin{pmatrix} \frac{\partial f}{\partial u} & \frac{\partial f}{\partial v} \\ \frac{\partial g}{\partial u} & \frac{\partial g}{\partial v} \end{pmatrix} = \begin{pmatrix} 1-v & -u \\ \alpha v & \alpha(u-1) \end{pmatrix}. \tag{4.10}$$

Evaluating the steady state $(u_s, v_s) = (0,0)$, we find that the Jacobian is

$$\boldsymbol{J} = \begin{pmatrix} 1 & 0 \\ 0 & -\alpha \end{pmatrix}, \tag{4.11}$$

with eigenvalues 1, $-\alpha$. Therefore the steady state $(0,0)$ is a saddle point.

Instead evaluating the steady state $(u_s, v_s) = (1,1)$, the Jacobian is

$$\boldsymbol{J} = \begin{pmatrix} 0 & -1 \\ \alpha & 0 \end{pmatrix}, \tag{4.12}$$

with eigenvalues $\pm i\sqrt{\alpha}$. Therefore the steady state $(1,1)$ is a centre.

Phase plane

We demonstrate how we can use our "recipe" outlined in section 4.1.2 to plot the phase plane in Figure 4.1 (step 1: plot nullclines; step 2: identify steady states; step 3: plot directions of trajectories; step 4: draw trajectories).

Analytic solution

It turns out that, for the Lotka-Volterra model, it is possible to find equations for the trajectories analytically. To do this, we divide equation (4.8) by equation (4.9), giving

$$\frac{\mathrm{d}u}{\mathrm{d}v} = \frac{u(1-v)}{\alpha(u-1)v}. \tag{4.13}$$

This equation can be solved via separation of variables, leading to

$$\int \frac{u-1}{u}\mathrm{d}u = \int \frac{1-v}{\alpha v}\mathrm{d}v, \tag{4.14}$$

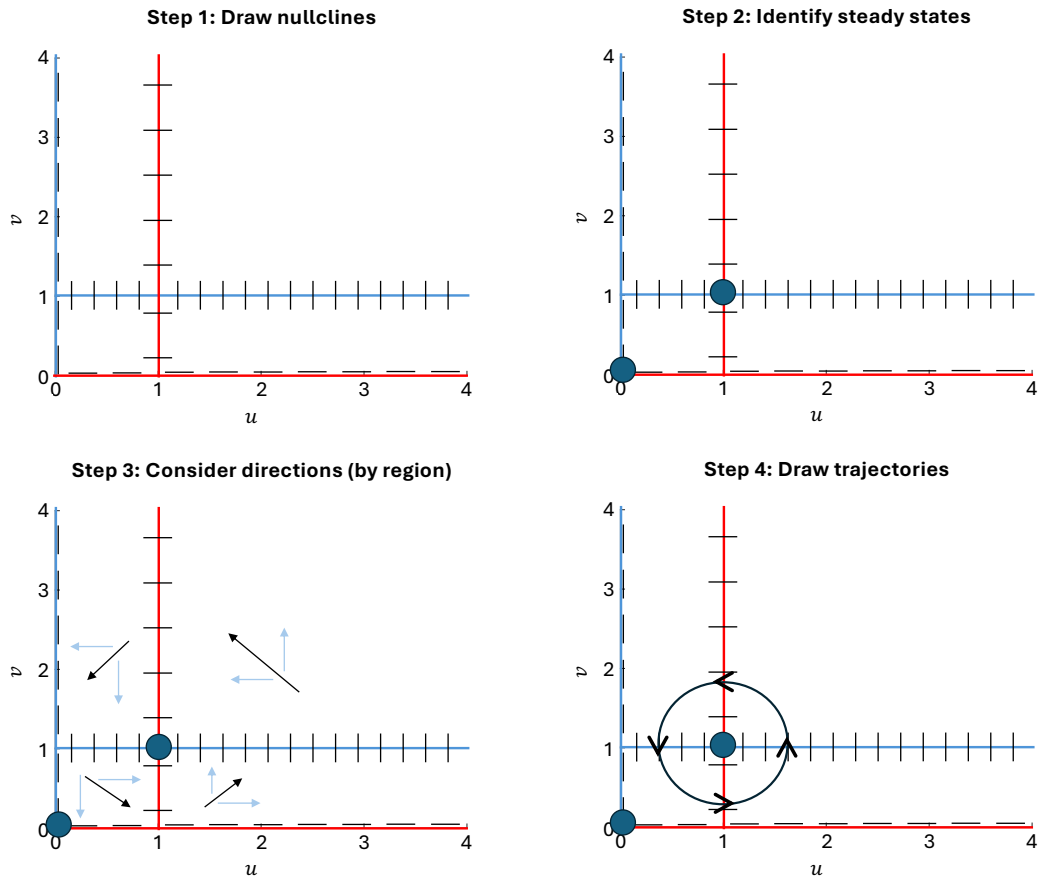

Figure 4.1: Schematic demonstrating how to sketch phase planes (here, for the non-dimensional Lotka-Volterra predator-prey model given by equations (4.8) and (4.9)). Step 1: Draw the nullclines (here, u nullclines are coloured blue and v nullclines are coloured red; the small black lines indicate the necessary direction - either horizontal or vertical - if trajectories cross the nullclines). Step 2: Identify the steady states (points where the u nullclines cross v nullclines). Step 3: Consider the directions of the trajectories in each region of the plane (by considering $\frac{\mathrm{d}u}{\mathrm{d}\tau}$ and $\frac{\mathrm{d}v}{\mathrm{d}\tau}$ in each region); Step 4: Draw on trajectories, based on the directions identified in step 3.

and so

$$\alpha u + v - \alpha \ln u - \ln v = C_1, \tag{4.15}$$

where C_1 is the constant of integration.

This can be rewritten as

$$\left(\frac{e^v}{v}\right)\left(\frac{e^u}{u}\right)^{\alpha} = C_2, \tag{4.16}$$

where C_2 is a constant (in fact, $C_2 = e^{C_1}$).

This enables us to plot the phase plane more accurately, by plotting (u, v) according to this equation for different values of C_2 (see left panel of Figure 4.2). This qualitatively agrees with the phase plane that we sketched in Figure 4.1.

It is also possible to sketch $u(\tau)$ and $v(\tau)$, as shown in the right panel of Figure 4.2. From this figure, we predict that u and v exhibit temporal oscillations, though not in phase. Hence we have a prediction; predators and prey populations oscillate out of phase. This makes sense

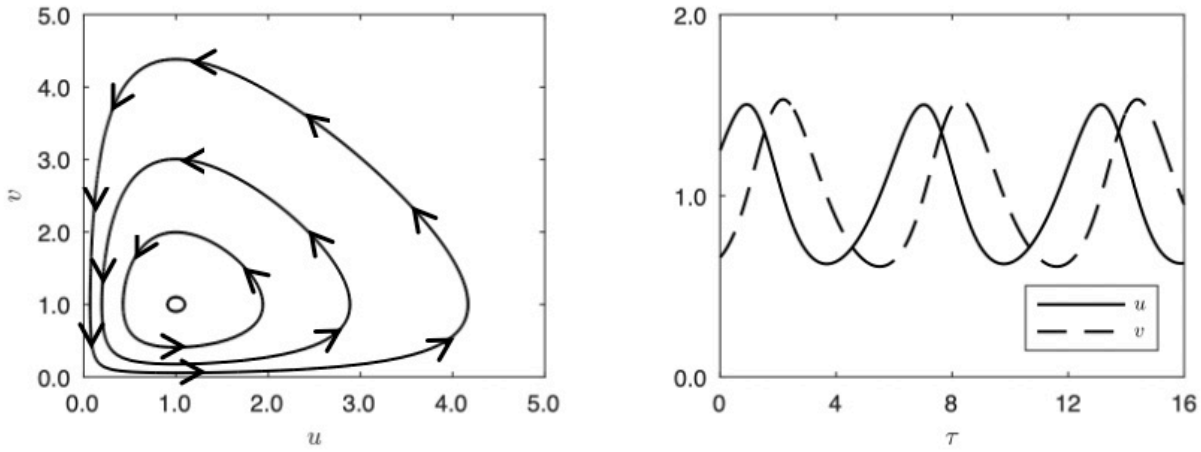

Figure 4.2: Dynamics of the non-dimensional Lotka-Volterra system for $\alpha = 1.095$ and $C_2 = 8.17, 11.02, 20.1$ and 54.6. The left-hand plot shows the dynamics in the (u, v) phase plane whilst the right-hand plot shows the temporal evolution of u and v.

ecologically, and indeed there are many examples of this in nature (e.g., a commonly studied example, due to the availability of a classic dataset from the 19th and 20th century, is oscillations in populations of lynx and hares).

4.2.2 Finite predation model

The Lotka-Volterra model assumes that, as $N \to \infty$, the rate of predation per predator becomes unbounded, as does the rate of increase of the size of the predator population. However, with an abundance of food, these quantities will saturate rather than become unbounded. Thus, a more realistic incorporation of an abundance of prey requires the Lotka-Volterra model to be refined.

One possibility is to instead use a model that we refer to here as the finite predation model, which (after non-dimensionalisation) is given by

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = u(1-u) - \frac{auv}{d+u}, \tag{4.17}$$

$$\frac{\mathrm{d}v}{\mathrm{d}\tau} = bv\left(1 - \frac{v}{u}\right), \tag{4.18}$$

in which a, b and d are positive constants. In these equations, the effect of predation per predator saturates at high levels of u.

Steady states

There is one steady state at which both the predators and prey can survive together, (u_s, v_s), satisfying

$$v_s = u_s \quad \text{where} \quad (1-u_s) = \frac{au_s}{d+u_s}. \tag{4.19}$$

Hence,

$$u_s = \frac{1}{2}\left[-(a+d-1) + \sqrt{(a+d-1)^2 + 4d}\right]. \tag{4.20}$$

The Jacobian at (u_s, v_s) is

$$\boldsymbol{J} = \begin{pmatrix} \frac{\partial f}{\partial u} & \frac{\partial f}{\partial v} \\ \frac{\partial g}{\partial u} & \frac{\partial g}{\partial v} \end{pmatrix}_{(u_s,v_s)} \tag{4.21}$$

where

$$\frac{\partial f}{\partial u}(u_s, v_s) = 1 - 2u_s - \frac{au_s}{d+u_s} + \frac{au_s v_s}{(d+u_s)^2} = -u_s + \frac{a(u_s)^2}{(d+u_s)^2}. \tag{4.22}$$

$$\frac{\partial f}{\partial v}(u_s, v_s) = -\frac{au_s}{d+u_s}, \tag{4.23}$$

$$\frac{\partial g}{\partial u}(u_s, v_s) = \frac{b(v_s)^2}{(u_s)^2} = b, \tag{4.24}$$

$$\frac{\partial g}{\partial v}(u_s, v_s) = b\left(1 - 2\frac{v_s}{u_s}\right) = -b. \tag{4.25}$$

The eigenvalues satisfy

$$\left(\lambda - \frac{\partial f}{\partial u}\right)\left(\lambda - \frac{\partial g}{\partial v}\right) - \frac{\partial f}{\partial v}\frac{\partial g}{\partial u} = 0 \quad \Longrightarrow \quad \lambda^2 - \left(\frac{\partial f}{\partial u} + \frac{\partial g}{\partial v}\right)\lambda + \left(\frac{\partial f}{\partial u}\frac{\partial g}{\partial v} - \frac{\partial f}{\partial v}\frac{\partial g}{\partial u}\right) = 0. \tag{4.26}$$

Hence

$$\lambda^2 - \alpha\lambda + \beta = 0 \quad \Longrightarrow \quad \lambda = \frac{\alpha \pm \sqrt{\alpha^2 - 4\beta}}{2}, \tag{4.27}$$

where

$$\alpha = -u_s + \frac{au_s^2}{(u_s+d)^2} - b, \quad \beta = b\left(u_s - \frac{au_s^2}{(u_s+d)^2} - (u_s - 1)\right). \tag{4.28}$$

Note that

$$\beta = b\left(1 - \frac{au_s^2}{(u_s+d)^2}\right) = b\left(1 - \frac{u_s(1-u_s)}{(u_s+d)}\right) = b\left(\frac{(u_s+d) - u_s + u_s^2}{u_s+d}\right) = b\left(\frac{d + (u_s)^2}{d+u_s}\right) > 0. \tag{4.29}$$

Thus, examining equation (4.27), if $\alpha < 0$ we have either a stable node (if $\alpha^2 - 4\beta > 0$) or stable spiral (if $\alpha^2 - 4\beta < 0$) at the steady state (u_s, v_s). If $\alpha > 0$, we have an unstable steady state at (u_s, v_s), which is either an unstable node (if $\alpha^2 - 4\beta > 0$) or an unstable spiral (if $\alpha^2 - 4\beta < 0$).

Phase plane

The u nullclines are given by

$$f(u, v) \equiv 0 \qquad \Longrightarrow \qquad u \equiv 0 \quad \text{and} \quad v = \frac{1}{a}(1-u)(u+d). \tag{4.30}$$

The v nullclines are given by

$$g(u, v) \equiv 0 \qquad \Longrightarrow \qquad v \equiv 0 \quad \text{and} \quad v = u. \tag{4.31}$$

A sketch of the nullclines and the behaviour of the phase plane trajectories for a case where the steady state is stable is shown in Figure 4.3.

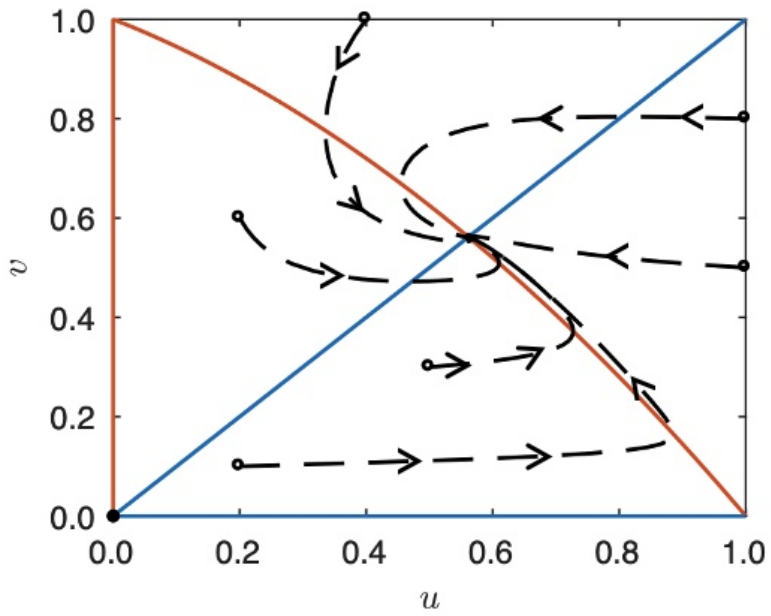

Figure 4.3: The (u, v) phase plane for the finite predation model in a scenario in which the co-existence steady state is stable. The u nullclines are plotted in orange and the v nullclines in blue. Trajectories for a number of different initial conditions are shown as dashed black lines.

4.3 Competition

4.3.1 Model of two competing populations

We consider a model of two competing populations. An example might be populations of red squirrels and grey squirrels, where both populations compete for the same resources. A typical model for their dynamics is

$$\frac{\mathrm{d}N_1}{\mathrm{d}t} = r_1 N_1 \left(1 - \frac{N_1}{K_1} - b_{12}\frac{N_2}{K_1}\right), \tag{4.32}$$

$$\frac{\mathrm{d}N_2}{\mathrm{d}t} = r_2 N_2 \left(1 - \frac{N_2}{K_2} - b_{21}\frac{N_1}{K_2}\right), \tag{4.33}$$

where K_1, K_2, r_1, r_2, b_{12}, b_{21} are positive constants. Let us associate N_1 with red squirrels and N_2 with grey squirrels in our example.

In particular, given a range of parameter values and some initial values for N_1 and N_2 at $t = 0$, we would like to know if the final outcome (i.e., the long-term behaviour of the populations) is one of the following possibilities:

- the reds become extinct, leaving the greys;
- the greys become extinct, leaving the reds;
- both reds and greys become extinct;
- the reds and greys co-exist.

This model can be non-dimensionalised to give

$$\frac{\mathrm{d}u_1}{\mathrm{d}\tau} = u_1(1 - u_1 - \alpha_{12}u_2), \tag{4.34}$$

$$\frac{\mathrm{d}u_2}{\mathrm{d}\tau} = \rho u_2(1 - u_2 - \alpha_{21}u_1), \tag{4.35}$$

where $\rho = r_2/r_1$.

Steady states

The steady states are

$$(u_{1,s}, u_{2,s}) = (0,0), \quad (u_{1,s}, u_{2,s}) = (1,0), \quad (u_{1,s}, u_{2,s}) = (0,1), \tag{4.36}$$

and

$$(u_{1,s}, u_{2,s}) = \frac{1}{1 - \alpha_{12}\alpha_{21}}(1 - \alpha_{12}, 1 - \alpha_{21}). \tag{4.37}$$

The coexistence steady state exists at positive population sizes if either: i) $\alpha_{12} < 1$ and $\alpha_{21} < 1$; or ii) $\alpha_{12} > 1$ and $\alpha_{21} > 1$.

The Jacobian is

$$\mathbf{J} = \begin{pmatrix} 1 - 2u_1 - \alpha_{12}u_2 & -\alpha_{12}u_1 \\ -\rho\alpha_{21}u_2 & \rho(1 - 2u_2 - \alpha_{21}u_1) \end{pmatrix}. \tag{4.38}$$

Steady state $(u_{1,s}, u_{2,s}) = (0,0)$**.**

$$\mathbf{J} - \lambda\mathbf{I} = \begin{pmatrix} 1 - \lambda & 0 \\ 0 & \rho - \lambda \end{pmatrix} \quad \Rightarrow \quad \lambda = 1, \rho. \tag{4.39}$$

Therefore $(0,0)$ is an unstable node.

Steady state $(u_{1,s}, u_{2,s}) = (1,0)$**.**

$$\mathbf{J} - \lambda\mathbf{I} = \begin{pmatrix} -1 - \lambda & -\alpha_{12} \\ 0 & \rho(1 - \alpha_{21}) - \lambda \end{pmatrix} \quad \Rightarrow \quad \lambda = -1, \rho(1 - \alpha_{21}). \tag{4.40}$$

Therefore $(1,0)$ is a stable node if $\alpha_{21} > 1$ and a saddle point if $\alpha_{21} < 1$.

Steady state $(u_{1,s}, u_{2,s}) = (0,1)$**.**

$$\mathbf{J} - \lambda\mathbf{I} = \begin{pmatrix} 1 - \alpha_{12} - \lambda & 0 \\ -\rho\alpha_{21} & -\rho - \lambda \end{pmatrix} \quad \Rightarrow \quad \lambda = -\rho, 1 - \alpha_{12}. \tag{4.41}$$

Therefore $(0,1)$ is a stable node if $\alpha_{12} > 1$ and a saddle point if $\alpha_{12} < 1$.

Steady state $(u_{1,s}, u_{2,s}) = \frac{1}{1-\alpha_{12}\alpha_{21}}(1 - \alpha_{12}, 1 - \alpha_{21})$**.**

$$\mathbf{J} - \lambda\mathbf{I} = \frac{1}{1 - \alpha_{12}\alpha_{21}} \begin{pmatrix} \alpha_{12} - 1 - \lambda & \alpha_{12}(\alpha_{12} - 1) \\ \rho\alpha_{21}(\alpha_{21} - 1) & \rho(\alpha_{21} - 1) - \lambda \end{pmatrix}. \tag{4.42}$$

Existence (at positive valued population sizes) and stability of this steady state depend on α_{12} and α_{21}.

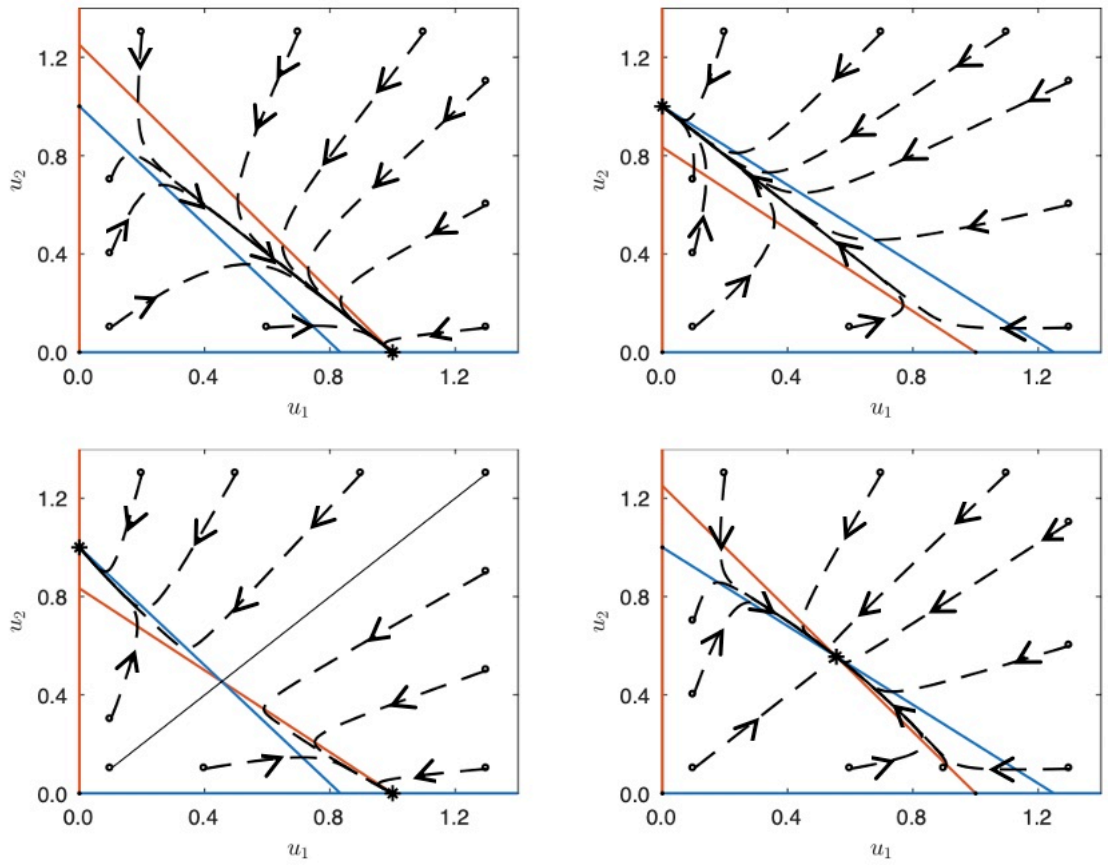

Figure 4.4: Dynamics of the two-species competition model (equations (4.34) and (4.35)). Top left: $\alpha_{12} = 0.8 < 1$, $\alpha_{21} = 1.2 > 1$ and the red squirrels (u_1) outcompete the grey squirrels (u_2). Top right: $\alpha_{12} = 1.2 > 1$, $\alpha_{21} = 0.8 < 1$ and the grey squirrels (u_2) outcompete the reds (u_1). Bottom left: $\alpha_{12} = 1.2 > 1$, $\alpha_{21} = 1.2 > 1$ and the species that survives is dependent on the initial conditions. Bottom right: $\alpha_{12} = 0.8 < 1$, $\alpha_{21} = 0.8 < 1$ and both species coexist. The stable steady states are marked with stars and $\rho = 1$ in all cases. The orange lines indicate u_1 nullclines while the blue lines indicate u_2 nullclines.

Phase plane

As suggested by the fact that the number and types of steady states differ depending on the values of α_{12} and α_{21}, the phase plane (and the long-term dynamics of the red and grey squirrel populations) varies for different values of α_{12} and α_{21} (i.e., different arrangements of the nullclines).

Careful construction of the phase planes using the standard "recipe" is left as an exercise. However, the qualitatively different cases are:

- $\alpha_{12} < 1 < \alpha_{21}$ (top left panel of Figure 4.4): the red squirrels outcompete the greys.
- $\alpha_{21} < 1 < \alpha_{12}$ (top right panel of Figure 4.4): the grey squirrels outcompete the reds.
- $\alpha_{12} > 1$ and $\alpha_{21} > 1$ (bottom left panel of Figure 4.4). The species that "wins" depends on the initial conditions.

- $\alpha_{12} < 1$ and $\alpha_{21} < 1$ (bottom right panel of Figure 4.4). Both the red squirrels and grey squirrels co-exist.

Numerical plots of phase planes demonstrating examples of these different cases are shown in Figure 4.4.

4.4 Mutualism

Mutualism, or symbiosis, involves two species gaining positively from their interaction. For example, oxpecker birds feed on large mammals (e.g. rhinos or zebras) by eating the parasites (e.g. ticks) on the mammals' bodies. The birds benefit by having easy access to food, and the mammals benefit as their parasite loads are controlled.

4.4.1 Model of mutualism

We consider a similar model to that of competition, but now with positive interactions,

$$\frac{\mathrm{d}N_1}{\mathrm{d}t} = r_1 N_1 \left(1 - \frac{N_1}{K_1} + b_{12}\frac{N_2}{K_1}\right), \tag{4.43}$$

$$\frac{\mathrm{d}N_2}{\mathrm{d}t} = r_2 N_2 \left(1 - \frac{N_2}{K_2} + b_{21}\frac{N_1}{K_2}\right), \tag{4.44}$$

where K_1, K_2, r_1, r_2, b_{12} and b_{21} are positive constants. The model can be non-dimensionalised to give

$$\frac{\mathrm{d}u_1}{\mathrm{d}\tau} = u_1(1 - u_1 + \alpha_{12}u_2), \tag{4.45}$$

$$\frac{\mathrm{d}u_2}{\mathrm{d}\tau} = \rho u_2(1 - u_2 + \alpha_{21}u_1). \tag{4.46}$$

In this model of mutualism, the non-axis nullclines have positive gradients. There are three steady states that have either u_1 or u_2, or both, equal to zero. An additional steady state with both u_1 and u_2 non-zero exists in some parameter regimes. The two possible behaviours displayed by the model are shown in Figure 4.5. We see that, when a non-zero steady state exists, it is stable, and the populations coexist. However, when this steady state does not exist, the populations grow unboundedly.

4.5 Summary

Now that we have reached the end of this chapter, given an ODE model describing the population dynamics of two interacting species, you should be able to:

- Find and classify the steady states.
- Sketch the phase plane by plotting the nullclines and considering the trajectories in different regions of the plane.
- Determine the type of interaction represented by a model (predator-prey, competition or mutualism).
- Interpret the model dynamics ecologically.

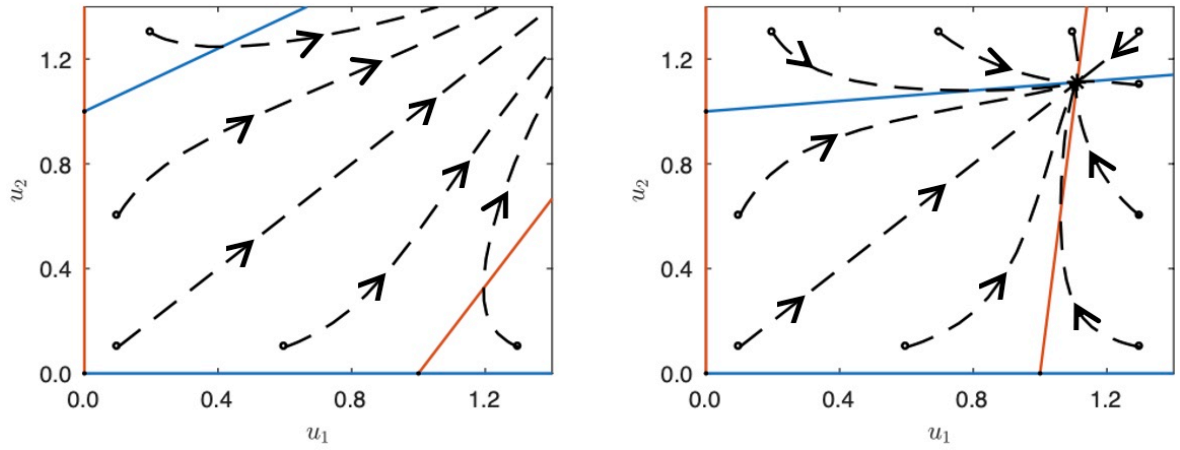

Figure 4.5: Dynamics of the non-dimensional mutualism system (equations (4.45) and (4.46)). The left-hand figure shows population explosion ($\alpha_{12} = 0.6 = \alpha_{21}$) whilst the right-hand figure shows population coexistence ($\alpha_{12} = 0.1 = \alpha_{21}$). The orange lines indicate u_1 nullclines and the blue lines indicate u_2 nullclines.

Chapter 5: Infectious disease modelling

Infectious disease modelling is an important sub-field of Mathematical Biology. Dynamic models of the type used to understand infectious disease outbreaks today were first developed early in the 21st century, but analyses of epidemiological data were undertaken long before that. For example, in the 18th century, Daniel Bernoulli analysed smallpox mortality data and explained how outbreaks could be controlled by variolation (i.e., inoculation of susceptible individuals with a small amount of material taken from a recently infected individual). In the 19th century, John Snow drew a map of the locations of cholera cases during the 1854 epidemic in London, finding that they all arose near the water pump in Broad Street. By doing this, Snow realised that cholera can be transmitted through water.

In the present day, infectious disease modelling analyses are increasingly used to guide public health measures during outbreaks. One of the first modern examples is the use of models to support decision making during the 2001 foot and mouth disease outbreak (a disease of animals). More recently, models were used during the 2014-16 Ebola epidemic in west Africa to generate projections of case numbers and to test different control measures. During the COVID-19 pandemic, models were used for a wide range of purposes, including estimating the level of transmission in different countries during different phases of the pandemic, testing possible vaccination strategies (e.g. prioritising vaccinating the oldest and most vulnerable individuals or younger individuals who tend to have more contacts with others) and planning the relaxation of control measures.

In this chapter, we will provide a brief introduction to two types of infectious disease model: compartmental models and renewal equation models. In the context of compartmental models, we will consider the most well-studied model: the SIR model. We will explain how key epidemiological quantities can be calculated: specifically, the basic reproduction number (R_0), the maximum number of individuals infected simultaneously, the final size, and the herd immunity threshold (the number of individuals who must be vaccinated for the disease to stop spreading). We will describe how the SIR model can be extended to include additional real-world realism, and investigate how a stochastic version of the SIR model can be used to calculate the probability that a disease will establish when it arrives in a new host population. In the context of renewal equation models, we will introduce the time-varying reproduction number and the end-of-outbreak probability.

5.1 Compartmental models (the SIR model)

In a compartmental model, individuals are categorised at any time according to their infection status. For example, in the Susceptible-Infectious-Removed (SIR) model, individuals are placed into three compartments:

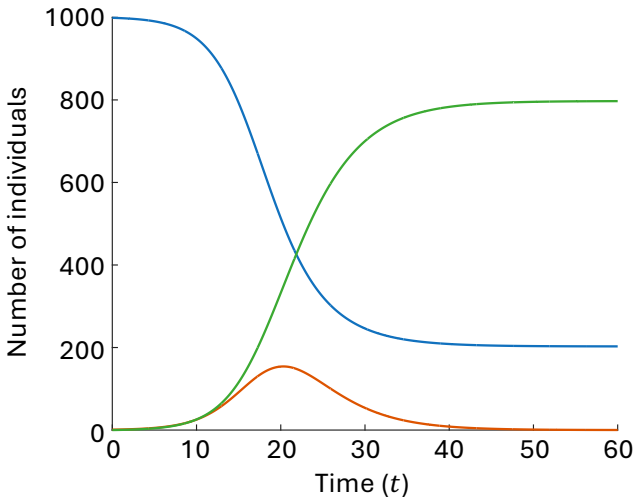

Figure 5.1: Numerical solution of the SIR model. $S(t)$, $I(t)$ and $R(t)$ are shown in blue, red and green, respectively. Parameter values: $\beta = 1/1500$, $\mu = 1/3$, $S(0) = 999$ and $I(0) = 1$.

- The susceptible compartment, S: individuals who can catch the disease;
- The infectious compartment, I: individuals who have caught the disease and can transmit it;
- The removed compartment, R: individuals who have recovered and are immune to the disease, or have died (in either case, they are no longer transmitting).

The SIR model is given by the following system of equations:

$$\frac{\mathrm{d}S}{\mathrm{d}t} = -\beta SI, \tag{5.1}$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \beta SI - \mu I, \tag{5.2}$$

$$\frac{\mathrm{d}R}{\mathrm{d}t} = \mu I. \tag{5.3}$$

We assume that, at the beginning of the outbreak,

$$S(0) = S_0, \quad I(0) = I_0, \quad R(0) = 0. \tag{5.4}$$

In the SIR model, the total infection rate is βSI. This expression is proportional to I since, if there are more infectious individuals, then we would expect the total infection rate to be higher. Similarly, it is also proportional to S, since if there are more individuals available for infection, we would also expect the total infection rate to be higher. The total removal rate is μI. This is because, if there are more infectious individuals, then we would expect the total removal rate to be higher.

An example numerical solution of the SIR model is shown in Figure 5.1. It can be seen that outbreaks fade out before the entire population has been infected (i.e., $S(\infty) > 0$).

We note that, in the SIR model as presented here,

$$\frac{\mathrm{d}}{\mathrm{d}t}(S + I + R) = 0 \quad \Longrightarrow \quad S + I + R = S_0 + I_0. \tag{5.5}$$

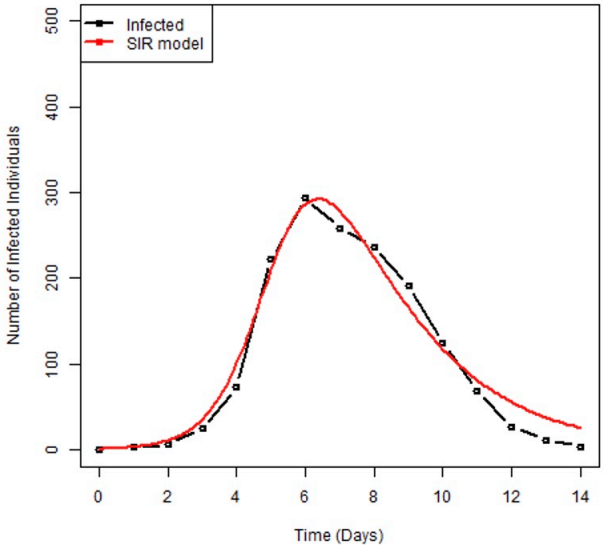

Figure 5.2: Data describing the number of infected individuals each day during the boys' boarding school influenza outbreak in the 1970s (black), alongside $I(t)$ obtained by solving the SIR model numerically (red) for an appropriate choice of β and μ.

In other words, the total population size is assumed to be constant, and we refer to the total population size $(S_0 + I_0)$ as N.

We note that there are a number of assumptions associated with the SIR model; for example, it is assumed that individuals mix homogeneously (each individual is equally likely to meet each other individual in the population) and it is assumed that individuals become infectious as soon as they are infected. Despite these assumptions, and despite the fact that the SIR model is a very simple model (there are only two tuneable parameters, β and μ), the SIR model can replicate real-world outbreak data closely. As an example, in Figure 5.2, we show data from an outbreak of influenza in a boys' boarding school in the north of England in the 1970s (black), alongside $I(t)$ obtained by solving the SIR model numerically (red) with β and μ chosen so that the model output and the data align closely.

5.2 Key epidemiological quantities

We now introduce some key epidemiological quantities, and illustrate how they can be calculated in the context of the SIR model.

5.2.1 Basic reproduction number (R_0)

Definition. The basic reproduction number, or R_0, is defined to be the expected number of infections generated by a single infected individual, in the absence of public health measures and if all of the other individuals in the population are susceptible.

To calculate R_0 using the SIR model, we note that, if individuals recover at rate μ, then the expected time for which a single infected individual is infectious is $1/\mu$. During that time, the individual generates new infections at rate $\beta S \approx \beta N$, since all individuals in the population are

assumed to be susceptible. Hence, R_0 is given by

$$R_0 = \text{infection rate} \times \text{infectious period} = \frac{\beta N}{\mu}. \tag{5.6}$$

In the SIR model, we also note that, early in an outbreak when $S \approx N$,

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \beta N I - \mu I = (\beta N - \mu)\, I, \tag{5.7}$$

which is positive if and only if $R_0 = \frac{\beta N}{\mu} > 1$. Hence, an outbreak will occur if and only if $R_0 > 1$.

It turns out that the threshold value of $R_0 = 1$ is a key concept in infectious disease modelling. If $R_0 > 1$, then early in an outbreak (when most people are susceptible) each infected individual is expected to infect more than one other person. As a result, an outbreak will occur (at least in the ODE version of the SIR model presented here). Conversely, if $R_0 < 1$, each infected individual is expected to infect fewer than one other, so the outbreak will fade out.

5.2.2 Maximum number infected

The maximum number of individuals who are ever infected simultaneously (i.e. the height of the peak in Fig 5.2) can be calculated for the SIR model (we denote the values of S and I at the peak of the outbreak by S_p and I_p, respectively). To do this, we divide equation (5.2) by equation (5.1) to obtain

$$\frac{\mathrm{d}I}{\mathrm{d}S} = -1 + \frac{\mu}{\beta S}. \tag{5.8}$$

Solving this equation and using the initial condition to find the constant of integration gives

$$I = -S + \frac{\mu}{\beta}\ln(S) + S(0) + I(0) - \frac{\mu}{\beta}\ln(S(0)). \tag{5.9}$$

Noting that $\frac{\mathrm{d}I}{\mathrm{d}S} = 0$ at the peak out the outbreak (whenever $R_0 > 1$), so that $S_p = \frac{\mu}{\beta}$, leads to

$$I_p = \frac{\mu}{\beta}\left(\ln\left(\frac{\mu}{\beta S(0)}\right) - 1\right) + S(0) + I(0). \tag{5.10}$$

The value of I_p increases as β increases. However, we note that public health measures may have unintended consequences. For example, an intervention that reduces β (such as partial isolation of infected individuals) would end up reducing I_p, as desired. However, this comes at the cost of the duration of the outbreak increasing (see, for example, Figure 5.3, in which we have plotted the temporal evolution of $I(t)$ for two different values of β; it can be seen there that the smaller value of β corresponds to a lower peak but a longer outbreak). During the COVID-19 pandemic, this phenomenon was referred to (in the context of social distancing strategies) as "flattening the curve".

5.2.3 Final size

Definition. The final size, or $R(\infty)$, is defined to be the total number of individuals infected over the duration of an outbreak.

If a public health policy maker is trying to choose between two different control strategies, one possibility is to choose the strategy that reduces the final size the most.

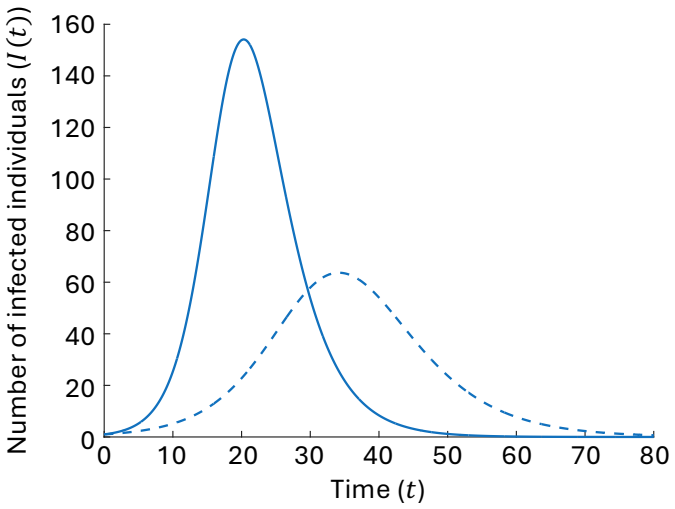

Figure 5.3: Temporal evolution of the number of infected individuals, $I(t)$, in the SIR model, for $\beta = 1/1500$ (blue) and $\beta = 3/4 \times 1/1500$ (blue dashed). This illustrates how a public health measure that reduces the value of β might reduce the peak of the outbreak, but at the cost of the outbreak lasting longer. Other parameter values: $\mu = 1/3$, $S(0) = 999$ and $I(0) = 1$.

An equation for the final size for the SIR model can be obtained from equation (5.9). In particular, we note that $I(t) \to 0$ as $t \to \infty$. Hence, taking $t \to \infty$ in equation (5.9) gives

$$-S(\infty) + \frac{\mu}{\beta}\ln(S(\infty)) + S(0) + I(0) - \frac{\mu}{\beta}\ln(S(0)) = 0. \tag{5.11}$$

Noting that all individuals end the outbreak either susceptible or removed, so that $S(\infty) + R(\infty) = S(0) + I(0)$, leads to the following implicit equation for $R(\infty)$:

$$R(\infty) = \frac{\mu}{\beta}\left(\ln(S(0)) - \ln(S(0) + I(0) - R(\infty))\right). \tag{5.12}$$

5.2.4 Herd immunity

Some public health measures aim to reduce transmission sufficiently that an outbreak cannot occur. If a perfectly protective vaccine is deployed, so that vaccinated individuals are guaranteed not to become infected, then in order to eradicate the disease it is not necessary to vaccinate all individuals. There is a critical proportion of the population, determined by R_0, that if vaccinated, will guarantee that an outbreak does not occur. The phenomenon whereby a critical threshold of the population is protected so that the entire population is guaranteed to avoid an outbreak is referred to as *herd immunity*. The critical proportion of the population is called the *herd immunity threshold*.

In the SIR model, the herd immunity threshold, denoted here by ν, can be calculated by considering a scenario in which $S = (1-\nu)N$ at the beginning of the outbreak. At that time,

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \beta(1-\nu)NI - \mu I = (\beta(1-\nu)N - \mu)\, I. \tag{5.13}$$

An outbreak will occur if $\frac{\mathrm{d}I}{\mathrm{d}t} > 0$ at the beginning of the outbreak. This corresponds to $\nu < 1 - \frac{1}{R_0}$. Hence, the herd immunity threshold (i.e. the target fraction of the population to vaccinate to prevent an outbreak) is $1 - \frac{1}{R_0}$.

5.2.5 Extensions to the SIR model

One of the benefits of compartmental modelling is that it is extremely straightforward to incorporate additional realism by simply adding more compartments. For example, in reality, individuals do not become infectious immediately after they become infected: there is a latent period following infection during which individuals are infected but not yet infectious. A latent period can be added to the SIR model by adding an Exposed (E) compartment between the S and I compartments, giving rise to the SEIR model:

$$\frac{\mathrm{d}S}{\mathrm{d}t} = -\beta SI, \tag{5.14}$$

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \beta SI - \gamma E, \tag{5.15}$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \gamma E - \mu I, \tag{5.16}$$

$$\frac{\mathrm{d}R}{\mathrm{d}t} = \mu I. \tag{5.17}$$

Alternatively, we may wish to model a disease for which individuals do not become immune when they recover, but instead because susceptible again. Examples include sexually transmitted infections such as gonorrhea or chlamydia. These diseases can be modelled using the SIS model:

$$\frac{\mathrm{d}S}{\mathrm{d}t} = -\beta SI + \mu I, \tag{5.18}$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \beta SI - \mu I. \tag{5.19}$$

A range of other real-world features of outbreaks can be incorporated into compartmental models simply by adding more compartments, including modelling diseases that are transmitted by vectors (e.g. dengue or malaria, which are transmitted by mosquitoes), spatial structure (for example, outbreaks across multiple cities; the SIR model can be employed within each city, with coupling between cities), age-structure and asymptomatic transmission. All of these extensions build on the basic SIR model that we have studied here.

5.2.6 Stochastic SIR model and the probability of a major outbreak

As we have seen, the classic SIR model is a system of ODEs. For a given set of parameter values, each time the model is solved numerically, the same smooth curves for $S(t)$, $I(t)$ and $R(t)$ are obtained. In real-world outbreaks, contacts between individuals and disease transmission are random processes. Linked to this, when a disease first arrives in a population, it is not guaranteed that it will establish in the population, even if R_0 is greater than one. For example, by chance, the first infected individual may recover without infecting anyone else, so that an outbreak does not occur.

An alternative model that accounts for randomness in transmission when simulating an outbreak is the *stochastic* SIR model (the SIR model that we studied above, which is made up of a system of ODEs, is sometimes called the *deterministic* SIR model). In the stochastic SIR model, if there are currently S susceptible individuals, I infected individuals and R removed individuals, then the next event is an infection event ($S \to S-1$ and $I \to I+1$) with probability $\frac{\beta IS}{\beta IS+\mu I}$ and the next event is a removal event ($I \to I-1$ and $R \to R+1$) with probability $\frac{\mu I}{\beta IS+\mu I}$. Three simulations of the stochastic SIR model, each generated starting with 999 susceptible individuals

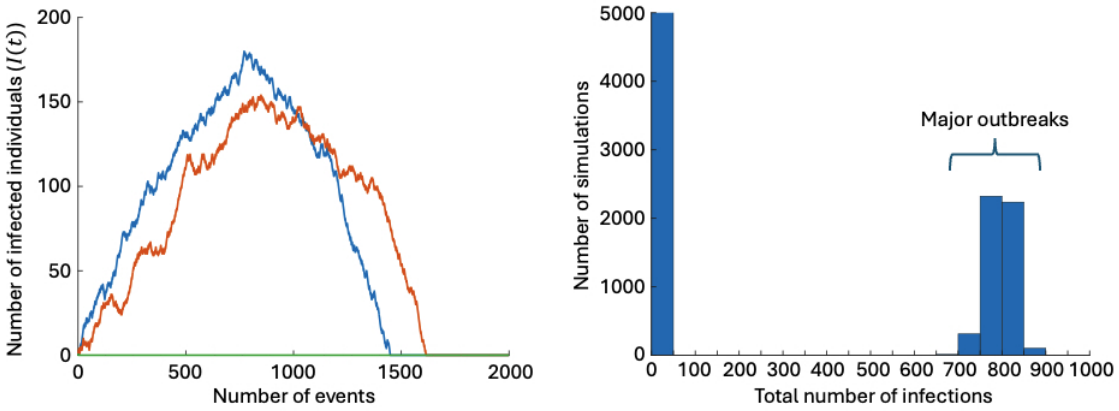

Figure 5.4: Left: Three simulations of the stochastic SIR model generated starting from a single infected individual. In the blue and red simulations, a major outbreak occurs, whereas the green simulation fades out without causing a major outbreak. Right: Bar chart showing the total number of infections, $R(\infty)$, in each of 10,000 simulations of the stochastic SIR model. Outbreaks can be divided in an obvious fashion into minor and major outbreaks. In both panels, in each simulation: $S(0) = 999, I(0) = 1; R(0) = 0, \beta = 0.0002, \mu = 0.1$. Note: the results shown here do not depend on the individual values of β and μ, but instead on the ratio β/μ.

and 1 infected individual, and each with identical parameter values, are shown in the left panel of Figure 5.4 (for each simulation, the line represents the value of I after each event). In the simulations shown in blue and red, a large outbreak occurs (these outbreaks are not identical, since the precise events that occur are not identical due to transmission being a random process governed by probabilities). In contrast, in the green simulation, the outbreak fades out without causing a large number of infections. In general, for R_0 larger than but not close to one, an outbreak will either fade out with few infections or a large number of infections will occur (see the right panel of Figure 5.4).

This gives rise to a concept known as the *probability of a major outbreak*. Starting from a single infected individual, the probability of a major outbreak is the probability that the disease does not simply fade out with few cases. The probability of a major outbreak therefore corresponds to the proportion of all simulations that belong to the cluster of large outbreaks in the right panel of Figure 5.4.

The probability of a major outbreak can be estimated as follows. First, we denote by q_i the probability that a major outbreak does not occur, starting from i infected individuals (with all other individuals in the population assumed to be susceptible). We then consider q_1, and condition on the outcome of the first event (either an infection event or a recovery event), so that by the Law of Total Probability,

$$q_1 = \frac{\beta \times 1 \times (N-1)}{\beta \times 1 \times (N-1) + \mu \times 1} q_2 + \frac{\mu \times 1}{\beta \times 1 \times (N-1) + \mu \times 1} q_0. \tag{5.20}$$

We note that q_0 is the probability of no major outbreak starting from no infected individuals, which is equal to one. We then make two approximations. First, we note that, in a sufficiently large population, $N - 1 \approx N$. Second, we note that the probability of no major outbreak arising starting from two infected individuals is similar to each infected individual independently failing to initiate a major outbreak, so that $q_2 \approx q_1^2$. Substituting these approximations into the equation

above, and dividing the numerator and denominator of each of the fractions in the equation above by μ, gives

$$q_1 = \frac{R_0}{R_0+1}q_1^2 + \frac{1}{R_0+1}, \tag{5.21}$$

since $R_0 = \frac{\beta N}{\mu}$.

This quadratic equation has solutions $q_1 = 1/R_0$ and $q_1 = 1$. Taking the smaller of these solutions, and noting that the probability of a major outbreak is $1 - q_1$, gives

$$\text{Probability of a major outbreak} = 1 - \frac{1}{R_0}, \tag{5.22}$$

whenever $R_0 > 1$ (and the probability of a major outbreak is zero whenever $R_0 \leq 1$).

Justification for taking the smaller solution for q_1 is beyond the scope of this course. However, it stems from the fact that q_1 is the extinction probability of a branching process (and you saw in Prelims Probability that the extinction probability is the smallest non-negative solution of the relevant equation).

5.3 Renewal equation models

Compartmental models such as the SIR model are ubiquitous in the field of infectious disease modelling. A challenge associated with using compartmental models, particularly those with a large number of compartments, is that it can be difficult to validate model simulations with real-world data. Doing so unambiguously would require modellers to be able to categorise all individuals in the population into the compartments of the model.

To avoid this complexity, an alternative type of model that can be used to simulate an infectious disease outbreak is a renewal equation. Renewal equation models simply track numbers of infections, rather than numbers of individuals in different model compartments.

A simple discrete-time renewal equation model is specified by

$$\mathbb{E}(I_t) = R_t \sum_{s=1}^{t-1} I_{t-s} w_s, \tag{5.23}$$

$$I_t \sim \text{Poi}(\mathbb{E}(I_t)). \tag{5.24}$$

We will unpick the various parts of this equation, but, in short, to simulate the number of infections on day t of the outbreak, you simply draw a number from a Poisson distribution with the given mean.

In this equation, R_t represents the time-varying reproduction number. This quantity is the expected number of infections generated by a single infected individual, based on the transmission conditions on day t (assuming that those conditions do not change going forwards). As such, the time-varying reproduction number is a measure of instantaneous transmissibility. At the beginning of an outbreak, when everyone in the population is susceptible (apart from the first case), R_t is equal to the basic reproduction number, R_0.

The values $w_1, w_2, w_3, \ldots$ represent the probability that the difference between infection times in infector-infectee transmission pairs is $1, 2, 3, \ldots$ days, respectively. In other words, if person A infects person B, then the probability that person B is infected the day after person A is w_1,

the probability that person B is infected two days after person A is w_2, and so on. The set of values $\{w_s\}_{s=1}^{\infty}$ is called the generation time distribution.

The form of the equation for $\mathbb{E}(I_t)$ can be explained as follows. To calculate the expected number of infections on day t, we first consider infections on day t arising from individuals infected on day $t-1$, I_{t-1}. Each of the I_{t-1} cases is expected to generate R_t infections each, with a proportion w_1 of those infections generated on day t. So, individuals infected on day $t-1$ are expected to generate $R_t I_{t-1} w_1$ infections on day t in total. Similarly, each of the I_{t-2} cases arising on day $t-2$ is expected to generate R_t infections each, with a proportion w_2 of those infections generated on day t. So, individuals infected on day $t-2$ are expected to generate $R_t I_{t-2} w_2$ infections on day t in total. Continuing this argument and summing over all past days gives the total number of infections that are to be expected on day t, namely $R_t \sum_{s=1}^{t-1} I_s w_s$.

Equations (5.23) and (5.24) can be used to simulate an outbreak. A common use of renewal equation models, that we do not explore further here, is estimating changes in R_t during outbreaks based on infection data (i.e., based on observations of I_t). These estimates were often reported in the news as the latest estimates of "the R number" during the COVID-19 pandemic.

5.3.1 End-of-outbreak probability

A key challenge for public health officials towards the end of an infectious disease outbreak is determining when the outbreak has finished. When a policy maker is confident that an outbreak is over, public health measures can be relaxed or removed, leading to both financial savings and permitting people to resume everyday behaviour that may have been restricted while the outbreak was ongoing.

Renewal equation models can be used to calculate the *end-of-outbreak probability*. This quantity represents, based on the observed data $\{I_s\}_{s=1}^{t-1}$, the probability that no more infections will occur from day t (the current day) onwards. Recalling that the probability mass function of a Poisson distributed random variable X with mean λ is given by $\mathbb{P}(X = k) = \frac{\lambda^k e^{-\lambda}}{k!}$, the end-of-outbreak probability is then

$$\mathbb{P}(I_t = 0, I_{t+1} = 0, \dots) = \prod_{k=t}^{\infty} e^{-R_k \sum_{s=1}^{k-1} I_{k-s} w_s}, \tag{5.25}$$

where we set $I_{k-s} = 0$ in this calculation whenever $k - s \geq t$.

An outbreak dataset and corresponding end-of-outbreak probability is plotted in Figure 5.5 in a scenario in which $R_t = 1.2$ for all values of t. In principle, a policy-maker might choose to declare the outbreak over when the end-of-outbreak probability exceeds a pre-determined threshold. For example, for the scenario considered in Figure 5.5, if the policy-maker declared the outbreak over when they are 95% sure that no more cases will arise (i.e., when the end-of-outbreak probability first exceeds 0.95), then they would declare the outbreak over when $t = 16$.

5.4 Summary

Based on the material in this chapter, you should be able to:

- Formulate simple compartmental epidemiological models, such as the (deterministic) SIR model.

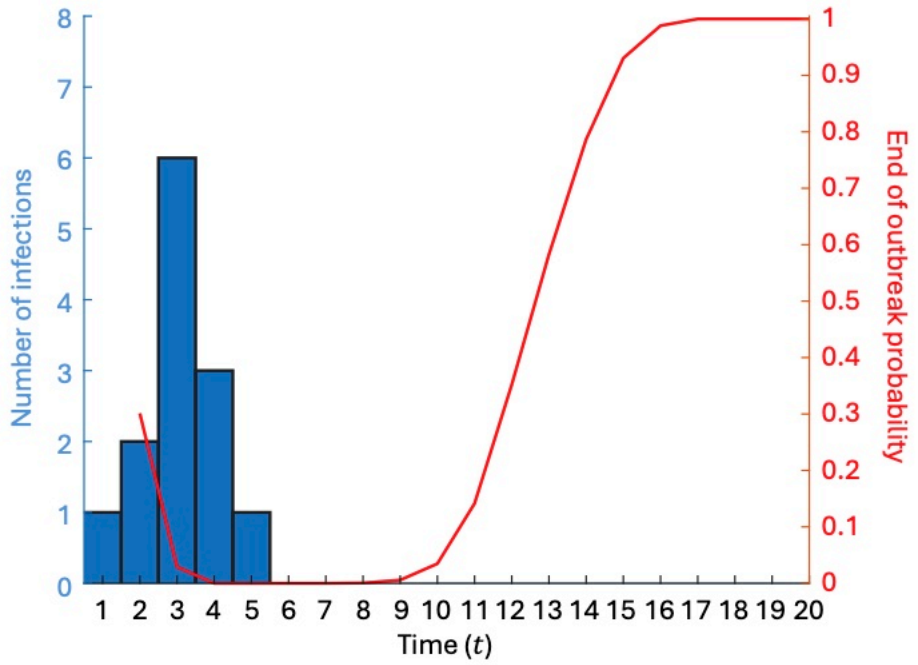

Figure 5.5: Calculation of the end-of-outbreak probability using a renewal equation model. Blue bars represent the number of infections each day, and the red line represents the end-of-outbreak probability estimate based on the infections arising prior to the current day. The end-of-outbreak probability was calculated assuming that $R_t = 1.2$ for all values of t and $\{w_s\}_{s=1}^{11} = \{0.05, 0.1, 0.2, 0.2, 0.15, 0.1, 0.1, 0.05, 0.02, 0.02, 0.01\}$ with $w_s = 0$ for $s \geq 12$.

- For the SIR model, calculate epidemiological quantities such as R_0, the maximum number infected simultaneously, the final size and the herd immunity threshold.
- Describe the stochastic SIR model and derive the probability of a major outbreak.
- Write down the basic renewal equation model and calculate the end-of-outbreak probability.

Chapter 6: Enzyme kinetics

In this chapter, we will consider developing and analysing models of biochemical reactions. Biochemical reactions are involved, for example, in immune responses to infections, metabolism and its control, and cell-signalling processes. Biochemical reactions are often controlled by enzymes. Enzymes are proteins that speed up chemical reactions without being consumed in the process.

As an example, red blood cells transport oxygen around the body. They contain a protein called haemoglobin, which is an enzyme. When blood passes through the lungs, haemoglobin combines with the oxygen, enabling the red blood cells to take up the oxygen. This process can also be reversed, so that the oxygen can be released by the red blood cells when they have moved around the body.

We will focus on a specific model, developed by Leonor Michaelis and Maud Menten early in the 20th century. In the model, a substrate combines with an enzyme to form a complex (and the reaction can also be reversed). The complex then splits into a product and the enzyme. Hence, overall, the substrate is transformed into a product (with the help of the enzyme). We will construct model equations using the Law of Mass Action, and non-dimensionalise the model. We will then undertake a phase plane analysis (as in Chapter 4). We will identify "fast" and "slow" parts of the dynamics, and solve the non-dimensionalised model to obtain equations governing the fast and slow dynamics. In the context of the slow dynamics, we will introduce the quasi steady state approximation (QSSA), and solving the fast dynamics will require us to perform a perturbation analysis. Finally, we will explain how the QSSA can be used to analyse more complex chemical reaction networks.

6.1 Michaelis-Menten model

In the Michaelis-Menten model, a substrate is turned into a product with the help of an enzyme. To do this, first the substrate and enzyme combine to form a complex (with rate constant k_1, say; see the Law of Mass Action below). This reaction is reversible, so the complex also splits into the substrate and enzyme (with rate constant k_{-1}). The complex also splits into the product and the enzyme (with rate constant k_2). Consequently, the system can be visualised as:

$$S + E \underset{k_{-1}}{\overset{k_1}{\rightleftharpoons}} C \qquad \text{and} \qquad C \xrightarrow{k_2} P + E, \tag{6.1}$$

where S is the substrate, E is the enzyme, C is the complex and P is the product. These reactions are representative of a large number of different enzyme systems.

To write down a model, we use the Law of Mass Action.

Definition. The Law of Mass Action A chemical reaction occurs at a rate that is proportional to the product of the concentrations of the reactants.

We denote the concentrations of S, E, C and P by s, e, c and p, respectively. Based on the Law of Mass Action, the rate that which the substrate and enzyme combine to form the complex is then k_1se. Considering all of the reactions in the Michaelis-Menten system, we obtain

$$\frac{\mathrm{d}s}{\mathrm{d}t} = -k_1se + k_{-1}c; \tag{6.2}$$

$$\frac{\mathrm{d}c}{\mathrm{d}t} = k_1se - k_{-1}c - k_2c; \tag{6.3}$$

$$\frac{\mathrm{d}e}{\mathrm{d}t} = -k_1se + k_{-1}c + k_2c; \tag{6.4}$$

$$\frac{\mathrm{d}p}{\mathrm{d}t} = k_2c. \tag{6.5}$$

We assume that initially we only have the substrate and enzyme present, and that there is substantially more substrate than enzyme initially, so the initial conditions are $s(0) = s_0$, $e(0) = e_0 \ll s_0$, $c(0) = 0$ and $p(0) = 0$.

We note that, by equation (6.5), the concentration of product depends only on the concentration of complex. Consequently, if we knew $c(t)$, then we would be able to find $p(t)$ by solving equation (6.5).

We also note that the concentration of enzyme is linked to the concentration of complex; adding equations (6.3) and (6.4) gives

$$\frac{\mathrm{d}}{\mathrm{d}t}(e + c) = 0 \quad \Longrightarrow \quad e(t) + c(t) = e_0. \tag{6.6}$$

Hence, we can find $e(t)$ using the formula

$$e(t) = e_0 - c(t). \tag{6.7}$$

We therefore focus on the amount of substrate and complex in this system (if we know $s(t)$ and $c(t)$, we could then find $p(t)$ and $e(t)$ if needed). Substituting (6.7) into equations (6.2) and (6.3) gives

$$\frac{\mathrm{d}s}{\mathrm{d}t} = -k_1(e_0 - c)s + k_{-1}c, \tag{6.8}$$

$$\frac{\mathrm{d}c}{\mathrm{d}t} = k_1(e_0 - c)s - (k_{-1} + k_2)c. \tag{6.9}$$

6.1.1 Non-dimensionalisation

Recall from Chapter 3.2 that, in this course, there are two possible scenarios in which you will be expected to non-dimensionalise a model: either the non-dimensional scalings will be given or the final non-dimensional equations will be given. We assume the former case here, and non-dimensionalise this model using the scalings:

$$\tau = k_1e_0t, \quad u = \frac{s}{s_0}, \quad v = \frac{c}{e_0}, \quad \lambda = \frac{k_2}{k_1s_0}. \tag{6.10}$$

Substituting these variables into equations (6.8) and (6.9), and also into the initial conditions $s(0) = s_0$ and $c(0) = 0$, gives

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = -u + \left(u + \left(\frac{k_{-1} + k_2}{k_1 s_0}\right) - \lambda\right) v, \tag{6.11}$$

$$\frac{e_0}{s_0}\frac{\mathrm{d}v}{\mathrm{d}\tau} = u - \left(u + \left(\frac{k_{-1} + k_2}{k_1 s_0}\right)\right) v, \tag{6.12}$$

where $u(0) = 1$ and $v(0) = 0$. We therefore also define $K = \frac{k_{-1}+k_2}{k_1 s_0}$ and $\epsilon = \frac{e_0}{s_0}$, noting that $\epsilon \ll 1$, to give the non-dimensionalised system

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = -u + (u + K - \lambda)v, \tag{6.13}$$

$$\epsilon\frac{\mathrm{d}v}{\mathrm{d}\tau} = u - (u + K)v. \tag{6.14}$$

6.1.2 Phase plane analysis

To visualise $u(\tau)$ and $v(\tau)$, we construct a phase plane.

The u-nullcline is given by

$$v = \frac{u}{u + K - \lambda}, \tag{6.15}$$

and the v-nullcline is given by

$$v = \frac{u}{u + K}. \tag{6.16}$$

We also consider the directions of trajectories in each region of the plane. We note that

$$\frac{\mathrm{d}v}{\mathrm{d}\tau} = \frac{1}{\epsilon}\left(u - (u + K)v\right), \tag{6.17}$$

and ϵ is small. Hence, unless $u - (u + K)v$ is small, which corresponds to being close to the v-nullcline, then v changes in time very quickly.

We can therefore consider the directions of trajectories in the phase plane, draw the trajectory starting from the initial condition $((u(0), v(0)) = (1, 0))$, and identify fast and (relatively) slow parts of the trajectory. The resulting phase plane is shown in the left panel of Figure 6.1.

From this phase plane, u and v can also be plotted as functions of τ; this is shown in the right panel of Figure 6.1.

6.1.3 Analytic solution

In Figure 6.1, it can be seen that there are two parts of the dynamics: initially, there is a fast part of the dynamics, as u remains approximately constant and v changes quickly. Then, there is a slow part of the dynamics, as both u and v decrease. We can consider these two parts separately, and solve the non-dimensional Michaelis-Menten model (i.e. equations (6.13) and (6.14)) in each part. When we consider the fast part, we will undertake a perturbation analysis (we consider small times that are only a small perturbation away from $t = 0$).

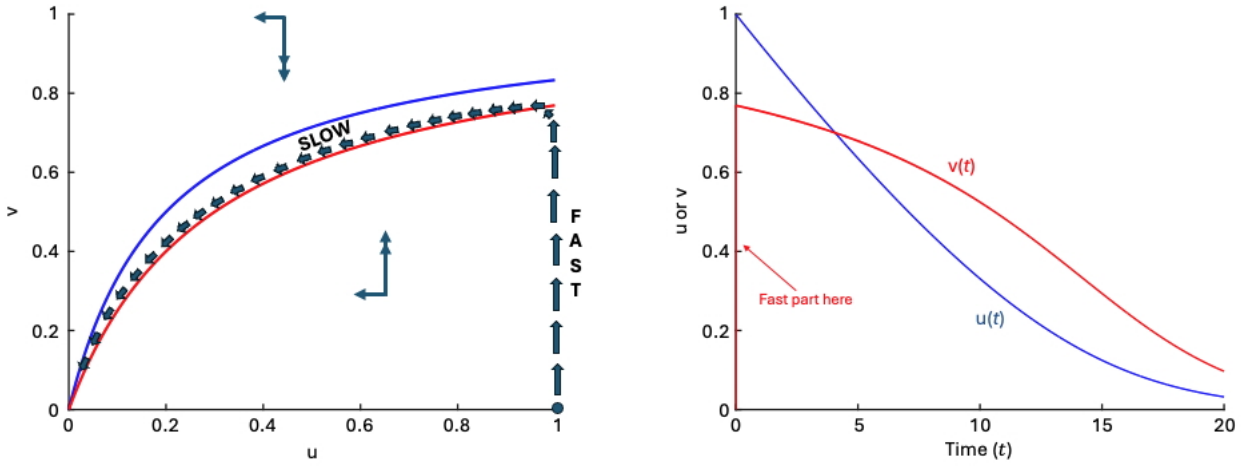

Figure 6.1: Left: Phase plane analysis for the non-dimensionalised Michaelis-Menten model. The blue line is the u-nullcline and the red line is the v-nullcline. The thick arrows mark out the trajectory starting from the initial condition $(u(0), v(0)) = (1, 0)$. Right: Temporal evolution of u and v based on the phase plane in the left panel. Parameter values: $K = 0.3$, $\lambda = 0.1$, $\epsilon = 0.001$.

Slow part

First, we focus on the dynamics arising for most of the trajectory (i.e. for all τ that are not exceptionally small).

We wish find a solution to the system

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = -u + (u + K - \lambda)v, \tag{6.18}$$

$$\epsilon\frac{\mathrm{d}v}{\mathrm{d}\tau} = u - (u + K)v. \tag{6.19}$$

To do this, we first apply the "Quasi steady state approximation" (QSSA).

Definition. Quasi steady state approximation. Consider an equation of the form $\epsilon\frac{\mathrm{d}v}{\mathrm{d}\tau} = g(u, v)$, where $\epsilon \ll 1$. Unless $\frac{\mathrm{d}v}{\mathrm{d}\tau}$ is large, then $g(u, v) \approx 0$. N.b. This part of the dynamics is relatively slow (compared to the fast part of the dynamics where $\frac{\mathrm{d}v}{\mathrm{d}\tau}$ is large), hence the name "quasi steady state".

The QSSA therefore allows us to set the left-hand-side of equation (6.19) to zero in the slow part of the trajectory (so that $\frac{\mathrm{d}v}{\mathrm{d}\tau}$ is not large). This implies that

$$v \approx \frac{u}{u + K}. \tag{6.20}$$

This can then be substituted into equation (6.18) to find

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = \frac{-\lambda u}{u + K}, \tag{6.21}$$

which can be solved via separation of variables to give

$$u + K\ln(u) = -\lambda\tau + C_1, \tag{6.22}$$

where C_1 is a constant. This is an implicit equation for $u(\tau)$. We know that, at the start of the slow part of the trajectory (so that τ is small, but not small enough to be in the fast part of the trajectory), $u \approx 1$ (since u does not change during the fast part of the trajectory; this can be seen from the phase plane, and will be confirmed in the following subsection when we analyse the fast part of the trajectory). As a result, $C_1 \approx 1$.

In summary, in the slow part of the trajectory, we have

$$u + K\ln(u) = -\lambda\tau + 1, \tag{6.23}$$

$$v = \frac{u}{u+K}. \tag{6.24}$$

Fast part

We now consider the fast part of the trajectory. Again,

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = -u + (u + K - \lambda)v, \tag{6.25}$$

$$\epsilon\frac{\mathrm{d}v}{\mathrm{d}\tau} = u - (u+K)v. \tag{6.26}$$

To "zoom in" on small values of τ, where the trajectory is fast, we rescale time by setting

$$\tau = \epsilon\sigma. \tag{6.27}$$

Then, $\frac{\mathrm{d}}{\mathrm{d}\tau} = \frac{1}{\epsilon}\frac{\mathrm{d}}{\mathrm{d}\sigma}$, and the model becomes

$$\frac{\mathrm{d}u}{\mathrm{d}\sigma} = -\epsilon u + \epsilon(u + K - \lambda)v, \tag{6.28}$$

$$\frac{\mathrm{d}v}{\mathrm{d}\sigma} = u - (u+K)v. \tag{6.29}$$

Now, since ϵ is small, we neglect terms of $\mathcal{O}(\epsilon)$ to obtain

$$\frac{\mathrm{d}u}{\mathrm{d}\sigma} = 0, \tag{6.30}$$

$$\frac{\mathrm{d}v}{\mathrm{d}\sigma} = u - (u+K)v. \tag{6.31}$$

This implies that u is constant (which we know to be one in the fast trajectory, from the initial condition $(u(0), v(0)) = (1, 0)$), so that

$$\frac{\mathrm{d}v}{\mathrm{d}\sigma} = 1 - (1+K)v. \tag{6.32}$$

Solving this equation via separation of variables, applying the initial condition $v(0) = 0$, and then re-inserting $\sigma = \frac{\tau}{\epsilon}$, gives rise to the final solution in the fast part of the trajectory,

$$u = 1, \tag{6.33}$$

$$v = \frac{1}{1+K}\left(1 - e^{-(1+K)\frac{\tau}{\epsilon}}\right). \tag{6.34}$$

Uptake of substrate

One quantity of interest is the rate at which the substrate is taken up in order to be transformed into the product. For the vast majority of the trajectory, we are in the slow part, so that (as we found following application of the QSSA),

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = \frac{-\lambda u}{u+K}. \tag{6.35}$$

This equation is in the form

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = -r(u). \tag{6.36}$$

The function $r(u)$ is referred to as the *uptake function*.

Note (please ignore if you have not taken Differential Equations 2): For those of you who have taken Differential Equations 2 (A6), the analysis undertaken above corresponds to an application of boundary layer theory (the fast part of the trajectory is inside the boundary layer near $\tau = 0$). In our analysis, we have identified the leading order inner and outer solutions, where the inner (fast) solution as $\sigma \to \infty$ matches the outer (slow) solution as $\tau \to 0$. Please note: Differential Equations 2 is not a pre-requisite for this course, and so understanding this analysis in these terms is not necessary here.

6.1.4 Extension to more complex biochemical reactions

In general, uptake of a substrate can occur with multiple enzymes simultaneously. In such a scenario, the non-dimensionalised rate equations have the form

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = f(u, v_1, v_2, \ldots, v_n), \tag{6.37}$$

$$\epsilon_i \frac{\mathrm{d}v_i}{\mathrm{d}\tau} = g_i(u, v_1, v_2, \ldots, v_n), \tag{6.38}$$

for $i = 1, 2, \ldots, n$.

In such a system, the quasi steady state approximation can be applied multiple times to give $g_i(u, v_1, v_2, \ldots, v_n) \approx 0$. These n equations can be solved simultaneously to find each v_i in terms of u, leading to a single equation describing the uptake of the substrate

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = f(u, v_1(u), v_2(u), \ldots, v_n(u)). \tag{6.39}$$

$$\tag{6.40}$$

This is again in the form

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = -r(u), \tag{6.41}$$

$$\tag{6.42}$$

where $r(u)$ is the uptake function.

6.2 Summary

Now that we have reached the end of this chapter, in addition to the techniques learnt in previous chapters you should be able to:

- Apply the Law of Mass Action to turn a set of biochemical reactions into a system of ODEs.
- Identify the "slow" and "fast" parts of trajectories (where relevant) when undertaking a phase plane analysis.
- Apply the QSSA to find and solve equations for the "slow" part of the trajectory arising in the Michaelis-Menten model (or similar biochemical reaction models).
- Rescale time, find and solve equations for the "fast" part of the trajectory arising in the Michaelis-Menten model (or similar biochemical reaction models).
- Find the uptake function describing the uptake of substrate in (the slow part of) a biochemical reaction model.

Chapter 7: Neuron signalling and excitable systems

In this chapter, we will study a model of neuron signalling as an example of an excitable system. In terms of the mathematical techniques that we will use, this chapter represents an application of phase plane analysis (as in Chapter 4).

In the body, information is transmitted in the form of electric currents through neurons. For example, the eyes receive signals from the outside world; this information is then transmitted through a network of neurons to the muscles, which then respond to the signals.

A neuron receives electric current through its dendrites (the receiving parts of the neuron). If sufficient electric current is received, then the neuron "fires". The current passes down the outer membrane of the axon (a long cylindrical tube that extends from the neuron). The current moves from one neuron to another via synapses and into the dendrites of the receiving neuron. Again, in the receiving neuron, if sufficient current is received through its synapses, then it "fires" and the process continues.

In the 1950s, Alan Hodgkin and Andrew Huxley studied the squid giant axon, which is one of the components of the nervous system in squid that allow them to move quickly through the water via short and fast movements. They performed experiments and constructed a mathematical model to explain the mechanisms underlying the initiation and transmission of electric current in the squid giant axon. In 1963, Hodgkin and Huxley received the Nobel Prize in Physiology or Medicine for their research.

The Hodgkin-Huxley model is made up of four differential equations. They solved these equations numerically and showed that the resulting solutions exhibited the same behaviour that was observed in experiments.

Rather than studying the full system, in this chapter we will focus on the Fitzhugh-Nagumo model, which captures the essence of the Hodkin-Huxley model in only two equations. Richard Fitzhugh constructed the model and Jinichi Nagumo built an electric circuit to reproduce the model's behaviour (both in the 1960s).

7.1 Fitzhugh-Nagumo model

We will study a specific form of the Fitzhugh-Nagumo model, given by

$$\epsilon \frac{\mathrm{d}v}{\mathrm{d}t} = f(v) - w + I_a, \tag{7.1}$$

$$\frac{\mathrm{d}w}{\mathrm{d}t} = bv - \gamma w. \tag{7.2}$$

In this model, v is the voltage (specifically, the difference in voltage between the axon and the surrounding medium; this quantity is measurable in experiments), I_a is the applied current, w

is a composite variable (arising from reducing the Hodkin-Huxley model to only two equations), and b and γ are positive-valued parameters. The parameter $\epsilon \ll 1$, so that v is a "fast" variable and w is a "slow" variable.

The function $f(v)$ is assumed to be of the form

$$f(v) = v(a - v)(v - 1), \tag{7.3}$$

in which $0 < a < 1$.

We consider two forms of this model: $I_a = 0$ and $I_a > 0$.

In principle, it is possible to study the Fitzhugh-Nagumo model analytically in a similar fashion to the Michaelis-Menten model, by considering the fast and slow parts of the trajectory (and applying the QSSA on the slow part of the trajectory). Here, however, we focus on undertaking phase plane analyses.

First, we introduce the concept of an excitable system, and demonstrate that this system can be excitable when $I_a = 0$. Second, we demonstrate that periodic trajectories may occur when $I_a > 0$.

7.1.1 Excitability ($I_a = 0$)

Definition. A system is *excitable* if a small perturbation from a stable steady state yields a fast return to the steady state, but a larger perturbation yields a large excursion before returning to the steady state.

In the scenario in which $I_a = 0$, we draw the nullclines and consider the directions of the trajectories in each region of the plane, as we did in the context of models of two interacting species in Chapter 4. There are two cases, as shown in Figure 7.1: i) If b/γ is large, there is only a single steady state (at $(0, 0)$), as shown in the top left panel of Figure 7.1; ii) If b/γ is small, there are three steady states, as shown in the top right panel of Figure 7.1.

We restrict our attention to the scenario in which b/γ is large and the only steady state is at $(0, 0)$.

As shown by the blue arrows in the middle left panel of Figure 7.1, if a small perturbation is made from the steady state in the positive v direction, then the system simply returns to the steady state.

However, if a large perturbation is made from the steady state in the positive v direction, then the system undergoes a large excursion before returning to the steady state (middle right panel of Figure 7.1). The system is *excitable*. This represents a situation in which the neuron fires; the resulting plot of v as a function of time is shown in the bottom left panel of Figure 7.1 and can be compared against experimental data.

We note here the existence of fast and slow dynamics, as in the Michaelis-Menten model. When the trajectory is far from the v-nullcline, then the system evolves quickly (and v changes quickly), whereas when the trajectory is close to the v-nullcline, then the system evolves more slowly.

7.1.2 Periodic firing ($I_a > 0$)

In the previous analysis, when $I_a = 0$ the phase planes illustrate that no periodic solutions are possible: trajectories simply tend to a steady state.

In contrast, when $I_a > 0$, periodic solutions may be possible. A scenario in which $I_a > 0$ leads to similar nullclines to those seen in Figure 7.1, but the cubic v nullcline is shifted vertically. This leads to four cases, shown in Figure 7.2.

In the case shown in the top right panel of Figure 7.2, in which the steady state lies between the turning points of the v nullcline, periodic trajectories arise. These are illustrated in Figure 7.3.

7.2 Summary

Now that we have reached the end of this chapter, in addition to the techniques learnt in previous chapters you should be able to:

- Use phase plane analyses to identify whether or not a system is excitable.
- Identify slow and fast parts of trajectories in relevant systems.
- Use phase plane analyses to assess whether periodic trajectories may be possible.

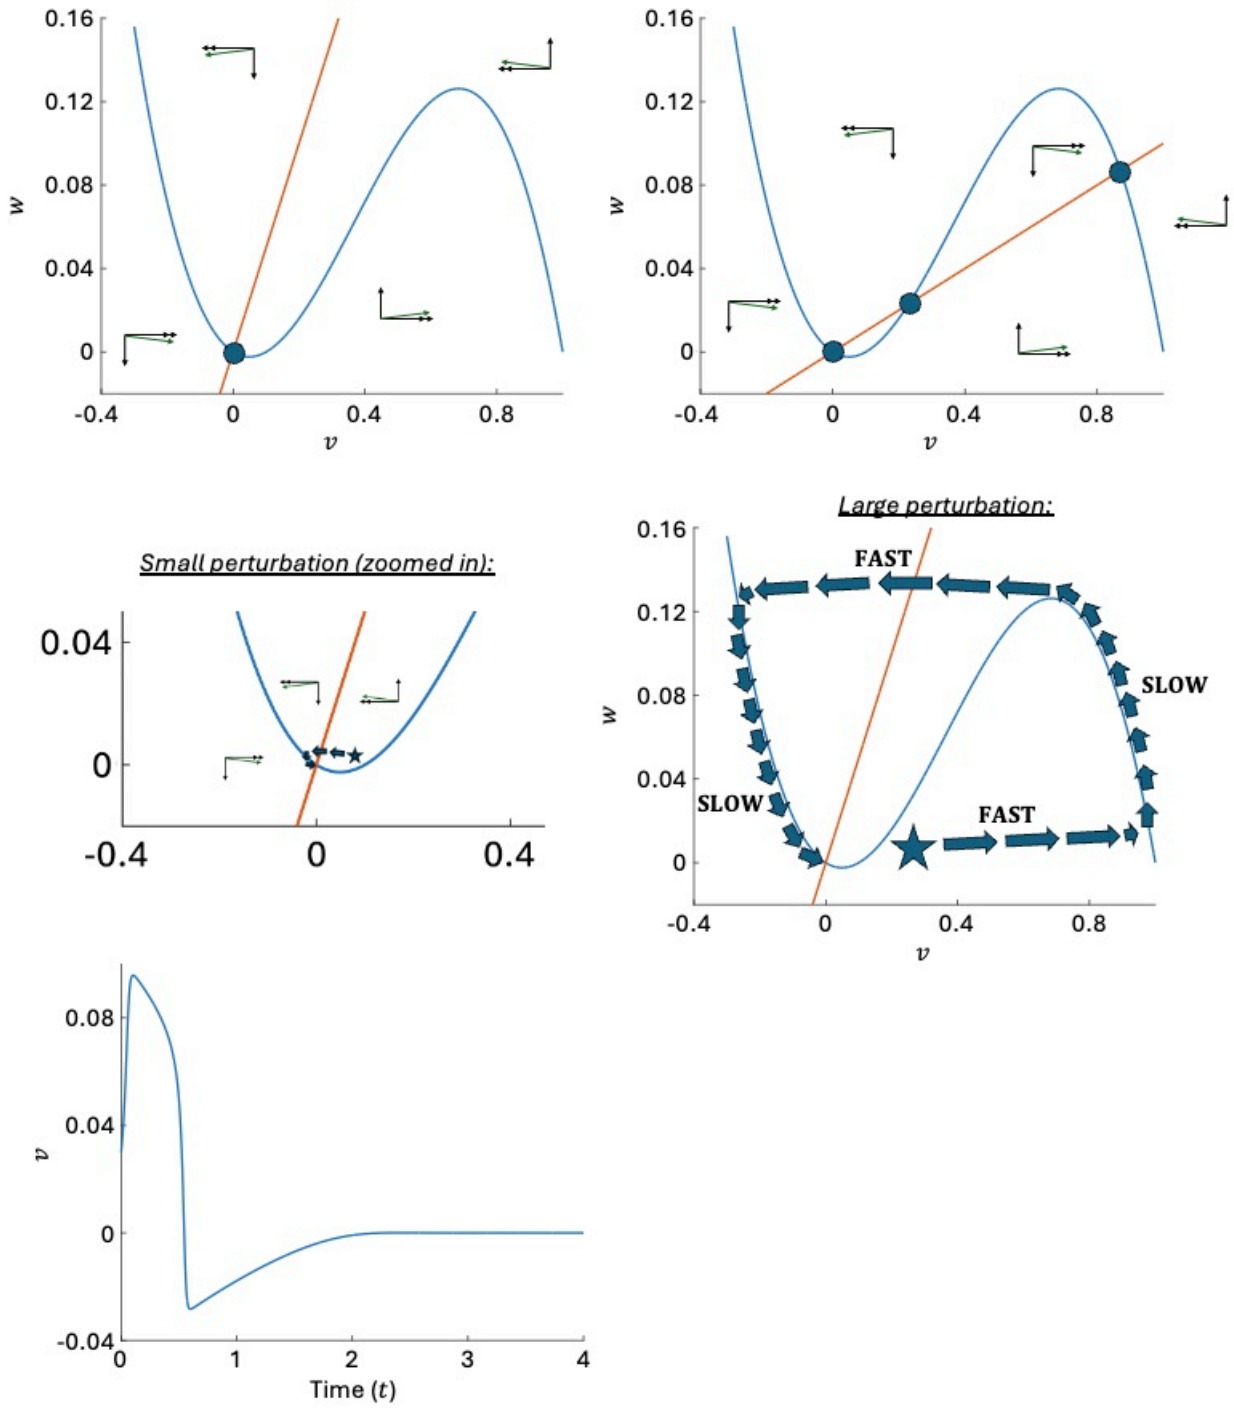

Figure 7.1: The Fitzhugh-Nagumo model when $I_a = 0$. Top left: Nullclines and directions of trajectories in each region when b/γ is large enough that there is only one steady state. Top right: Nullclines and directions of trajectories in each region when b/γ is small enough that there are three steady states. Middle left: When b/γ is large, a small perturbation from the steady state in the positive v direction leads to a trajectory that returns to the steady state quickly. Middle right: When b/γ is large, a large enough perturbation from the steady state in the positive v direction leads to a trajectory that involves a large excursion before returning to the steady state. Fast and slow parts of the trajectory can be seen. The middle left and right panels illustrate that the system is excitable. Bottom left: The temporal evolution of $v(t)$ corresponding to the scenario illustrated in the middle right panel (i.e. with a sufficiently large perturbation in the positive v direction from the steady state at $(0, 0)$). Parameter values: $\epsilon = 0.01$, $a = 0.1$, $\gamma = 1$ and $I_a = 0$. In the top right panel, $b = 0.1$; in all other panels, $b = 0.5$.

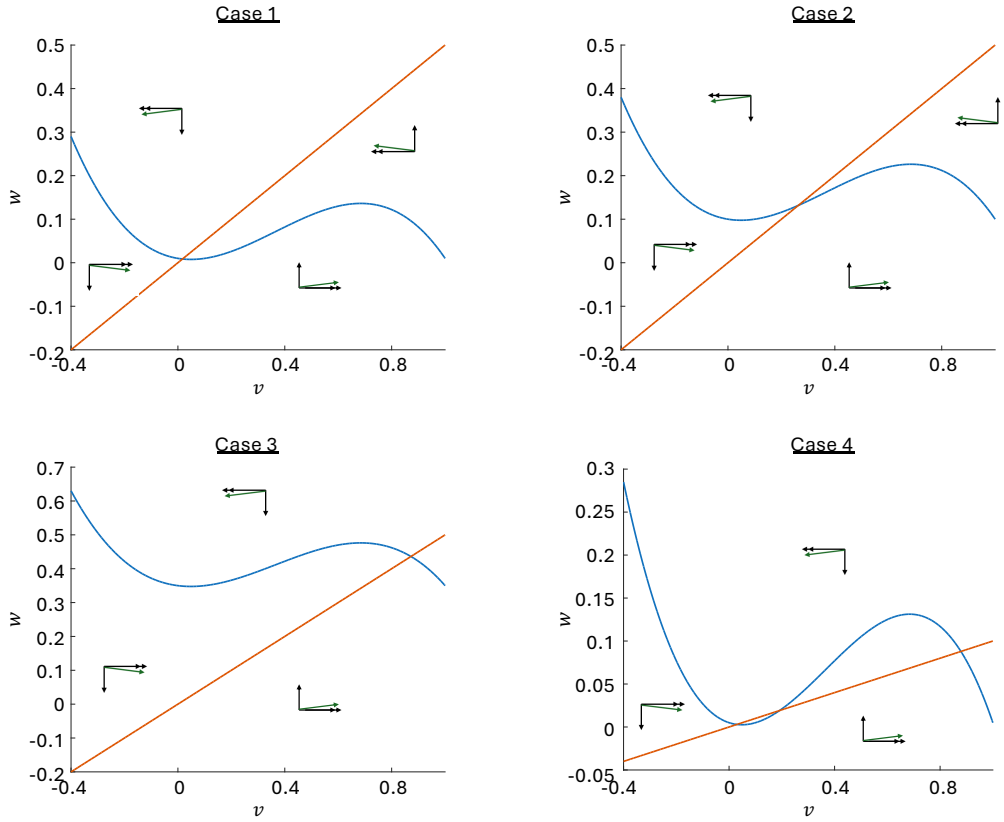

Figure 7.2: Nullclines in the Fitzhugh-Nagumo model for different values of $I_a > 0$. Top left: Case 1 ($I_a = 0.01$ and $b = 0.5$). Top right: Case 2 ($I_a = 0.1$ and $b = 0.5$). Bottom left: Case 3 ($I_a = 0.35$ and $b = 0.5$). Bottom right: Case 4 ($I_a = 0.005$ and $b = 0.1$). Other parameter values (which apply to all panels): $\epsilon = 0.01$, $a = 0.1$ and $\gamma = 1$.

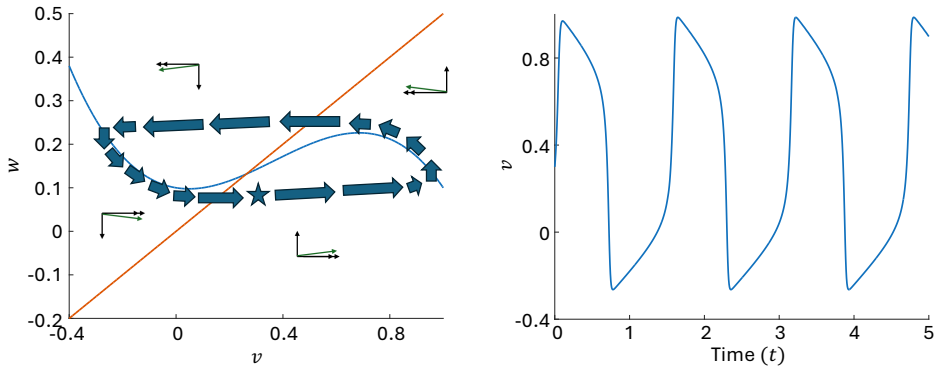

Figure 7.3: Periodic trajectories in the Fitzhugh-Nagumo model (case 2 in Figure 7.2). Left: Phase plane analysis. Right: Corresponding plot of $v(t)$ as a function of t.

Final remark: Congratulations on reading to the end of these notes! I hope that you enjoyed this course, and that you are interested in learning more mathematical biology in future. Best of luck in the Part A examinations! ☺