

Part II Mathematical Biology - Lent 2017

Lecturer: Prof. Julia Gog (jrg20@cam.ac.uk)

(this version September 2019 - some typos corrected)

Contents

0	Introduction	3
1	Deterministic systems, no spatial structure	5
1.1	Single population models	5
1.1.1	Simple birth and death models	5
1.1.2	Delay models	8
1.1.3	Populations with age structure	19
1.2	Discrete time	25
1.2.0	Revision: 1-D stability in difference equations (maps)	25
1.2.1	The logistic map	26
1.2.2	Higher order discrete systems	30
1.3	Multi-species models	35
1.3.0	Revision: 2-D stability in continuous time	35
1.3.1	Competition models	36
1.3.2	Predator-prey models	40
1.3.3	Chemical kinetic models	46
1.3.4	Epidemic models	50
1.3.5	Excitable systems	59

2	Stochastic systems	64
2.0	Preliminaries	64
2.0.0	Revision: discrete probabilities and generating functions	64
2.0.1	Why bother?	65
2.0.2	The first step	66
2.1	Discrete population sizes	68
2.1.1	Single populations	68
2.1.2	Extinction	77
2.1.3	Multiple populations	80
2.2	Continuous population sizes	84
2.2.1	Fokker-Planck for a single variable	84
2.2.2	Multivariate Fokker-Planck	88
3	Systems with spatial structure	97
3.0	Preliminaries	97
3.1	Diffusion and growth	99
3.1.1	Linear diffusion in finite domain	99
3.1.2	Linear diffusion in infinite domain	103
3.1.3	Nonlinear diffusion	112
3.2	Travelling waves in reaction-diffusion systems	116
3.2.0	General F	116
3.2.1	Fisher's equation	117
3.2.2	Bistable systems	126
3.3	Spatial instabilities	130
3.3.1	Chemotaxis	130
3.3.2	Turing instabilities	135

0 Introduction

Acknowledgements

Much of the content and the starting point for these notes comes from Prof. Peter Haynes's version of the notes from 2012. Any typos or errors are, however, likely to be the fault of this author.

Practicalities

These are intended as an archived version of the notes from the Part II Mathematical Biology course, as lectured by me (Julia Gog) in Lent 2017, within Part II of the Mathematical Tripos, University of Cambridge. For Part II students in future years, these notes might be useful as an extra resource perhaps for reading ahead over the summer, but please do not use them as a replacement for attending lectures or working from the present lecturer's resources. Lectures are good, you should go to them: you'll always learn a bit more, see something in a different way or build more intuition for what is really going on.

There are exercises in green boxes throughout these notes. These were set in lectures and intended as being additional to the usual examples sheets. These exercises should be fairly doable without any help from a supervisor, and I recommend doing them as you work through the notes. There exist a full set of 'solutions': you *might* find these wherever you found these notes: these can be used to check you are working along the right lines, or to see what was intended.

For anyone who uses the notes, I hope these notes are interesting and useful. Most of all, I hope they inspire you to explore mathematical biology further. If you have any comments or corrections, please do email me: jrg20@cam.ac.uk. I might not respond immediately (or at all, sorry!), but feedback will be useful in updating and amending these for future use.

Preparation for this course

Part II Dynamical Systems is 'helpful' for parts of this course but certainly not essential. If you did not do Dynamical Systems, then it might be wise to do a little revision of parts of Ia Differential Equations: stability of equilibria of discrete and continuous time systems (Jacobians, saddles/focus/node, phase-plane diagrams). Indeed '*Ordinary Differential Equations*' by Robinson (see schedules for Ia Differential Equations) chapters 32 and 33 ('coupled nonlinear equations' and 'ecological models') will put you right on track for this course. The middle part of the course on stochastic systems will use some knowledge from Ia Probability, including generating functions. It would be a

good idea to revise separable solutions from Ib Methods for the last part of the course on diffusion.

Interesting reading

None of these are essential to follow this course, but should be of interest:

- J.D. Murray *Mathematical Biology (3rd edition)* (see schedules) - the classic text on mathematical biology, covering a range of applications
- D. Neal *Introduction to Population Biology* - much overlap with this course in mathematical detail, but explores the biological principles in rather more depth and includes many real examples. Should be completely readable by you during or after this course.
- *Mathematics is biology's next microscope, only better; Biology is mathematics' next physics, only better* - article by Joel E. Cohen in PLoS Biology 2004 DOI: 10.1371/journal.pbio.0020439

1 Deterministic systems, no spatial structure

1.1 Single population models

1.1.1 Simple birth and death models

The simplest model?

Let $x(t)$ be population size as a function of time t . Assume that the number of offspring produced per individual per unit time is a constant $b > 0$. Similarly assume that the death rate (number of deaths per unit time per individual) is a constant $d > 0$.

$$x(t + \delta t) = x(t) + b x \delta t - d x \delta t$$

Divide by δt and take the limit as $\delta t \rightarrow 0$.

$$\frac{dx}{dt} = (b - d)x = rx \quad \text{where} \quad r = b - d.$$

Solution is $x(t) = x_0 e^{rt}$, where $x(0) = x_0$, so the population grows indefinitely if $r > 0$ and decays towards zero (implying extinction) if $r < 0$.

Exercise 1: In the case when $r < 0$, find the half-life of the population

Exercise 2: Actually, this simple model is pretty good for invasions of new populations. Suppose a new disease is discovered and there are 1000 cases last week and 1500 cases this week, roughly when did the disease first appear?

Note that in a deterministic system, only the difference between b and d matters, e.g. $b = 21$, $d = 20$ gives entirely the same dynamics as $b = 1,000,001$, $d = 1,000,000$. These will differ in an analogous stochastic model (the ones with higher rates will fluctuate wildly).

Birth and death rates depend on population size Rather than constant, allow the number of offspring per individual per unit time to depend on population size, $a(x)$, and similarly the death rate $b(x)$. Then we have:

$$\frac{dx}{dt} = [a(x) - b(x)]x$$

Again, only the difference between birth and death rates matter in the deterministic system.

Typically, one might expect the birth rate per capita to decrease and/or death rate to increase for very large population size, as resources become scarce. For example here we could take the birth rate to be constant ($b(x) = B$) and the death rate to be proportional to population size ($d(x) = D x$):

$$\frac{dx}{dt} = [B - D x]x$$

By rescaling the population size and renaming parameters, we have the *logistic equation*:

$$\frac{dx}{dt} = \alpha x(1 - x)$$

Exercise 3: Find the rescaling.

For $x < 1$, births outnumber the deaths and the population grows. For $x > 1$, the opposite occurs and the population shrinks. The equilibrium population size (scaled) is one.

The logistic model

$$\frac{dx}{dt} = \alpha x(1 - x)$$

This is easy to solve:

$$\int \frac{1}{x(1-x)} dx = \int \frac{1}{x} + \frac{1}{1-x} dx = \log \left| \frac{x}{1-x} \right| + C = \alpha t$$

So putting $x = x_0$ at $t = 0$ we have:

$$\frac{x}{1-x} = \frac{x_0}{1-x_0} e^{\alpha t}$$

Which rearranges to:

$$x = \frac{x_0 e^{\alpha t}}{(1-x_0) + x_0 e^{\alpha t}}$$

Reassuringly, our steady population size is there: $x_0 = 1$ gives $x(t) = 1$. Also, it is always sensible to check zero initial conditions: $x_0 = 0$ gives $x(t) = 0$.

Exercise 4: show that the solution to the logistic equation can be rewritten for some t_0 as:

$$x = \begin{cases} \frac{1}{2} + \frac{1}{2} \tanh\left(\frac{1}{2}\alpha(t - t_0)\right) & \text{for } x_0 < 1 \\ \frac{1}{2} + \frac{1}{2} \coth\left(\frac{1}{2}\alpha(t - t_0)\right) & \text{for } x_0 > 1 \end{cases}$$

Note that for positive initial population size ($x_0 > 0$), $x \rightarrow 1$ as $t \rightarrow \infty$ (from above if $x_0 > 1$ and from below if $x_0 < 1$). There is a stable equilibrium, achieved for all positive initial conditions. The zero equilibrium $x = 0$, is unstable.

One-dimensional stability recap

Consider:

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

The steady-states are the values of x^* for which $f(x^*) = 0$. These may be interchangeably referred to as equilibria, fixed points, steady states or constant solutions. Stability is determined by behaviour near the fixed point, which can be found by linearisation around x^* . Set $x(t) = x^* + \epsilon(t)$. Then:

$$\frac{dx}{dt} = \frac{d\epsilon}{dt} = f(x^* + \epsilon) = \underbrace{f(x^*)}_{=0} + \epsilon f'(x^*) + \underbrace{\mathcal{O}(\epsilon^2)}_{\text{ignore}}$$

Hence:

$$\frac{d\epsilon}{dt} \simeq f'(x^*)\epsilon \quad \text{which has solution} \quad \epsilon(t) \simeq \epsilon_0 \exp[f'(x^*)t].$$

So ϵ , the perturbation away from x^* grows if $f'(x^*) > 0$ (unstable) and shrinks if $f'(x^*) < 0$ (stable).

In practice, just check sign of f' at fixed points. For simple biological models, this can usually be done easily by plotting f .

Exercise 5: check stability of the fixed points of the logistic model

1.1.2 Delay models

So far, we have $x'(t)$ depending on $x(t)$, i.e. the instantaneous current population size. Of course this is not always realistic. For example, offspring are not really produced instantaneously, there may be a significant gestation period, or time for eggs to hatch. Even then, new offspring may need further time to mature to adulthood, before they can in turn produce offspring. So, we might want $x(t)$ to denote adults, and births and/or deaths may depend on the population size at some past time point. In physiological models, there is often some form of delay, for example heart rate does not respond instantly to exercise. In biochemical signalling, there can be many steps between a trigger and effect, which can sometimes be modelled relatively simply as a time lag.

end of
lecture 1

Mathematically, this leads to *delay-differential equations* (DDEs). Here is an example, the *Hutchinson-Wright* equation, which can be viewed as an extension to the logistic equation:

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

where the delay time T is a new parameter in the model (assume $T > 0$, note $T = 0$ was logistic equation).

We can analyse its dynamics with much the same ideas as before: find the interesting fixed points and look at their stability by considering a small perturbation. Clearly $x(t) = 1$ is still the non-trivial steady state. Now set $x(t) = 1 + \epsilon(t)$ and sub in:

$$\begin{aligned} \frac{d\epsilon}{dt} &= \alpha(1 + \epsilon(t))(-\epsilon(t - T)) \\ \frac{d\epsilon}{dt} &= -\alpha\epsilon(t - T) + O(\epsilon^2) \end{aligned}$$

And drop $O(\epsilon^2)$ from here.

This is still linear, so reasonable to seek a solution of the form $\epsilon = \epsilon_0 e^{st}$:


$$s = -\alpha e^{-sT} \quad (1)$$

We would like know the solutions for s . We see that if $T = 0$, then this just returns $s = -\alpha$, which corresponds to the stable fixed point of the logistic equation. If $T > 0$, then we need to look a bit more carefully.

First we might reasonably seek real s solutions. Rearranging:

$$sT e^{sT} = -\alpha T$$

Consider the shape of the LHS as a function of sT . It has a single minimum at $sT = -1$ when the LHS is equal to $-e^{-1}$. So, there are negative real roots for $\alpha T < e^{-1}$ and no real roots otherwise.



If we look at the solution near $sT = 0$, for small αT , the gradient is approximately 1, so we have $sT \approx -\alpha T$ so this is a continuation of the solution $s = -\alpha$, which is what we would have got with the logistic equation.

So far we have only considered real roots for s , but we might (correctly) suspect there could be complex roots of 1 for s . What would this mean? Our perturbation would follow $\epsilon_0 e^{st}$, so a complex solution would just give a solution that grows or decays but with oscillations (think back to complimentary functions in second order linear ODEs). We are now interested in the sign of the real part of s . If $Re(s) > 0$ we say it is unstable, if $Re(s) < 0$ we say it is stable. It is not usually possible to solve explicitly for s , but we can see now that it would be sensible to find when stability might change, i.e. $Re(s) = 0$.

Now let's seek complex roots of (1) by setting $s = \sigma + i\omega$ (where σ and ω are the real and imaginary parts of s). Sub in:

$$\sigma + i\omega = -\alpha e^{-s\sigma} e^{-is\omega} = -\alpha e^{-s\sigma} [\cos(\omega T) - i \sin(\omega T)]$$

Take real and imaginary parts:

$$\begin{aligned} \sigma &= -\alpha e^{-\sigma T} \cos(\omega T) && \text{real part} \\ \omega &= +\alpha e^{-\sigma T} \sin(\omega T) && \text{imaginary part} \end{aligned}$$

Seek a solution with $\sigma = 0$. Things simplify quite a bit:

$$\begin{aligned} 0 &= -\alpha \cos(\omega T) \\ \omega &= +\alpha \sin(\omega T) \end{aligned}$$

Squaring and adding gives $\omega^2 = \alpha^2$ so $\omega = \pm\alpha$. This is not too surprising: we should expect complex conjugate pairs of solutions. Could limit calculations to $\omega > 0$ if it helped, and just remember the complex conjugates are also there. In any case, subbing in either solution to the second equation, gives the same outcome:

$$\sin(\alpha T) = 1 \quad \text{so } \alpha T = \frac{\pi}{2}, \frac{5\pi}{2}, \frac{9\pi}{2}, \dots$$


and no need to worry about negative solutions, as both $\alpha > 0$ and $T > 0$. So thinking about increasing T up from zero, we have a complex root switch real part sign many times. We are interested in the first one: $\alpha T = \pi/2$.

This is optional, but as this is the first example, we will check that we really do have stable solutions when $0 < T < \frac{\pi}{2\alpha}$. It turns out to be sensible to split into two cases according to modulus of ω

- For $|\omega| > \alpha$, from considering modulus in the equation for the real part we see that $\exp(-\sigma T) > 1$ hence $\sigma < 0$.
- For $|\omega| \leq \alpha$, $|\omega T| \leq \alpha T < \pi/2$ so $\cos(\omega T) > 0$. From the equation for the real part, we see $\sigma < 0$ again.

So either way, we have negative σ and hence stable solutions. Note we have not actually found any values for s , but we have shown they will have negative real part in this range.

Numerical simulation is consistent with $0 < T < \frac{1}{\alpha e}$ solutions decay exponentially to the fixed point; for $\frac{1}{\alpha e} < T < \frac{\pi}{2\alpha}$ solutions decay and oscillate to the fixed point, and for $T > \frac{\pi}{2\alpha}$ the solution is unstable and heads to a cycle. This is typical: delay-differential equation models often lead to oscillatory solutions.




(See Mathematica: Delay Logistic Equation)

Under the carpet

Treat this note as starred. If you are happy with DDEs already, then skip it. If you are concerned that something might have been swept under the carpet here, you are right, so read on. What we have actually done is

- Found where there are real solutions for s and shown they are negative.
- For first range of T showed that any solution for s has negative real part
- Found all the values of $T > 0$ where a solution has real part zero

We can actually know more about the solutions for s of $sTe^{sT} = -\alpha T$ if we read up on ‘Lambert W-Functions’. There are many solutions. Mathematica has a built-in function that can be used to give them numerically, and we can plot their real part as a function of T (set $\alpha = 1$ for simplicity).



$\sigma = \text{Re}(s)$ against T for $\alpha = 1$ for top 10 solutions
(same plot each side, just different vertical scale)

The root with the largest real part (top line on graph) actually corresponds to that largest real solution to start with, and you can see the sharp change of direction as it becomes complex. The vertical zoom-in on the right shows more clearly that successive (pairs of) solutions pass upwards into positive σ .

For $T = 0$ we only had one value of s (namely $s = -\alpha$). This was enough to determine the linear behaviour of any small perturbation: $\epsilon = \epsilon_0 \exp(-\alpha t)$, and we'd just need to put in the appropriate single constant ϵ_0 . Now, for a perturbation, we should specify not just $\epsilon(0)$ but also $\epsilon(t)$ for the interval $t \in [-T, 0]$. And the resultant dynamics will be as a sum of these types of solution with different s :

$$\epsilon(t) = \sum_i a_i e^{s_i t}$$

where the α_i are determined by the initial conditions. The s_i with the largest real part will end up dominating as t increases.

So, actually we have got the dynamics right from the simple approach we first took: the real solution dominated when it existed, then we had an oscillatory but decaying solution until we found the lowest T where things could lose stability. For math bio, treat this simple approach as sufficient.

Be a bit careful when rescaling DDEs

This is just a word of caution about rescaling delay differential equations with respect to time. In short, you must remember to rescale any time lag also. In long, we will use the above as an example:

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

There are two parameters here, α and T . It is tempting to try and get rid of α by rescaling time. We set $\hat{t} = \alpha t$ to cancel out with the α :

$$\frac{d\epsilon(t)}{dt} = \alpha \frac{d\epsilon}{d\hat{t}} = -\alpha\epsilon(t - T)$$

so

$$\frac{d\eta(\hat{t})}{d\hat{t}} = -\epsilon(t - T) = -\eta(\alpha(t - T)) = -\eta(\hat{t} - \alpha T)$$

and finally:

$$\frac{d\eta(\hat{t})}{d\hat{t}} = -\eta(\hat{t} - \alpha T).$$

So really we have not eliminated α but we have compounded our two parameters to a single parameter combination αT .

In general, be aware that the lag needs to rescale with time also. It is not usual in practice to write out all of these steps. It is usually acceptable to reuse the original variable name (ϵ here), but the change was made explicit just this once.

Exercise 6: Find the equivalent of equation (1) for this rescaled DDE. (It turns out to be slightly different, but it ought to give us the same conditions for stability.)


DDE Example: Blowflies

This example stems from classic experimental work by Nicholson and others in the 1950s on the Australian sheep blowfly. Populations of flies were kept in the lab and population size was tracked over time, showing some quite spectacular fluctuations despite the available food and other external factors being kept steady. The full life cycle of these flies is a few weeks (eggs, larval stages, then adult). Mathematical biologists have modelled this using delay differential equations.

The unusual thing here is that the number of eggs produced by adult flies is very strongly regulated by population size, in fact we assume that the per capita rate of egg production is exponentially decreasing with population size. This means that the total number of eggs produced is no longer monotonic increasing with population size, but now is unimodal with a peak at $x = x_0$.

Per capita egg production if population has size x : $P e^{-x/x_0}$
 Total egg production from population of size x : $p(x) = P x e^{-x/x_0}$

end of
lecture 2



Now assume that eggs turn into adults after a delay t_D , and that the per capita death rate δ does not depend on population size, and we have our model:

$$\frac{dx(t)}{dt} = \underbrace{P x(t - t_D) \exp\left(-\frac{1}{x_0} x(t - t_D)\right)}_{\text{new eggs at time } t - t_D} - \underbrace{\delta x(t)}_{\text{death}} \quad (2)$$

Note that $x(\text{blah})$ denotes x evaluated at blah (as opposed to multiplied by), and similarly for \hat{x} below.

This system has four parameters¹: P , t_D , x_0 and δ . By strategic rescaling, this can be reduced to two. As usual, we can rescale time to absorb a parameter, and here we'll choose the death rate. Set $\hat{t} = \delta \times t$ to make $\frac{d}{dt} = \delta \frac{d}{d\hat{t}}$.

We can also rescale x to tidy the exponent: set $\hat{x} = x/x_0$. We will consider \hat{x} as a

¹and as usual in mathematical biology, assume everything is positive unless you have a good reason to think otherwise

function in our rescaled time: \hat{t} . So we are setting this

$$x(t) = x_0 \hat{x}(\hat{t}) = x_0 \hat{x}(\delta t),$$

so in particular

$$x(t - t_D) = x_0 \hat{x}(\delta \times (t - t_D)) = x_0 \hat{x}(\hat{t} - \delta t_D).$$

Hence making these changes and also dividing through by x_0 :

$$\frac{d\hat{x}(\hat{t})}{d\hat{t}} = \frac{P}{\delta} \hat{x}(\hat{t} - \delta t_D) e^{-\hat{x}(\hat{t} - \delta t_D)} - \hat{x}(\hat{t})$$

From this we can see that the model really only depends on two parameter combinations. Set $a = \delta t_D$ and $b = Pt_D$ (turns out to be sensible to think of them both as increasing in t_D , so use these and then drop the hats to get the system in a suitable form to analyse:

$$\frac{dx(t)}{dt} = \frac{b}{a} x(t - a) e^{-x(t-a)} - x(t) \quad (3)$$

Exercise 7: get from equation 2 to equation 3 (without using notes!)

To find any equilibria, we solve for $x(t) = x^*$ where x^* is constant:

$$0 = \frac{b}{a} x^* e^{-x^*} - x^*$$

Which gives $x^* = \log \frac{b}{a}$ as the non-trivial solution. Assume $b > a$ so that this solution is positive.

For stability, look at dynamics close to this fixed point, i.e. ϵ small²:

$$x(t) = x^* + \epsilon y(t)$$

Sub this into equation 3:

$$\begin{aligned} \epsilon y'(t) &= \frac{b}{a} (x^* + \epsilon y(t - a)) \underbrace{e^{-x^*}}_{=a/b} e^{-\epsilon y(t-a)} - (x^* + \epsilon y(t)) \\ &= (x^* + \epsilon y(t - a)) e^{-\epsilon y(t-a)} - (x^* + \epsilon y(t)) \\ &= (x^* + \epsilon y(t - a)) (1 - \epsilon y(t - a)) - (x^* + \epsilon y(t)) + \mathcal{O}(\epsilon^2) \end{aligned}$$

²We don't really need ϵ and $y(t)$ (lectures just used $\epsilon(t)$). Just need some function of time which is assumed to be small, but sometimes it is easier to put ϵ in as a constant to explicitly keep track of what is small.





And the order 1 terms cancel (which is reassuring, as it was supposed to be a fixed point) and then taking just order ϵ :

$$y'(t) = (1 - x^*)y(t - a) - y(t)$$

Exercise 8: Continue from here to get to:

$$\begin{aligned}\sigma &= (1 - x^*)e^{-a\sigma} \cos(a\omega) - 1 \\ \omega &= (1 - x^*)e^{-a\sigma} (-\sin(a\omega)).\end{aligned}$$

Exercise 9: Show the system is stable for $b < e^2 a$.



(See Mathematica: Blowflies, and keep $b \gg a$)


Physiological Example: Breathing

This is a simple model of respiration (breathing) where we focus on one particular function of breathing: to remove carbon dioxide from the blood. We assume there is a feedback system where the volume of breath depends positively on the level of carbon-dioxide in the blood.

$$\begin{aligned} V &= \text{air inhaled (per breath)} \\ C(t) &= \text{concentration of carbon dioxide in the blood at time } t \end{aligned}$$

And we use a Hill³ function to model how V depends on C :

$$V(C) = V_{\max} \frac{C^m}{a^m + C^m}$$



Hill function for various m : red $m = 1/2$, black $m = 1$, blue $m = 2$. They all pass through the same point at $C = a$: half of the maximum V_{\max} .

However, there will be a lag between the level of carbon dioxide being detected and the time until the breath volume is adjusted. This will involve a series of chemical reactions, and communication with the brainstem, but we do not need to know all the details, just that there is some time difference, call it T . Also assume that the amount of carbon dioxide breathed out is proportional both to its current concentration and to the breath volume, with multiplicative constant b . Finally, assume that carbon dioxide in the blood is added at a constant rate p (from other physiological processes around the body). This gives us the following equation:

³Archibald Hill 1886-1977: Trinity mathematician and Nobel Prize winning Physiologist, among many other things: http://www.nobelprize.org/nobel_prizes/medicine/laureates/1922/hill-bio.html

$$\begin{aligned}\frac{dC}{dt} &= \underbrace{+p}_{\text{CO}_2 \text{ in to blood}} - \underbrace{b C(t) V (C(t-T))^m}_{\text{breathed out}} \\ &= p - b C(t) V_{\max} \frac{C(t-T)^m}{a^m + C(t-T)^m}\end{aligned}$$

end of
lecture 3

And we go ahead and rescale in the same way as before. We can pick up a good rescaling for C by tidying the Hill function: set $\hat{C} = C/a$. Again, we can choose to eliminate something else by rescaling time, so go for making that first term into 1. Set $\hat{t} = \frac{p}{a}t$ so that $\frac{d}{dt} = \frac{p}{a} \frac{d}{d\hat{t}}$.

$$\frac{d\hat{C}}{d\hat{t}} = 1 - \frac{abV_{\max}}{p} \hat{C}(\hat{t}) \frac{\hat{C}(\hat{t} - pT/a)^m}{1 + \hat{C}(\hat{t} - pT/a)^m}$$

Again, note how the time lag also is rescaled as we are now working with \hat{C} as a function of \hat{t} . Now we can see there are essentially two new parameter combinations emerging here, $\tau = pT/a$ and $\alpha = abV_{\max}/p$, so work in terms of those and also we can drop the hats at this point⁴:

$$\frac{dC}{dt} = 1 - \alpha C(t) \frac{C(t-\tau)^m}{1 + C(t-\tau)^m}$$

This could be tidied a little further by defining $f(x) = x^m/(1+x^m)$:

$$\frac{dC}{dt} = 1 - \alpha C(t) f(C(t-\tau))$$

Seek a constant equilibrium solution:

$$0 = 1 - \alpha C^* f(C^*)$$

$$1 = \alpha C^* f(C^*)$$

The function f increases from zero to one. While we can't write down an explicit solution for this (at least not for general m in the Hill function), we can easily see that the right-hand side is unboundedly monotonically increasing, starting from zero, so there will be a unique solution C^* . Also, as $f < 1$ we know that C^* will satisfy $C^* > 1/\alpha$, which will turn out to be useful later.

For stability, as usual we set $C = C^* + \epsilon y(t)$.

⁴in lectures, we just reused T , but done with τ here

$$\begin{aligned}
 \epsilon y'(t) &= 1 - \alpha (C^* + \epsilon y(t)) f(C^* + \epsilon y(t - \tau)) \\
 &= 1 - \alpha (C^* + \epsilon y(t)) (f(C^*) + \epsilon y(t - \tau) f'(C^*)) + \mathcal{O}(\epsilon^2) \\
 &= 1 - \alpha [C^* f(C^*) + \epsilon y(t) f(C^*) + C^* \epsilon y(t - \tau) f'(C^*)] + \mathcal{O}(\epsilon^2) \\
 &= \underbrace{1 - \alpha C^* f(C^*)}_{=0} - \epsilon \alpha [y(t) f(C^*) + C^* y(t - \tau) f'(C^*)] + \mathcal{O}(\epsilon^2)
 \end{aligned}$$

And so to order ϵ (in other words, linearising):

$$y'(t) = -\underbrace{\alpha f(C^*)}_A y(t) - \underbrace{\alpha C^* f'(C^*)}_B y(t - \tau)$$

for some constants A and B .

Exercise 10 : show that

$$A = \frac{1}{C^*}, \quad B = \left(1 - \frac{1}{\alpha C^*}\right) \frac{m}{C^*}.$$

and check that B is positive

So, in essence:

$$y'(t) = -A y(t) - B y(t - \tau)$$

Now we have an equation which is linear in y , though it is still a delay equation, so we try a solution of the form $y = e^{st}$:

$$s = -A - B e^{-s\tau}$$

We can immediately see that for $\tau = 0$ that $s = -A - B$, so C^* is stable. For more general parameters, we explore this as usual by setting $s = \sigma + i\omega$ (where we always take σ and ω to be real).

$$\begin{aligned}
 \sigma &= -A - B e^{-\sigma\tau} \cos \omega\tau \\
 \omega &= B e^{-\sigma\tau} \sin \omega\tau
 \end{aligned}$$

Note that the equations are symmetric in $\pm\omega$, which is not surprising: the roots for s should be in complex conjugate pairs. We could restrict attention to $\omega > 0$.

There's a few things we could argue through now. With a bit of work, we could show for small τ (actually not very small needed), we should show $\sigma > 0$ impossible, so still stable.

Actually, we might be interested in the shape parameter of the Hill function m . We can show that if m is small, then B will be small (need to be a bit careful as m is implicitly

in c^* , but that can be bounded). If B is small, we can also see $\sigma > 0$ impossible. Now think about larger m , as it increases from small. To try and find a boundary when the real part goes through zero (which should give us the edge of instability), set $\sigma = 0$:

$$\begin{aligned} 0 &= -A - B \cos \omega\tau \\ \omega &= B \sin \omega\tau \end{aligned}$$

Which rearranges to:

$$-A\tau \tan(\omega\tau) = \omega\tau \quad (4)$$

$$B^2 = A^2 + \omega^2 \quad (5)$$

Equation 4 has a root for $\omega\tau$ in $(\pi/2, \pi)$, and its exact value will depend on $A\tau$. Call it $g(A\tau)$. There are other roots further on, but this lowest one turns out to be the one of interest. Equation 5 gives an implicit expression that must be satisfied by A , B and τ :

$$B^2 = A^2 + \tau^{-2}g(A\tau)^2$$

Exercise 11: Consider varying the parameter m . Show that this boundary for instability means that:

$$1 + \frac{\pi^2 C^{*2}}{4\tau^2} \leq m^2 \left(1 - \frac{1}{\alpha C^*}\right)^2 \leq 1 + \frac{\pi^2 C^{*2}}{\tau^2}$$

And in particular, this means that this critical m is greater than 1.


1.1.3 Populations with age structure

So far, we have just been considering the population size (usually x). However, when thinking about birth and death rates, it would often be important to consider the age of individuals. Now, set $n(a, t)$ to be the number of individuals at time t who are age a ⁵. The *total* population at time t can be found by integrating over all ages:

$$N(t) = \int_0^\infty n(a, t) da$$

Now we set up birth and death rates as functions of age. Let $b(a)$ be the birth rate *from* individuals of age a . Let $\mu(a)$ be the death rate *of* individuals of age a .

⁵Strictly, this is a density function in a . So this should really be $n(a, t)\delta a$ is the number of individuals aged between a and $a + \delta a$. We don't usually need to say all of this though.



Example for $b(a)$ and $\mu(a)$. Birth rate might be highest from a particular age group, while the death rate may increase with age.

Now start to build the equations that govern n . Consider how a chunk⁶ of population ages as time increases from t to $t + \delta t$. They will age by δt , but a small number might have die in that time:

$$n(a + \delta t, t + \delta t) = n(a, t) - \mu(a)\delta t n(a, t) + \mathcal{O}(\delta t^2) \quad (6)$$

The left hand side can be expanded by Taylor series, again to first order in δt :

$$n(a + \delta t, t + \delta t) = n(a, t) + \delta t \frac{\partial n}{\partial a}(a, t) + \delta t \frac{\partial n}{\partial t}(a, t) + \mathcal{O}(\delta t^2)$$

Sub this in to equation 6, cancel the $n(a, t)$, divide by δt and take limit $\delta t \rightarrow 0$:

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

Or we usually write:

$$\boxed{\frac{\partial n}{\partial t} + \frac{\partial n}{\partial a} = -\mu(a) n(a, t)} \quad (7)$$

With an initial condition in time, this is most of the story, but we also need a boundary condition in age, i.e. the newborns that appear at age zero:

$$\boxed{n(0, t) = \int_0^\infty b(a) n(a, t) da} \quad (8)$$

⁶c.f. fluids courses where we start with a blob of fluid. In fact this derivation is very similar if we replace space with age. The material derivative in fluids ($\frac{D}{Dt} = \frac{\partial}{\partial t} + u \cdot \frac{\partial}{\partial x}$) is like the left hand side in equation 7 if we think about age rather than space, and our blob moves (ages) with velocity 1 in time. Note: don't try to think too hard about a fluids equivalent for the the boundary condition at age 0: that would be taking the analogy too far!

The age ∞ might look worrying, but in practice for any sensible model, the birth rate times population size is zero for a greater than some age, or at least mathematically the product tends to zero sufficiently fast for this integral to always be sensible.

Wave-like solutions

(Note: would be a good idea to have read through this once and understood the ideas, but otherwise treat it as starred)

The left hand side of equation 7 may remind you of characteristic or wave solutions from earlier courses (such as IB Methods). This works very easily for $\mu(a) = 0$ and ignoring age boundary condition. Any general function of $n(a, t) = g(a - t)$ solves $\frac{\partial n}{\partial t} + \frac{\partial n}{\partial a} = 0$. This would be just a fixed population age distribution, just drifting upwards in age with time (which is fine only for no deaths).

We can extend this to account for the age-dependent death rate, and it can easily be checked that equation 7 is solved by:

$$n(a, t) = \text{Exp} \left[- \int_0^a \mu(s) ds \right] g(a - t)$$

The exponential term represents the probability of *surviving* to age a , where s is a dummy variable for age in the integration. We don't know much about g yet. This again is a wave-like solution, with some decay with age to account for deaths.

If we are given an initial condition i.e. $n(a, 0) = n_0(a)$ then that is enough to specify $g(x)$ for $x > 0$:

$$n(a, 0) = n_0(a) = e^{-\int_0^a \mu(s) ds} g(a)$$

$$\Rightarrow g(a) = \frac{n_0(a)}{e^{-\int_0^a \mu(s) ds}}$$

Effectively, this is working out the number of births at time a ago, by consider those age a and scaling it up to account for the proportion that died before time zero.

To determine $g(x)$ for $x < 0$, we should use the newborn boundary condition (equation 8), which gives:

$$g(-t) = \int_0^\infty b(a) e^{-\int_0^a \mu(s) ds} g(a - t) da$$

And then split integral range at $a = t$ so when $a - t$ changes sign so we can use the initial condition in the second term:

$$\begin{aligned}
 g(-t) &= \int_0^t b(a) e^{-\int_0^a \mu(s) ds} g(a-t) da + \int_t^\infty b(a) e^{-\int_0^a \mu(s) ds} g(a-t) da \\
 &= \int_0^t b(a) e^{-\int_0^a \mu(s) ds} g(a-t) da + \int_t^\infty b(a) e^{-\int_0^a \mu(s) ds} \left[e^{+\int_0^{a-t} \mu(s) ds} n_0(a-t) \right] da \\
 &= \int_0^t b(a) e^{-\int_0^a \mu(s) ds} g(a-t) da + \int_t^\infty b(a) e^{-\int_{a-t}^a \mu(s) ds} n_0(a-t) da
 \end{aligned}$$

The first term is the births from those who were born after $t = 0$. The second term represents offspring from those individuals who were part of the initial condition population. They were age $a - t$ initially at $t = 0$, then they needed to survive from age $a - t$ to age a (which is the exponential term), and then give birth.

Under all sensible models with sensible initial conditions⁷, this second term tends to zero as time increases (original population dies or doesn't contribute to birth rate). Then in principle $g(-t)$ can be determined by using the values for $g(x)$ for $-t < x < 0$. Essentially, this means that for large enough time, we can ignore the details of the initial condition (i.e. bin the second term above). It is this irrelevance of initial condition that we are exploiting in the next section.

(end of starred section)

Normal mode solutions

Without worrying about initial conditions (see previous section), we look for a 'normal mode' solution to equations 7 and 8. Set:

$$n(a, t) = r(a) e^{\gamma t}$$

Where r is the general shape of the population distribution, and the whole thing is scaled up or down in time, according to the exponent γ .

Sub it into equation 7:

$$\begin{aligned}
 \gamma r(a) e^{\gamma t} + r'(a) e^{\gamma t} &= -\mu(a) r e^{\gamma t} \\
 r'(a) &= -(\mu(a) + \gamma) r(a)
 \end{aligned}$$

Which we can solve by integrating with respect to a :

$$r(a) = r(0) e^{-\gamma a} e^{-\int_0^a \mu(s) ds}$$

⁷'sensible' would certainly demand that the mean number of offspring is finite

The same exponential of an integral appears as in the previous section. Again, it is the probability of surviving from birth to age a . So we now have:

$$n(a, t) = r(0) e^{\gamma(t-a)} e^{-\int_0^a \mu(s) ds}$$

We can substitute the normal mode solution with this expression for r into equation 8:

$$r(0) e^{\gamma t} = \int_0^\infty b(a) r(0) e^{\gamma(t-a)} e^{-\int_0^a \mu(s) ds} da$$

Cancelling, and defining ϕ :

$$1 = \int_0^\infty b(a) e^{-\gamma a} e^{-\int_0^a \mu(s) ds} da := \phi(\gamma)$$

So if we can find a gamma that satisfies $\phi(\gamma) = 1$, we have a valid normal mode solution. How does ϕ depend on γ ? We can see that it must be a decreasing function of γ and that it can be made as small as we like by taking very large γ , and as large as we like by making γ more negative⁸. So, there will be a unique root for $\phi(\gamma) = 1$. The question is whether it is for positive or negative gamma, which determines whether our population will grow or decay (shrink). Check $\phi(0)$ to see which side of 1 it is:


$$\phi(0) = \int_0^\infty b(a) e^{-\int_0^a \mu(s) ds} da$$

This is just the probability of being still alive at age a , the birth rate at that age, and then integrated over all ages. This must be the average number of births from one individual, i.e. the mean number of offspring from one individual.

$$\begin{aligned} \phi(0) > 1 &\Rightarrow \text{solution in } \gamma > 0 &\rightarrow \text{growth} \\ \phi(0) < 1 &\Rightarrow \text{solution in } \gamma < 0 &\rightarrow \text{decay} \end{aligned}$$

⁸ Actually integral usually blows up for some $\gamma < 0$, but all we need to know is that there is a root to $\phi(\gamma) = 1$

end of
lecture 4



Some example curves for $\phi(\gamma)$. The red corresponds to a population that will decay, and the blue to a population that will grow.

Exercise 12: Try this method for the case where births and deaths don't actually depend on age, i.e. $b(a) = b$ and $\mu(a) = \mu$.

- *Find an expression for $\phi(\gamma)$.*
- *Find $\phi(0)$ and check it makes biological sense.*
- *Find γ that solves $\phi(\gamma) = 1$ (can't always do this explicitly, but can in this case)*
- *Find the condition for population growth/decay.*

1.2 Discrete time

There is often good reason to consider time in discrete steps in biological models (as opposed to the continuous time and differential equations in previous sections). For example, it may be natural to consider some particular period such as a year as that is the natural life cycle for many species (annual plants, insects such as the monarch butterfly), or even if lifecycle is longer, it may make sense to think in terms of discrete step of a year (e.g. hibernating mammals). Another natural timescale is a day. Another use of discrete time is when consider a model with generations of population, and it might be natural to model the timestep as per generation.

However, we know (e.g from Part Ia Differential Equations) that discrete systems⁹ can behave in a complicated way. In particular a first-order system can be ‘chaotic’.

1.2.0 Revision: 1-D stability in difference equations (maps)

Consider the system where the number of individuals in the next generation is a function of the number in the current generation. Call this function f , so

$$x_{n+1} = f(x_n). \quad (9)$$

A fixed point x^* is a solution to

$$x^* = f(x^*).$$

To analyse stability, look at perturbations from the fixed point, i.e. set $x_n = x^* + \epsilon_n$ (where we think of ϵ_0 as small). Subbing this in to equation 9:

$$\begin{aligned} x^* + \epsilon_{n+1} &= f(x^* + \epsilon_n) \\ x^* + \epsilon_{n+1} &= f(x^*) + f'(x^*) \epsilon_n + \mathcal{O}(\epsilon^2) \\ \epsilon_{n+1} &= f'(x^*) \epsilon_n + \mathcal{O}(\epsilon) \end{aligned}$$

So the perturbation just grows geometrically, with ratio $f'(x^*)$.

In practice, we don't usually present all of this, just go from finding x^* to considering the modulus of $f'(x^*)$:

$$\begin{array}{ll} |f'(x^*)| < 1 & x^* \text{ is a stable fixed point} \\ |f'(x^*)| > 1 & x^* \text{ is an unstable fixed point} \end{array}$$

⁹or difference equations, or maps: means the same thing


For the higher order discrete systems (see below), it is possible to come up with analogous tests for $f(x_n, x_{n-1}, \dots)$, but in practice it is usual to just do the perturbation explicitly.

1.2.1 The logistic map

You will have encountered this example in other courses: it is probably the most studied discrete map. However, you should now note that it was first proposed as a population model, and the need to understand biologically-motivated systems has driven forward the mathematical area of chaos theory. The basic equation is wonderfully simple, but as you probably already know, the resultant dynamics are wonderfully rich:

$$x_{n+1} = \alpha x_n(1 - x_n) = f(x_n)$$

This just contains one parameter¹⁰. To make sure the map is from the interval $[0, 1]$ to itself, we restrict attention to $\alpha \in [0, 4]$.



The logistic map, for various α .

The fixed points satisfy $x = f(x)$, so seek:

$$\begin{aligned} f(x) - x &= 0 \\ \alpha x(1 - x) - x &= 0 \\ -x[\alpha x - \alpha + 1] &= 0 \end{aligned}$$

And we see that $x = 0$ is always a fixed point. In addition, we have $x^* = 1 - 1/\alpha$, which is in range when $\alpha > 1$.

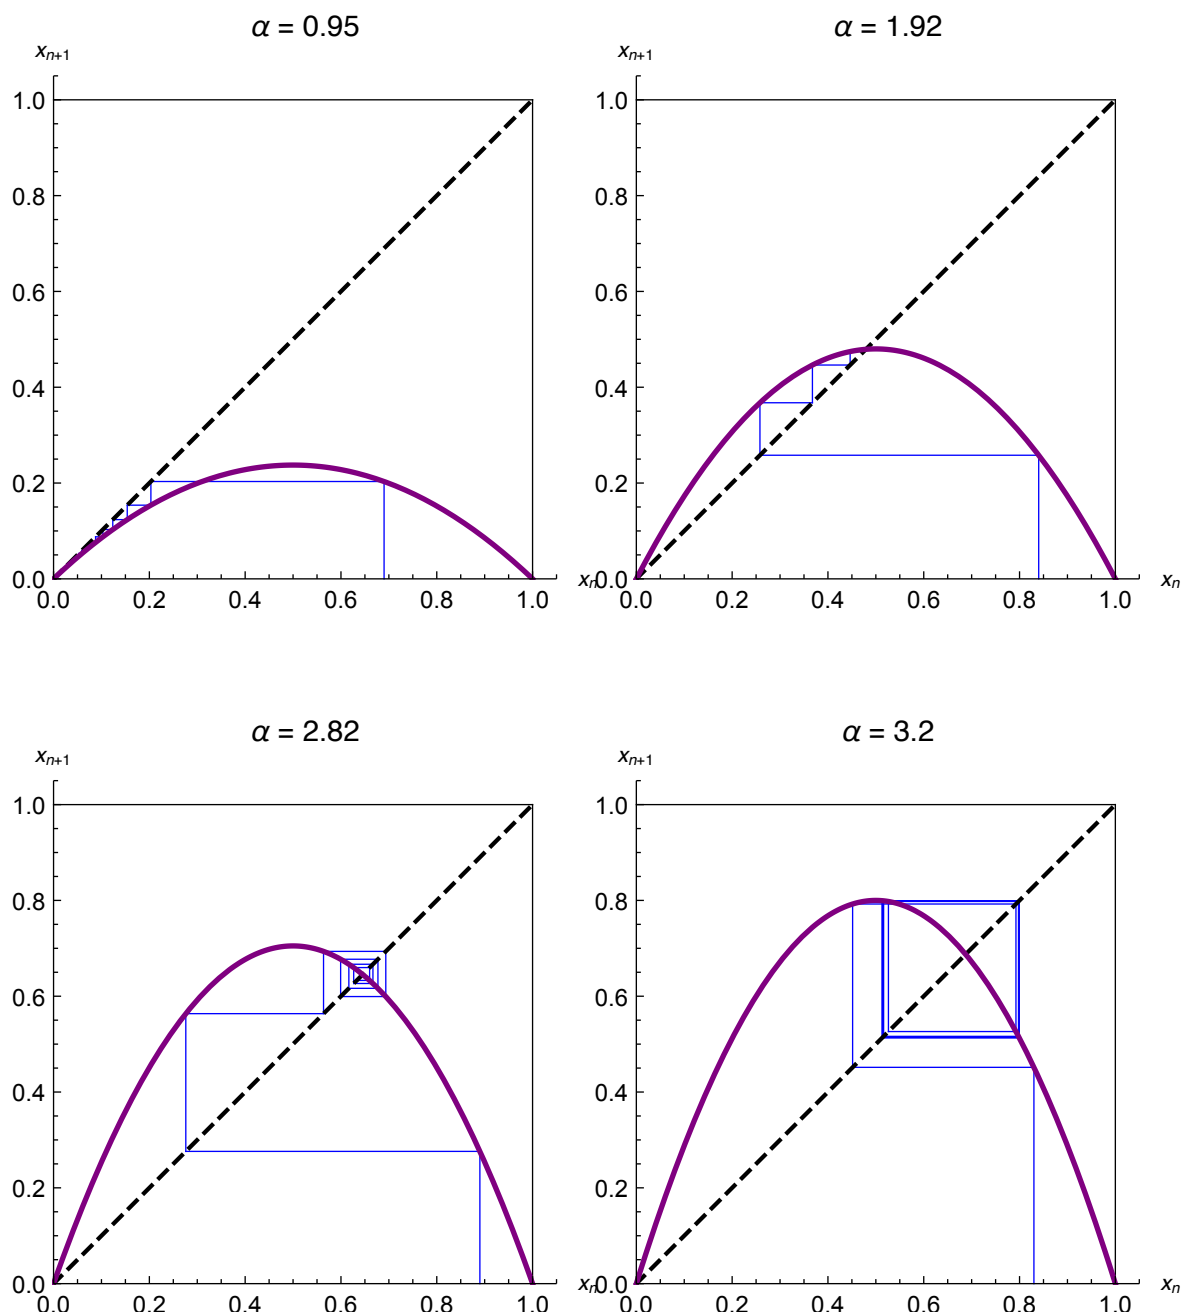
For stability, we have

$$f'(x) = \alpha(1 - 2x).$$

¹⁰we're calling it α here, but often it is r , μ or a and occasionally it is rescaled by a factor of 4

For $x^* = 1$, $f'(x^*) = \alpha$ so it is stable for $\alpha < 1$ and unstable for $\alpha > 1$.

For $x^* = 1 - 1/\alpha$, $f'(x^*) = 2 - \alpha$. So it is stable for $1 < \alpha < 3$ (doesn't exist for $\alpha < 1$) and unstable for $3 < \alpha$. We can break this down a little further to see when $f'(x^*)$ is positive and negative: it changes sign at 2, so $2 < \alpha < 3$ is stable, and nearby perturbations are jumping either side of the fixed point (and tending in). We can also note that $f'(x^*)$ goes through -1 at $\alpha = 3$, and this corresponds to a period-doubling bifurcation.



The logistic map, for various α . The blue lines show where successive iterates go, starting from an initial condition (line from the x-axis). See mathematica: logistic map.

At this point, it is a good idea to look at some numerical outputs, for example the cobweb figures here or see the mathematica file. What we can observe is that the

behaviour is straightforward to $\alpha < 3$. For $0 < \alpha < 1$ all trajectories head to the origin. For $1 < \alpha < 2$ all trajectories head into $x^* = 1 - 1/\alpha$ from one side (which side depends on initial conditions). For $2 < \alpha < 3$, trajectories still head to x^* but now in an oscillatory way, resulting in the blocky spirals in the cobweb diagram¹¹.

For α just a bit bigger than there, there looks to be a stable period-2 cycle, that is a pair of points where the map jumps from one to the other. Mathematically, x_1 and x_2 such that:

$$f(x_1) = x_2, \quad f(x_2) = x_1 \quad \text{and} \quad x_1 \neq x_2$$

Such x must satisfy $f^2(x) = x$, so we seek these solutions:

$$\begin{aligned} f^2(x) - x &= 0 \\ \alpha f(x)(1 - f(x)) - x &= 0 \\ \alpha^2 x(1 - x)(1 - \alpha x(1 - x)) - x &= 0 \\ -\alpha^3 x^4 + 2\alpha^3 x^3 - \alpha^2(1 + \alpha)x^2 + (\alpha^2 - 1)x &= 0 \end{aligned}$$

Then at this point it looks a bit hopeless, but there is actually a way forward: any solution to $f(x) = x$ will also be a solution to $f^2(x) = x$. So, we should be able to factorise out $f(x) - x$, in fact exactly the expression we solved to find the fixed points. So working carefully, we can factorise out $x[\alpha x - \alpha + 1]$:

$$-x[\alpha x - \alpha + 1](\alpha^2 x^2 - \alpha(\alpha + 1)x + (\alpha + 1)) = 0$$

We don't want these fixed point solutions though ($x_1 \neq x_2$ for a period-2 point), so can cancel these factors, leaving:

$$\alpha^2 x^2 - \alpha(\alpha + 1)x + (\alpha + 1) = 0 \tag{10}$$

And hence:

$$x_1, x_2 = \frac{1}{2\alpha} \left((1 + \alpha) \pm \sqrt{(1 + \alpha)(\alpha - 3)} \right)$$

These can be assigned to x_1 and x_2 either way around. Looking at the square root, these exist for $\alpha > 3$. When they appear at $\alpha = 3$ they start at $x = 2/3$, i.e. where the fixed point is ($x^* = 1 - 1/\alpha$).

Exercise 13: Check that $f(x_1) = x_2$ and $f(x_2) = x_1$

We can also check the stability of this period-2 cycle by considering f^2 (function applied twice) and checking if its derivative has modulus bigger than one or not, but there is a

¹¹actually this must be why they are called 'cobweb diagrams', from the way oscillatory fixed points look like sort of like orb webs from spiders, but maybe only like that in minecraft.

nice technique to make this simple in terms of algebra. Rather than going back to the big expression for $f^2(x)$, just use the chain rule:

$$\frac{d}{dx}f^2(x) = \frac{d}{dx}f(f(x)) = f'(f(x)) f'(x)$$


And at x_1 (or x_2), this gives is just $f'(x_1)f'(x_2)$. We have that $f'(x) = \alpha(1 - 2x)$ so

$$f'(x_1)f'(x_2) = \alpha(1 - 2x_1)\alpha(1 - 2x_2) = \alpha^2[1 - 2(x_1 + x_2) + 4x_1x_2]$$

and we can even read off the sum and product of x_1 and x_2 from the quadratic we solved to find them (equation 10).

Exercise 14: Show that the period-2 cycle of the logistic map becomes unstable at $\alpha = 1 + \sqrt{6} \approx 3.45$

What happens for higher values of α ? See Dynamical Systems course for more details, but in brief: there are a series of *period-doubling bifurcations*. We've found the first two, where the fixed point (a.k.a period-1) becomes unstable and a period-2 appears. Then the period-2 becomes unstable, and a period-4 will appear. This keeps going as α is increased, but these all accumulate at a certain value ($\alpha_\infty \approx 3.5699$). For $\alpha > \alpha_\infty$ there are windows of 'chaos', but also windows where things settle to stable periodic orbits.



Bifurcation diagram, focusing on $\alpha > 3$. For each value of α , this is made just by picking random initial conditions for x and then iterating the map forward many times, throw the first few hundred iterates away, and then plot the results. Then keep doing this for lots of values of α .

1.2.2 Higher order discrete systems

So far we have considered the case when the value of x in the next time step or generation only depends on the current value. For many situations in biology, this is not enough: the dynamics may depend on earlier times also. We could develop some general theory, much like the box above but now for $x_{n+1} = f(x_n, x_{n-1}, \dots, x_{n-p})$, but this is neither useful in practice nor informative. Instead we will explore a few examples, and see techniques which can be applied more generally.

Example: discrete version of breathing

Recall the physiological example from earlier where breathing is regulated to adjust for varying levels of carbon dioxide in the blood, with some lag. We can formulate a similar model in discrete time:

$$\underbrace{V_{n+1}}_{\text{Breath volume next step...}} = \underbrace{f(C_{n-k})}_{\text{... depends on CO}_2 \text{ } k \text{ steps ago}} = \alpha C_{n-k}$$

One could imagine a more general f , but here we just consider a linear example. Then the equation for carbon dioxide change in blood:

$$\underbrace{C_{n+1} - C_n}_{\text{Change in CO}_2 \text{ in blood}} = \underbrace{M}_{\text{CO}_2 \text{ added}} - \underbrace{\beta V_{n+1}}_{\text{breathed out}}$$

and as usual α , β and M are real, constant and positive. Note that this is actually not just different from earlier breathing model by being discrete: the model for CO_2 being breathed out is now just proportional to V and does not depend on C at all¹². The steps in n just represent some time step, e.g. minutes or breaths.

We can collapse this into a single variable:

$$C_{n+1} = C_n + M - \alpha\beta C_{n-k}. \quad (11)$$

Then seek constant solutions $C_n = C_*$:

$$C_* = C_* + M - \alpha\beta C_* \implies c_* = \frac{M}{\alpha\beta}.$$

Now investigate stability of this steady state by perturbing¹³ as $C_n = C_* + \delta_n$ and sub in to (11):

$$C_* + \delta_{n+1} = C_* + \delta_n + \underbrace{M - \alpha\beta C_*}_{=0} - \alpha\beta \delta_{n-k}$$

¹²Arguably, the earlier model is more sensible, where the higher the concentration in the blood, the higher the amount of CO_2 gets exchanged in lungs per volume of breath and expelled. This model is actually different, and chosen here just for convenient linearity later.

¹³not bothering with ϵ now, just think of δ_n as small

$$\delta_{n+1} - \delta_n + \alpha\beta\delta_{n-k} = 0$$

Normally we would then linearise in small δ , but this is already in right form here. Also, we could simplify slightly by making some compound parameter instead of $\alpha\beta$ but that is not essential.

Now we explore some different values of k . This is the time lag in steps between blood levels of CO_2 taking a value and the breathing volume adjusting. First, for $k = 0$ we simply have

$$\delta_{n+1} = (1 - \alpha\beta)\delta_n$$

and then it is clear that the steady state C_* is stable for $0 < \alpha\beta < 2$ and unstable for larger $\alpha\beta$.

Next, try $k = 1$:

$$\delta_{n+1} - \delta_n + \alpha\beta\delta_{n-1} = 0$$

To solve this linear difference equation, seek $\delta_n = p^n$ solutions:

$$p^2 - p + \alpha\beta = 0 \quad \Rightarrow \quad p_{\pm} = \frac{1}{2} \pm \sqrt{\frac{1}{4} - \alpha\beta}.$$

and the general solution is a linear combination of these geometric solutions:

$$\delta_n = Ap_+^n + Bp_-^n.$$

For $0 < \alpha\beta < \frac{1}{4}$, both p_{\pm} are real and $\in (0, 1)$, so p^n decays for both, and hence C_* is stable.

For $\frac{1}{4} < \alpha\beta$, both p_{\pm} are complex¹⁴. This actually doesn't change our approach very much: we still need to know when solutions grow or decay and hence $|p| < 1$ or otherwise:

$$p_{\pm} = \frac{1}{2} \pm i\sqrt{\alpha\beta - \frac{1}{4}} \quad \Rightarrow \quad |p|^2 = \left(\frac{1}{2}\right)^2 + \left(\alpha\beta - \frac{1}{4}\right) = \alpha\beta$$

hence for $\frac{1}{4} < \alpha\beta < 1$ the steady state is stable (and a perturbation decays in an oscillatory manner) and for $1 < \alpha\beta$ it is unstable. In summary, the steady state C_* is stable for $0 < \alpha\beta < 1$ and unstable for larger values. As expected, the longer lag decreases the range for stability (more parameter values are unstable).

This example worked out without too much difficulty as we could solve explicitly for p . This will be unlikely to work as we go up to higher order and therefore get something trickier than a quadratic to solve. Sticking with $k = 1$, we can explore an alternative strategy. For small $\alpha\beta$ we could see that our roots for p all had modulus less than one, so all we need to do is to imagine turning up $\alpha\beta$ until the first time a root goes

¹⁴The p_{\pm} are complex conjugates of each other. We will have real initial conditions for C and therefore δ . This will make A and B complex conjugates and this will give real δ_n for all n .

unstable. At the moment when this happens, a root will have modulus exactly one¹⁵. So, seek $p = e^{i\theta}$ for some $\theta \in [0, 2\pi)$.

$$p^2 - p + \alpha\beta = 0 : \quad e^{2i\theta} - e^{i\theta} + \alpha\beta = 0$$

Then taking real and imaginary parts:

$$\cos 2\theta - \cos \theta + \alpha\beta = 0 \quad (12)$$

$$\sin 2\theta - \sin \theta = 0 \quad (13)$$

Start with the simpler one, the imaginary part (13):

$$\sin 2\theta = \sin \theta \implies 2\theta = \theta + 2n\pi \quad \text{or} \quad 2\theta = (\pi - \theta) + 2n\pi$$

for $n \in \mathbb{Z}$. Restricting attention to $\theta \in [0, 2\pi)$:

$$\theta = 0 \quad \text{or} \quad \frac{\pi}{3}, \frac{3\pi}{3}, \frac{5\pi}{3}$$

The real part will supply the corresponding values of $\alpha\beta$:

$$\alpha\beta = 0 \quad \text{or} \quad 1, -2, 1$$


And as we are looking for the first positive¹⁶ $\alpha\beta$, we see this is $\alpha\beta = 1$. This corresponds to $\theta = \pi/3$ or $\theta = 5\pi/3$. Equivalently, this is

$$p_{\pm} = e^{\pm i\pi/3} = \frac{1}{2} \pm i \frac{\sqrt{3}}{2}$$

which matches up with our first approach. Note that $p^6 = 1$ for both p_+ and p_- , so small perturbations satisfy:

$$\delta_{n+6} = Ap_+^{n+6} + Bp_-^{n+6} = Ap_+^n + Bp_-^n = \delta_n$$

so have period 6 at the boundary (at least to linear order), and close to period 6 just above and below.




Typical solutions for $k = 1$ for various $\alpha\beta$ near to threshold for stability, starting from a perturbation away from steady state.

¹⁵Actually expect this to happen in complex conjugate pairs. We are implicitly assuming that roots are continuous in our model parameters (which does not seem too unreasonable in math bio). We are also assuming that the offending root(s) actually goes right *through* modulus one: technically it would be possible to just reach modulus one and then go back inside the unit circle. That would just be mean.

¹⁶actually the solution at zero corresponds to a root $p = 1$ exactly at $\alpha\beta = 0$, and this root edges just below one as we make $\alpha\beta$ small but positive. So this is moving into the region for stability, not moving out!

Example: Multi-generation model

Another set of problems which leads to higher order discrete systems is when the time steps are population generations and multiple generations need to be considered to find the next generation size. For example, a type of annual plant produces γ seeds in the summer and then dies. Those seeds stay in the ground over the winter. The next summer, each seed¹⁷ has probability σ of successfully germinating and growing into a new adult plant. Failing that, the seed will germinate the summer after that with probability τ . Assume they cannot germinate after that.



We turn this wordy description into equations¹⁸, accounting for number of adult plants at season n by x_n , seeds that have been waiting one year as $s_n^{(1)}$ and two years as $s_n^{(2)}$.

Then

$$\begin{aligned} x_n &= \sigma s_n^{(1)} + \tau s_n^{(2)} \\ s_n^{(1)} &= \gamma x_{n-1} \\ s_n^{(2)} &= (1 - \sigma) s_{n-1}^{(1)} \end{aligned}$$

Then we can go from this system in multiple variables to one in a single variable, and here the only sensible choice is x_n :

$$x_n = \sigma \gamma x_{n-1} + \tau(1 - \sigma) \gamma x_{n-2}. \quad (14)$$

Indeed, arguably it would be possible to go straight here from the description in words. Note $\gamma > 0$ and $\sigma, \tau \in [0, 1]$ as they are proportions.

Note that equation (14) is linear in x and the steady state solution is $x_* = 0$. In theory we now propose a perturbation, sub it in, linearise, but of course in this case we just get the same equation back again. So, directly go for $x_n = p^n$:

$$p^2 - \sigma \gamma p - \tau(1 - \sigma) \gamma = 0$$

¹⁷Think of this as proportion of seeds, as there are lots of seeds lots of seeds

¹⁸Indeed, this is half of the art of math bio in practice in research, except the wordy descriptions tend to be a lot more vague on crucial details and then there's a lot of decisions for the modeller to make.

$$p_{\pm} = \frac{\sigma\gamma}{2} \pm \frac{1}{2}\sqrt{\sigma^2\gamma^2 + 4\gamma\tau(1-\sigma)}$$

As the contents of the square root is positive, these roots are real. We can also see that the roots are positive and negative with the positive root having the larger modulus (i.e. $0 < -p_- < p_+$). Solutions will be of the form

$$x_n = A_1 p_+^n + A_2 p_-^n$$

and the first term will dominate as n increases. In fact, we just need to check p_+ for stability¹⁹.

Seeking $p_+ = 1$ and a little algebra we arrive at

$$\gamma[\sigma + (1 - \sigma)\tau] = 1$$

It is worth stepping back into original meaning of the parameters at this point. Considering this, the square bracket is the proportion of seeds that *ever* germinate, and the prefactor γ is the number of seeds produced by each adult plant ever. So $\gamma[.]$ is the mean number of offspring per plant. Call the whole thing K :

$$K = \gamma[\sigma + (1 - \sigma)\tau]$$

It is not surprising that it determines the boundary for stability:

$$\begin{aligned} \text{mean offspring} = K &< 1, & x = 0 & \text{ is stable} \\ \text{mean offspring} = K &> 1, & x = 0 & \text{ is unstable} \end{aligned}$$

It is often the case that conditions on the boundary for stability has an intuitive explanation in terms of the original biological model. It is worth looking out for these as they are a good check that the algebra has come out correctly and the answer is sensible.

end of
lecture 6

¹⁹If it is greater than one, then x_n grows eventually regardless of p_- . If it is less than one, then p_- also has modulus less than one, so x_n decays towards zero eventually.

1.3 Multi-species models

The dynamics of interacting populations (or biological substances) gives rise to the most interesting models in mathematical biology. In the next few sections, we will work through a series of examples, illustrating more general principles and techniques.

1.3.0 Revision: 2-D stability in continuous time

Consider the system:

$$\begin{aligned}\frac{du}{dt} &= f(u, v) \\ \frac{dv}{dt} &= g(u, v)\end{aligned}$$

A fixed point (u^*, v^*) satisfies $f(u^*, v^*) = 0$ and $g(u^*, v^*) = 0$.

To explore stability, consider a small perturbation to that fixed point. Set $u(t) = u^* + \xi(t)$, $v(t) = v^* + \eta(t)$ and expand in small ξ, η :

$$\frac{d\xi}{dt} = f(u^* + \xi, v^* + \eta) = \underbrace{f(u^*, v^*)}_{=0 \text{ at FP}} + \xi \left. \frac{\partial f}{\partial u} \right|_{\text{FP}} + \eta \left. \frac{\partial f}{\partial v} \right|_{\text{FP}} + \mathcal{O}(\xi^2, \xi\eta, \eta^2)$$

Similarly:

$$\frac{d\eta}{dt} = g(u^* + \xi, v^* + \eta) = \xi \left. \frac{\partial g}{\partial u} \right|_{\text{FP}} + \eta \left. \frac{\partial g}{\partial v} \right|_{\text{FP}} + \mathcal{O}(\xi^2, \xi\eta, \eta^2)$$

So the local dynamics comes down to the Jacobian:

$$\begin{pmatrix} \dot{\xi} \\ \dot{\eta} \end{pmatrix} = \begin{pmatrix} \frac{\partial f}{\partial u} & \frac{\partial f}{\partial v} \\ \frac{\partial g}{\partial u} & \frac{\partial g}{\partial v} \end{pmatrix} \bigg|_{\text{FP}} \begin{pmatrix} \xi \\ \eta \end{pmatrix}$$

where the 2×2 matrix is the Jacobian. In practice, just find the Jacobian and start from there. We might sometimes want the eigenvectors to help draw phase-diagrams, but usually we just want the eigenvalues.

Or even more basic, we just need to know the sign of the real parts of the eigenvalues. Let T be the trace and D the determinant of the Jacobian evaluated at some fixed point. Then (for 2-D), the eigenvalues are:

$$\lambda = -\frac{1}{2}T \pm \frac{1}{2}\sqrt{T^2 - 4D}$$

or alternatively the other way $T = \lambda_1 + \lambda_2$ and $D = \lambda_1\lambda_2$. So

$$\begin{aligned}D < 0 & : \text{ saddle} \\ D > 0, T < 0 & : \text{ stable} \\ D > 0, T = 0 & : \text{ centre} \\ D > 0, T > 0 & : \text{ unstable}\end{aligned}$$

We could then subdivide the stable and unstable cases according to focus or node (eigenvalues complex or real) by checking $T^2 - 4D$, but often in math bio we do not need to do this.

1.3.1 Competition models

The classic example is

$$\begin{aligned}\dot{N}_1 &= r_1 N_1 \left(1 - \frac{N_1}{K_1} - b_{12} \frac{N_2}{K_2} \right) \\ \dot{N}_2 &= r_2 N_2 \left(1 - \frac{N_2}{K_2} - b_{21} \frac{N_1}{K_1} \right).\end{aligned}$$

where there are two species N_1 and N_2 . Each species alone has simple logistic dynamics (see lecture 1) with linear growth, and a negative quadratic term which means each species has some stable equilibrium size (K_1 and K_2 respectively: the carrying capacities for each species alone). The terms with the b_{12} and b_{21} are the interactions between the species: they each slightly ‘harm’ the other. An analysis of this system is one of the questions on examples sheet 1. The net outcome is not too surprising: if the negative interaction terms are not too big, then the two species will coexist at some stable equilibrium value. If the interaction terms are too strong, then the two species cannot stably coexist, and one or other species wins out.

Here we study instead a different competition system, motivated by recent research on controlling the spread of dengue. Dengue is a virus that causes disease in humans. Rather than being transmitted directly from human to human, it requires an intermediate vector: a mosquito. If the mosquito bites and takes blood from someone who is infected, the mosquito could go on to infect anyone they bite later.

You can read a lot more about the ideas on the Eliminate Dengue website²⁰, but in brief: *Wolbachia* are a type of bacteria that can infect a huge range of insect species. Researchers have developed a strain of *Wolbachia* that can infect the kind of mosquitos that can carry dengue. The bacteria seems to block transmission of dengue virus, so we would like to see if this would be a viable way to control dengue in practice. Recent research suggests that the same approach will work for Zika virus, but there is more work to be done on this. For whichever virus, our question is the same: if we introduce some mosquitos that carry *Wolbachia* into the wild, will eventually all mosquitos carry *Wolbachia*?

In mosquitos, *Wolbachia* is only transmitted vertically, which means to offspring mosquitos (as opposed to ‘horizontal transmission’: to general others in same species). If a female mosquito is infected, her eggs will certainly be infected, regardless of the carrier status of the male. She will also produce fewer eggs than usual. Here’s the weird bit:

²⁰<http://www.eliminatedengue.com/>

if the female is uninfected but the male is infected, the eggs will not be viable at all, so no offspring at all. Now we start pulling this into a mathematical formulation.

Let x be the number of uninfected female mosquitos, and y be the number of infected female mosquitos. Assume that the uninfected mosquitos have a per capita death rate d and a bonus per capita death rate of ϵ times the total number of female mosquitos (competition). The infected mosquitos have shorter lifespans on average, which we model here as a higher death rates by having an additional factor μ with $d(\mu > 1)$.

We do not need to explicitly track the males: just assume their infection state is in proportion to the females (which you can check is the case by building full equations if you really want), so a proportion $x/(x+y)$ uninfected. Also assume there are enough males around for all eggs to be fertilised. Suppose that in a purely uninfected population, the rate of viable eggs for female mosquitos being produced is r per female capita. If the female is infected, assume that they produce λ times the normal number of eggs, so $\lambda < 1$.

Summarising the four possible crosses (female-male infection status combinations):

Cross	Frequency	Egg rate	State
$F \times M$	$x \cdot \frac{x}{x+y}$	r	Uninfected
$F \times \textcolor{green}{M}$	$x \cdot \frac{y}{x+y}$	0	\nexists
$\textcolor{green}{F} \times M$	$y \cdot \frac{x}{x+y}$	λr	Infected
$\textcolor{green}{F} \times \textcolor{green}{M}$	$y \cdot \frac{y}{x+y}$	λr	Infected

Where the $\textcolor{green}{F}$ means infected and F is uninfected females, similarly for $\textcolor{green}{M}$ and M for males.

Bringing this together, we get the following system:

$$\begin{aligned}\dot{x} &= r x \frac{x}{x+y} - d x - \epsilon x(x+y) \\ \dot{y} &= \lambda r y \frac{x}{x+y} + \lambda r y \frac{y}{x+y} - \mu d y - \epsilon y(x+y)\end{aligned}$$

By rescaling time by a factor r and both x and y by a factor ϵ/r , the system can be slightly tidied²¹:

$$\begin{aligned}\dot{x} &= x \left[\frac{x}{x+y} - \frac{d}{r} - (x+y) \right] \\ \dot{y} &= y \left[\lambda - \mu \frac{d}{r} - (x+y) \right]\end{aligned}$$

²¹There's a wording ambiguity about which direction is meant by 'rescale by a factor', but interpret here in the direction that tidies things up!

Looking at each population alone (i.e. forcibly setting $x = 0$ or $y = 0$), each is just logistic, and we can pick out the equilibrium single population sizes:

$$\text{Uninfected only, } y = 0 : \quad \dot{x} = x \left(1 - \frac{d}{r} - x \right) = x (x_0 - x)$$

$$\text{Infected only, } x = 0 : \quad \dot{y} = y \left(\lambda - \mu \frac{d}{r} - y \right) = y (y_0 - y)$$

where $x_0 = 1 - d/r$ and $y_0 = \lambda - \mu d/r$ are the equilibrium sizes of each population alone.

Now take some sensible assumptions, thinking about the case that we wish to model. We are imagining an uninfected resident population, so they are viable, i.e. $r > d$ so $x_0 > 0$. We also want to consider the case when the purely infected population is viable, otherwise the proposed introduction is hopeless anyway, so assume $\lambda r > \mu d$, i.e. $y_0 > 0$. We have already assumed that infection reduces the number of eggs ($\lambda < 1$) and infection increases the death rate ($\mu > 1$). So putting all these together, we have $0 < y_0 < x_0 < 1$.

This has demonstrated another strategy for rescaling parameters: discover a meaningful parameter combination by considering the relatively simple equilibrium points such as each population alone, then see if the system can be written out nicely in terms of these quantities. It works well for this example. Organising the parameters to be in terms of x_0 and y_0 , the full system becomes:

$$\begin{aligned} \dot{x} &= x \left[x_0 - \frac{y}{x+y} - (x+y) \right] \\ \dot{y} &= y \left[y_0 - \frac{x}{x+y} - (x+y) \right] \end{aligned}$$

There are four non-negative fixed points: no mosquitos $(0, 0)$, purely uninfected $(x_0, 0)$ or purely infected $(0, y_0)$, and an interested mixed state one (x_1, y_1) , where setting the square brackets to zero and working through:


$$x_1 = y_0(1 - x_0 + y_0) \quad \text{and} \quad y_1 = y_0(x_0 - y_0)$$

and as $0 < y_0 < x_0 < 1$, both x_1 and y_1 are positive.


Exercise 15: Check the Jacobian at (x_1, y_1) and show that it corresponds to a saddle

The next stage is to find the null clines and then put together the phase diagram. Nullclines are just when $\dot{x} = 0$ and when $\dot{y} = 0$. On these, the trajectories are pure

vertical and pure horizontal respectively. But more usefully, these curves also divide up the phase plane into regions where the direction of trajectories is purely in one quadrant (e.g. up and left), which helps pulling the picture together.



Finally, we may answer our original question: what happens if we start at the uninfected equilibrium and introduce some infected mosquitoes? This is sketched below: it really depends where the dividing line is between basins of attraction of the two stable fixed points. The dividing line (separatrix) goes through the saddle point (indeed it is the stable manifold of the saddle point).



A schematic of the different outcomes of an introduction of infected mosquitoes. The shaded area in each case is the region where trajectories would head towards $(x_0, 0)$, drawn for two different examples. The red dot is where we are on the phase diagram after an introduction of a certain number of infected mosquitoes. Clearly, if we want all mosquitoes infected eventually, we would like to be in the righthand regime.

So introducing a very small number of mosquitoes is not enough, there has to be quite a few infected mosquitoes brought in. However, on the plus side, once the infected

population is established, it won't easily revert back.

Exercise 16: Imagine you are a mathematical modeller advising on this project. The experimentalists can work to change the strain of Wolbachia so as to make it less damaging to the mosquitos by softening the effect on the death rate (decrease μ towards 0) or the egg production rate (increase λ towards 1). We would like to make it so that a small introduction of infected mosquitos would be enough to make all mosquitos infected eventually. Would you advise them to concentrate their efforts on λ or μ or a combination of the two? (Hint: consider the impact on x_0 and y_0 and then the nullclines).

We come back to this example later in the course, when we consider the spatial effects.

end of
lecture 7

1.3.2 Predator-prey models

No course in math bio would be complete without this iconic system: the Lotka-Volterra model of predator-prey dynamics. The prey population (size N) would grow on its own, and the predator population (size P) would decay on its own. The interaction is predation which happens at rate proportion to the product of both population sizes: mass-action. The predation harms the prey and benefits the predators. This leads to the system

$$\begin{aligned}\frac{dN}{dt} &= aN - bNP = N(a - bP) \\ \frac{dP}{dt} &= cNP - dP = P(cN - d)\end{aligned}$$

where a , b , c and d are all positive. This can be rescaled to

$$\begin{aligned}\dot{u} &= u(1 - v) \\ \dot{v} &= -\alpha v(1 - u)\end{aligned}$$

Exercise 17: Carry out this rescaling and show that $\alpha = d/a$ (and also note $\alpha > 0$).

There are two fixed points: $(0, 0)$ and $(1, 1)$. The general Jacobian is given by

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

So evaluating at the fixed point at the origin:

$$J_{(0,0)} = \begin{pmatrix} 1 & 0 \\ 0 & -\alpha \end{pmatrix}$$

And checking the eigenvalues²² we see that the origin is a saddle. This is not a surprise as we geared the model so the prey would group on their own and the predators would decay.

The non-trivial fixed point gives

$$J_{(1,1)} = \begin{pmatrix} 0 & -1 \\ \alpha & 0 \end{pmatrix} \quad T = 0, D = \alpha \quad \lambda = \pm i\sqrt{\alpha}$$


which corresponds to a centre. This is neither stable nor unstable to linear order, and in theory requires further work to determine non-linear stability (higher-order terms).

Nullclines In sketching phase diagrams, fixed point analysis gives us the dynamics close to equilibrium values, and we are left to join up the picture in between. Very often, it is useful to divide up space into regions where \dot{u} and \dot{v} are positive or negative. This means finding the *nullclines*: curves where one or other variable is unchanging.

For the Lotka-Volterra system:

$$\dot{u} = 0 : \quad u = 0 \quad \text{or} \quad v = 1$$

$$\dot{v} = 0 : \quad v = 0 \quad \text{or} \quad u = 1$$



This is starting to suggest that trajectories might be cycles: closed curves so the solution is periodic. However neither the Jacobian nor the diagram with the nullclines has conclusively shown that we have cycles. Luckily, we can explicitly find the trajectories for this system. Start by removing the time dependence to just think about curves in u, v space:

$$\frac{du}{dv} = \frac{\dot{u}}{\dot{v}} = \frac{-u(1-v)}{\alpha v(1-u)}$$

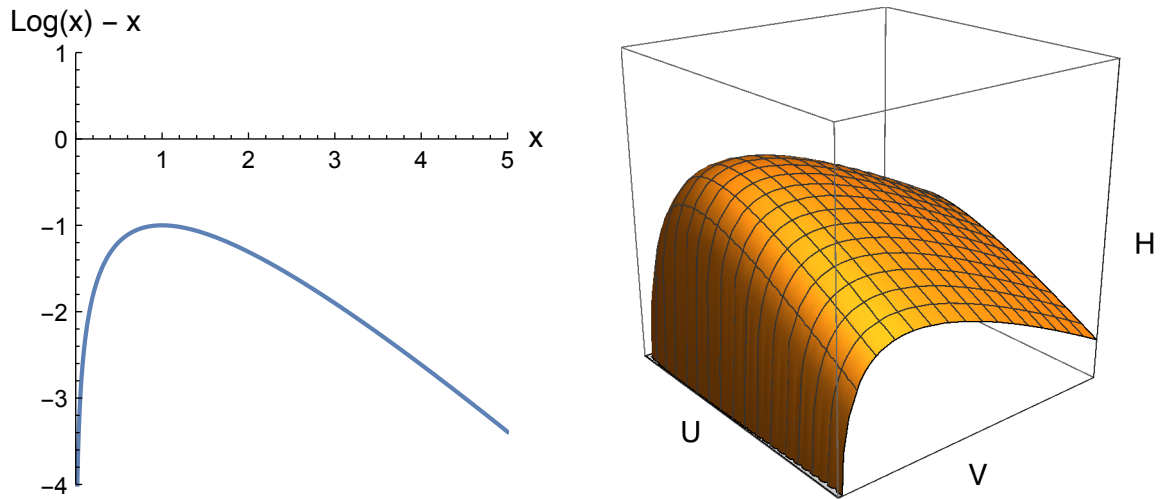
²²The matrix is diagonal, so just read off the eigenvalues

Then integrate:

$$\int \frac{\alpha(1-u)}{u} du = - \int \frac{1-v}{v} dv$$


$$\alpha(\log u - u) = -(\log v - v) + C$$

In other words, we have $H(u, v) = \alpha(\log u - u) + (\log v - v)$ is a constant on trajectories²³.



Contours of constant $h(u, v)$, start by considering $\log x - x$. Then put into 3D by noting it is the sum of these in u and v .

Then plot contours of constant H :



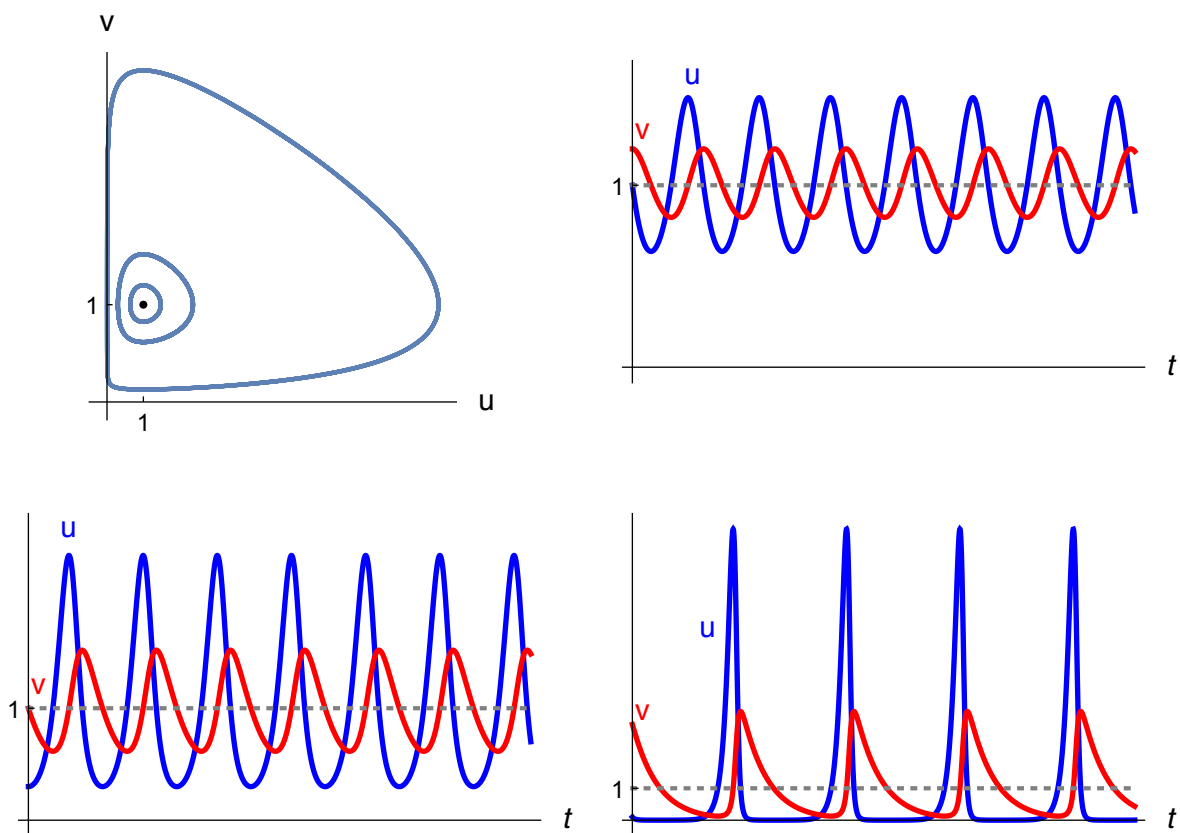
And we know trajectories must remain on constant contours, so just add some arrows and we are done.

²³In language of dynamics systems this is like a Hamiltonian.

For the Lotka-Volterra system, we can also show that the *time-average* of the period orbits are all equal to the values of the population size at the fixed point. These were both normalised to 1 here, but it still holds for the system before rescaling of course. Starting from $\dot{u} = u(1 - v)$ and integrating over one period:

$$\begin{aligned} \int \frac{1}{u} du &= \int 1 - v dt \\ \underbrace{[\log u]}_{=0 \text{ as same } u \text{ at each end}} &= \underbrace{\int 1 dt}_{\text{period } T} - \underbrace{\int v dt}_{T \times \text{average } v} \end{aligned}$$

So the time-average value of v is 1. A similar calculation starting from \dot{v} shows that the time-average of u is also 1. This is true no matter which contour of H we are on.



Solutions of the Lotka-Volterra system for assorted starting conditions, with the trajectories shown also together in the u, v plane. Note that the vertical scale changes between plots, but the dotted line marks 1, the time-average of both populations. It can also be seen from these solutions that the prey (blue) lead and the predators (red) follow.


The effect of fishing Returning to the unscaled original system, we can ask the question: what is the effect of fishing on the system? We simulate this by introducing terms representing fishing of both populations at a constant per capita rate (but potentially different constants for predators and prey):

$$\begin{aligned}\dot{N} &= N(a - bP) - f_N N = N((a - f_N) - bP) \\ \dot{P} &= P(cN - d) - f_P P = P(cN - (d + f_P))\end{aligned}$$

And this is just the same system as before with just a tweak of parameters: a is replaced by $a - f_N$ and d is replaced by $d + f_P$ (taking care with the signs). So now no further serious work is needed: just look at the effect of shifting the parameters. The fixed point moves:

$$\text{Originally at } \left(\frac{d}{c}, \frac{a}{b} \right) \quad \text{Now: } \left(\frac{d + f_P}{c}, \frac{a - f_N}{b} \right)$$

i.e. a higher number of prey and lower number of predators. We also know from above that the time-average of the cycles will similarly move.



This result is remarkably simple, but maybe slightly surprising: fishing is depleting both populations, yet the net result is actually an increase in the prey population. The effects of fishing are concentrated in the predators. In general, when we have multiple populations in balance with each other and their environment (an ecosystem) external interventions can have non-obvious effects. Mathematical models can help us understand these²⁴. Vito Volterra used exactly the system we have studied here to understand the changes in fish population in the Adriatic. During war time, fishing was lower than usual but the proportion of the catch that were predators (sharks etc.) went up. This is rather elegantly explained by the above analysis.

The cycles are not robust! That we have cycles is pretty special. If we change the structure of the Lotka-Volterra system in almost any way²⁵ then the cycles will break. Here we look at the example of adding logistic-style quadratic terms, which represent

²⁴This is where mathematical biology has been historically most successful: at explaining observed phenomena.

²⁵The parameter tweak of fishing was not a structural change.

interaction within each species alone: intra-specific competition adds to the death rate. Working now in the rescaled system:

$$\begin{aligned}\dot{u} &= u(1-v) - \epsilon_u u^2 \\ \dot{v} &= -\alpha v(1-u) - \alpha \epsilon_v v^2\end{aligned}$$

Assume both rates $\epsilon_u, \epsilon_v \ll 1$ so the change to the system is small²⁶. Note the factor of α in new term in the second equation: this is just to make the algebra tidier, as the system can now be written:

$$\begin{aligned}\dot{u} &= u[1-v-\epsilon_u u] \\ \dot{v} &= -\alpha v[1-u+\epsilon_v v]\end{aligned}$$

Look for fixed points, paying attention only to $u \geq 0$ and $v \geq 0$.

Exercise 18: Show that there are three fixed points now.

- The usual trivial fixed point $(0, 0)$
- Prey population only (at their logistic equilibrium): $(\epsilon_u^{-1}, 0)$
- One with both populations present: (u^*, v^*) where

$$u_* = \frac{1 + \epsilon_v}{1 + \epsilon_u \epsilon_v} \quad \text{and} \quad v_* = \frac{1 - \epsilon_u}{1 + \epsilon_u \epsilon_v}$$

and hence $u_* > 1$, $v_* < 1$.

Use the Jacobian²⁷ to analyse the stability of this fixed point.

$$\begin{aligned}J_{(u_*, v_*)} &= \begin{pmatrix} \underbrace{[1 - v_* - \epsilon_u u_*]}_{=0} + u_*(-\epsilon_u) & -u_* \\ \alpha v_* & -\alpha \underbrace{[1 - u_* + \epsilon_v v_*]}_{=0} - \alpha v_* \epsilon_v \end{pmatrix} \\ &= \begin{pmatrix} -\epsilon_u u_* & -u_* \\ \alpha v_* & -\alpha v_* \epsilon_v \end{pmatrix} = \begin{pmatrix} - & - \\ + & - \end{pmatrix}\end{aligned}$$

And hence $T < 0$ and $D > 0$ and this is always a stable fixed point.

²⁶Or at least it is small for sensible u, v , i.e. on scale of original fixed point

²⁷Three useful tricks: (i) keep brackets intact as long as possible: for example in finding the fixed point you will have solved $(1 - v - \epsilon_u u) = 0$ so keep it together when differentiating and then it will be effortlessly zero at the fixed point; (ii) don't sub in the expressions for u_* and v_* too soon, just keep them as u_* and v_* as long as possible; (iii) often we don't even need the full Jacobian evaluated: it is enough to see the signs of the entries, sometimes.

Exercise 19: Continue this example by using nullclines.

Hint: think about the competition terms and how their main effect is to ‘push’ large u and v inwards.


Exercise 20: For those who did dynamical systems^a: construct a (strict) Lyapunov function to show (u^, v^*) attracts all trajectories that start in $u, v > 0$.*

Hint: $H(u, v)$ was constant before, so seems like a good place to start.

^afor those that didn’t but are unlikely to be put off by my warnings: good for you! All you need here is that a Lyapunov function $V(u, v)$ is such that it is strictly decreasing along trajectories in this domain, except at the fixed point (where it is constant). The function V should be continuous and zero at the fixed point, positive everywhere else. If you find such a function, trajectories have V decreasing and bounded below, so V tends to some constant, so trajectories tend to a some set where $\dot{V} = 0$, which here is just the fixed point.


end of
lecture 8

1.3.3 Chemical kinetic models



The big yellow blob is the bacterium, the blue things are the receptors and the pink dots are the nutrients. OK so a bit of imagination required!

Imagine a single bacterium. It has receptors ready to pick up nutrients. The receptors start empty (E) and pick up nutrients in a substrate (S). Receptors then carry the nutrient so there is a complex (C) consisting of a reception and a nutrient particle. This goes into the bacterium where the nutrient is dropped off to be final product (P). We can describe this in terms of chemical kinetic notation as follows:



with three rates: k_1 is the rate that E and S combine, k_2 is the rate at which C can fail and fall apart back into E and S, and k_3 is the rate at which C is internalised and drops off P and becomes an empty E again.

To set up differential equations, suppose that everything is large numbers so that we can sensibly talk about concentrations. Let s be the concentration of S, and similarly for e , p and c . Then in differential equations, our system can be described as

$$\begin{aligned}\dot{s} &= -k_1 e s + k_2 c \\ \dot{e} &= -k_1 e s + k_2 c + k_3 c \\ \dot{c} &= +k_1 e s - k_2 c - k_3 c \\ \dot{p} &= +k_3 c\end{aligned}$$

For initial conditions, suppose that we start with all the nutrients outside and all receptors empty, so $s(0) = s_0$, $e(0) = e_0$, $c(0) = 0$ and $p(0) = 0$, for some $s_0, e_0 > 0$. Also suppose that there are a lot more nutrients than receptors, so that receptors are used many times before the substrate is fully depleted: $s_0 \gg e_0$.

From thinking about the original scenario, we can immediately see that there must be some invariant quantities. In particular, the number of receptors stays constant if we include both carrying (C) and empty (E). Also the nutrients are constant if we consider them outside (S), being carried (C) and inside (P).

We could also see this from the equations by spotting combinations that add to give no change:

$$\dot{e} + \dot{c} = 0, \quad \dot{s} + \dot{c} + \dot{p} = 0$$

so these are constant, and we can determine the constants by initial conditions:

$$e + c = e_0, \quad s + c + p = s_0.$$

We can use these to go from four differential equations down to two, substituting for e and p (in principle at least, p doesn't actually occur anywhere). This is our reduced system:

$$\begin{aligned}\dot{s} &= -k_1(e_0 - c)s + k_2 c &= -k_1 e_0 s + (k_1 s + k_2)c \\ \dot{c} &= +k_1(e_0 - c)s - k_2 c - k_3 c &= +k_1 e_0 s - (k_1 s + k_2 + k_3)c\end{aligned}$$

This is already much better, but we can rescale further by strategic use of our initial conditions. We set $u = s/s_0$ and $v = c/e_0$ and then these have nice interpretations: u is the proportion of total nutrient still available (so u is one initially) and v is the proportion of receptors that are currently carrying nutrient (so v is zero initially). We could also rescale time to adsorb one of the rates.

Exercise 21: By a suitable rescaling, turn the system into:

$$\begin{aligned} u' &= -u + (u + \mu - \lambda)v & u(0) &= 1 \\ \epsilon v' &= +u - (u + \mu)v & v(0) &= 0 \end{aligned}$$

where

$$\lambda = \frac{k_3}{k_1 s_0}, \quad \mu = \frac{k_2 + k_3}{k_1 s_0}, \quad \epsilon = \frac{e_0}{s_0} \ll 1.$$


We see that the v dynamics are faster than the u dynamics (by an order of $1/\epsilon$). So we can see what the solutions will be like without doing too much more work. Do the v dynamics first imagining u to be almost fixed. We can see that it is exponential decay down to a constant:

$$v = \frac{u}{u + \mu}$$


where the decay is fast (takes order ϵ time). So assume that has happened, and as we now slowly change u , the fast dynamics of v ensures it keeps moving quickly back to its equilibrium value again (which is itself a function of u). Substitute this value of v into the u' equation:

$$u' = -u + (u + \mu - \lambda) \frac{u}{u + \mu} = -\frac{\lambda u}{u + \mu}$$

so we can see that u just decays towards $u = 0$.



This is for $\epsilon = 0.1$ which is not super small, and even then it is clear the trajectories quickly move to $v = u/(u + \mu)$. The red trajectory is for the initial condition $u = 1, v = 0$.



Same but for $\epsilon = 0.01$, when the fast-slow dynamic is clearer: trajectories move vertically first until they are near the $v = u/(u + \mu)$ curve, then they move slowly along that curve

Thinking back to the original system, the v dynamics being fast just means that the proportion of receptors occupied quickly settles to some value and then slowly readjusts as the nutrient is depleted. This was all a consequence of our original assumption of there being much more nutrient out there than the total number of receptors.

This system is actually a classic model for enzyme kinetics used by biochemists (Michaelis-Menten).


Exercise 22: Work back to the original system and show that the inverse of the rate of nutrient uptake is linear in the inverse of substrate. In other words:

$$\left(\frac{dp}{dt}\right)^{-1} = A + B s^{-1}$$

for some constants A and B .

In principle this gives a test for whether a set of observations are consistent with Michaelis-Menten kinetics. In practice, one should probably be more than mildly concerned about issues with inverting small numbers, when those small numbers are hard-to-measure experimental quantities.

Exercise 23: Translate this to differential equations:



Also, find an invariant sum. You can do this by considering the reactions above directly, or by inspecting your differential equations (I think it is usually easier to spot these from the reactions directly.)

Note that $2C$ is shorthand for $C + C$, so reaction rates will involve c^2 . Similarly $3C$ and so on.

1.3.4 Epidemic models

There is a long and fruitful tradition of using mathematical models to capture how infectious disease spreads through a population. The most classic model (first published over a century ago) is so-called ‘SIR’ model. Here, the population is divided into three groups: susceptible (S), infected/infectious (I) and recovered/removed (R). This compartmental model is a little different to the population models we have seen already. In predator-prey and competition models that we have considered so far, the groups are really separate and it is not possible to move from one to another: you’re there from birth to death. Here, it is possible to move between S , I and R . This type of model is usually called a ‘compartmental model’ where there are flows and interaction between compartments. It works for populations as well as chemical kinetics (see below) among many other things.

Here we have two processes causing a flow between compartments: infection takes individuals from S to I and recovery takes them from I to R . For now, we do not include any births or deaths: we are thinking of some short outbreak of disease in a closed population. The full system is:

$$\begin{aligned} \dot{S} &= -\beta IS \\ \dot{I} &= +\beta IS - \nu I \\ \dot{R} &= +\nu I \end{aligned}$$


The total rate of infection is mass-action: proportional to the product of S and I . We have seen interaction terms like this before, for example in predation ($N \times P$). It corresponds to the frequency of the two groups encountering each other. Or another way to think about it is that the rate *per capita* of infection is proportional to the number of infected I , the number available to infect is S , so the total rate is proportional to $S \times I$. We call coefficient β the transmission rate. The recovery term is simpler, just simple rate per individual: ν .

It is easy to see that the total population size remains constant here²⁸. We then denote total population size as $N = S + I + R$, and see $\dot{N} = 0$. So if we know N (think of as an extra parameter), then knowing any two of S , I and R is enough to determine the third. So, we might as well drop R from the system, and actually we don't even need to sub $R = N - I - S$ in anywhere as it does not occur in the other equations:

$$\begin{aligned}\dot{S} &= -\beta IS \\ \dot{I} &= +\beta IS - \nu I\end{aligned}$$

We can reconstruct R if we even need it, which we rarely do in practice. Usually I is most visible, perhaps followed by S .

We can solve for S and I numerically without much difficulty, and here is a typical output when we start with a few infecteds and the rest susceptible:



When can an epidemic happen?

An 'epidemic' is simply when the cases increase, i.e. $\dot{I} > 0$. So initially $t = 0$ and

$$\dot{I}(0) = [\beta S(0) - \nu]I(0)$$

and assume that one or a few infected are introduced initially so that $I(0) > 0$. The we just need to check the sign of $[\beta S(0) - \nu]$ (where $S(0)$ is just the initial number of susceptibles). Imagine that we are starting from nearly everyone susceptible, and only a very small number infected, then $S(0) \approx N$. We see that an epidemic is possible if and only if

$$\frac{\beta}{\nu}N > 1.$$

²⁸Disturbingly, this all still works for fatal diseases: the R compartment serves equally well for those who have recovered and are now fit and well and now immune to the disease, as those who are in fact dead. Either way, they can't get the disease again. We just call them all 'removed', which always seems a bit uncaring.

As always, we should try and interpret this inequality. Looking at the lefthand side, the trickiest bit is the $1/\nu$. We know ν is the recovery rate per individual. One way to imagine this is as a random process with event happening at rate ν if it hasn't happened already. This leads to an exponential distribution of time to recovery. The *expected* time to recovery is $1/\nu$, i.e. the duration of a infection. This is multiplied by β , which is the rate that one individual infects each susceptible²⁹. Finally, N is roughly the number of susceptible individuals around. So, time infected, rate of on infected infecting each contact, number of contacts: the product is total number infected, on average, from one infection.

Actually this idea holds more generally in disease modelling, and this number is known as the *reproductive ratio*, or R_0 ³⁰. Note this is an unfortunate piece of historical notation as this R_0 has nothing to do with the R of *SIR*. It is defined as *the mean number of secondary cases in an otherwise susceptible population*. In other words, if we drop one infected into our susceptible population, how many others will they directly infect? For this classic *SIR* model, we have found:

$$R_0 = \frac{\beta}{\nu} N.$$

and our threshold for an epidemic is $R_0 > 1$. This of course is also the threshold in more general models also, once we have identified R_0 (which is not always easy to do in practice³¹).

end of
lecture 9

Vaccination We actually have enough to say quite a bit about the potential for vaccination as a control measure, without going into the differential equations in any further detail. Suppose that we intervene before a possible epidemic to vaccinate a proportion p of the population, and assume the vaccine gives total protection against the disease. This is equivalent³² to moving a proportion p of the population straight into R and leaving $S(0) \approx (1 - p)N$. Our condition for an epidemic to be possible is now

$$\frac{\beta}{\nu} S(0) = \frac{\beta}{\nu} (1 - p)N = (1 - p)R_0 > 1.$$


We can rearrange this to find the threshold for p in terms of R_0 : $p < 1 - 1/R_0$ for an epidemic to be possible:

²⁹a sort of *per capita per capita*

³⁰pronounced R-nought

³¹There's plenty more to read and learn about R_0 , and this paper by Heffernan *et al.* is a good place to start: <http://doi.org/10.1098/rsif.2005.0042>

³²All of this is for a perfect vaccine, so being vaccinated means you truly in R . One can easily extend this to think about imperfect vaccines.



If R_0 is less than one, there's no chance of an epidemic anyway, so no need to vaccinate. For $R_0 > 1$, there is some $p_c \in (0, 1)$ which is the vaccination threshold. It is not too hard to understand this intuitively, and to see why it depends on R_0 only. We'd like to vaccinate enough individuals so that the disease cannot spread through the population. We'd like to make the *effective* reproduction ratio less than 1. If we introduce one infected and they would infect R_0 others, we'd like to vaccinate at least $R_0 - 1$ of those R_0 . That's just the proportion p_c .

This is all interlinked with the idea of *herd immunity* which is the idea that having many individuals immune can even protect other individuals. For each person vaccinated, it protects not just that person, but all the people they would have gone on to infect, and even later generations of infection onwards. These ideas are not new to mathematicians, and can be thought about more generally in terms of percolation theory.

Epidemic and final size

We can go further with this simple model and learn something of the trajectory of an epidemic once it does start. Here, it does not make sense to do a traditional fixed point analysis: the whole the epidemic is itself transient behaviour.

At the end of an epidemic, the susceptibles are *not* fully depleted: S seems to plateau to some positive value. Actually, we can understand this by considering the $S - I$ phase-plane. The whole of $I = 0$ is a line of fixed points, but it is not so much these we are interested in as the trajectories that arrive at them. We can explicitly solve for I as a function of S in this case:

$$\frac{dI}{dS} = \frac{\dot{I}}{\dot{S}} = \frac{+\beta IS - \nu I}{-\beta IS} = -1 \frac{\nu}{\beta S} = -1 + \frac{N}{R_0} \frac{1}{S}$$


And integrating with respect to S :

$$I = -S + \frac{N}{R_0} \log S + C$$

for some constant C which is determined by initial conditions (now thinking about I and S as functions of time again):

$$[I(t) - I(0)] = -[S(t) - S(0)] + \frac{N}{R_0} \log \left[\frac{S(t)}{S(0)} \right] \quad (15)$$

The phase plane for this system should be restricted to the triangle in the $S - I$ plane where $S \geq 0$, $I \geq 0$, and $S + I \leq N$. We can see from above that $dI/dS > -1$ so the trajectories never go as steep downwards as the diagonal line. If we now consider things as a function of time, we see that trajectories can start anywhere, and head to the $I = 0$ line.



Solutions $I(S)$. To make this a phase diagram for the system in time, put some arrows on the curves (all will have S decreasing). The trajectory in red is the one which has $S \approx N$ and $I \approx 0$ initially, corresponding to an epidemic starting after a few infecteds introduced into an otherwise susceptible population.

Exercise 24: Show that trajectories have a maximum number of infecteds as $S = N/R_0$, and explain why this is. Hint: it might be useful to use the idea of ‘effective’ reproduction ratio (often labelled R_{eff} or r) which is the equivalent of R_0 but at some general value of S).

Suppose we know $I(0)$ and $S(0)$, can we work out where exactly on the $I = 0$ line that the trajectory will head to? In other words, can we work out what value S will

tend to? It turns out we can, at least implicitly. Consider in particular the case when we introduce a small number of infecteds into an otherwise susceptible population ($I(0) \ll N$, $S(0) \approx N$). Let time go to infinity, and we know $I(t) \rightarrow 0$. Suppose $S(t)$ tends to some value which we write as σN (which is the definition of σ). We see that σ is the proportion of the population that escape the epidemic and actually never get infected. We can find an implicit equation for σ by substituting all our initial and final values into 15:


$$0 = -[\sigma N - NB] + \frac{N}{R_0} \log \left[\frac{\sigma N}{N} \right]$$

and rearranging:

$$\sigma - \frac{1}{R_0} \log \sigma = 1. \quad (16)$$

Exercise 25: Start from the equations for \dot{S} and \dot{I} and get to equation 16 (without any notes).

Note that N has cancelled out: our escape proportion σ only depends on R_0 . This equation always has $\sigma = 1$ as a solution. This corresponds to nothing happens: final values and initial values the same, i.e. no epidemic. This is not the interesting solution. For $R_0 > 1$, there is a solution for $\sigma \in (0, 1)$, and that is the one that we want³³.



This was an important early result of disease modelling (about a century ago): epidemics burn out leaving some proportion still susceptible. Before the concept of herd

³³Actually for R_0 securely above 1, $\sigma \approx e^{-R_0}$, so proportion escaping gets very small, but non-zero still.

immunity, this would have been a non-trivial result. Without this understanding, epidemiologists would need to invoke some special reason for the epidemic ending before it got everyone, for example some people are naturally protected, or the virus has changed, or the weather changed and stopped the epidemic. What this very basic model showed was that none of this was necessary: epidemics will just burn themselves out before they get to everyone.

Exercise 26: Show that if $R_0 = 1 + \epsilon$, then $\sigma = 1 - 2\epsilon + \mathcal{O}(\epsilon^2)$.

Extensions to the SIR model

There are a lot of ways to extend the SIR model ³⁴. Indeed it is a whole research area. It is tempting just to add more and more detail in the hopes of making something more ‘realistic’, but in doing that, we lose model parsimony. Simple models are good as they are tractable enough to draw some general insights (such as R_0 above, which is a useful concept far more generally). But sometimes there are some parts of the dynamics that we really need to add to understand some particular problem.

One more example here: we have looked at the simple SIR model which is a one-off epidemic. What about something like measles in the UK before vaccination? It seemed to cause large epidemics every two years or so. To consider this longer term dynamic, we absolutely must add in some host turnover, i.e. natural births and deaths (not due to the disease), to keep a trickle of new susceptibles coming into the system. Easiest way to do that is a constant birth/death rate per capita (μ):

$$\begin{aligned}\dot{S} &= -\beta IS - \mu S + \mu N \\ \dot{I} &= +\beta IS - \nu I - \mu I\end{aligned}$$

And if we did put the \dot{R} equation, it would have $-\mu R$ representing deaths. The births are all going into S : we assume everyone is born susceptible.

We can see by checking the \dot{I} equation as above that we have a slightly modified expression for R_0 in this model:

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

³⁴And maybe one way to simplify: if there is no recovery, then we have the SI model. This might be appropriate for an extreme zombie invasion. As before, the total population size is constant so $S = N - Z$. The dynamics is now just 1D, and is given by

$$\dot{Z} = \beta Z(N - Z)$$

and hence zombies are logistic. Surprisingly, there is a whole book on the ‘Mathematical Modelling of Zombies’ by ‘Robert Smith?’ (yes, that is really his name).

(which is reassuring the same as before if $\mu = 0$). Now we are going to focus on acute (short) infection, i.e. the infectious period is much shorter than host lifetime. Or put another way, the recovery rate is far higher than the natural death rate: $\nu \gg \mu$.

The births and deaths mean that there is a balance to the epidemic dynamics: susceptibles are depleted by the epidemic, but replenished by natural population turnover. So, now there is actually a fixed point with disease present ($I^* > 0$):

$$S^* = \frac{\mu + \nu}{\beta} = \frac{N}{R_0}, \quad I^* = \frac{\mu(N - S^*)}{\beta S^*} = \frac{\mu}{\beta}(R_0 - 1)$$

and note attempt to write things in terms of R_0 whenever we see likely-looking expressions appear. We really should not be surprised by the factor of $R_0 - 1$. It means this fixed point only makes sense when $R_0 > 1$, which seems sensible.

The Jacobian is given by

$$J = \begin{pmatrix} -\beta I - \mu & -\beta S \\ \beta I & [\beta S - \nu - \mu] \end{pmatrix}$$

so evaluating at fixed point (noting that once again the square bracket is zero):

$$J_{(S^*, I^*)} = \begin{pmatrix} -\mu R_0 & -(\mu + \nu) \\ \beta(R_0 - 1) & 0 \end{pmatrix}$$

and for once, we are actually going to find the eigenvalues, as there is more we can learn here than just stability or otherwise. The eigenvalues λ satisfy

$$\lambda^2 - \text{tr}(J)\lambda + \det(J) = \lambda^2 + \mu R_0 \lambda + \mu(\mu + \nu)(R_0 - 1)$$

which has solutions


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

This looks like a terrible mess for a moment, and then we realise that the square root contains μ^2 and $\mu\nu$ terms. We were considering the case when $\nu \gg \mu$, so this simplifies (taking factor of -1 out of square root to get i):

$$\lambda \approx -\frac{1}{2}\mu R_0 \pm i \underbrace{\sqrt{\mu\nu(R_0 - 1)}}_{=\omega}.$$

The eigenvalues are a complex pair with negative real part, so the fixed point is a stable focus (inwards spirals). This is not too hard to think about if, we suppose the epidemics

happen and then host turnover refills S gradually. From these considerations, we can draw a phase diagram, and typical dynamics against time:



I have cheated here to make the diagrams a bit clearer. On the left, I have scaled S down somewhat to fit it on same plot. On right, I've taken a different $\nu : \mu$ ratio, to move the fixed point away from the S axis, so swirly dynamics is clearer.

This is like a damped pendulum, settling down to some endemic equilibrium. Looking at these epidemics as we settle down, we can work out their approximate period from the imaginary part of the eigenvalues. The dynamics will be like $\cos(\omega t)$, and some decaying exponential. So the period T is

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

and it is of note that $(\mu\nu)^{-\frac{1}{2}}$ is just the geometric mean of host lifetime and disease infectious period. The $R_0 - 1$ is back again. If R_0 is only just above threshold, the period between epidemics is large. If R_0 is huge, then short period (bit like a pendulum in very strong gravity).

To finish the measles story, here we have period epidemics settling down. But real life is not like this model in one crucial way. The dynamics of measles in UK before vaccination involved much of the spread being in school-age children. And of course then, we should think about how schools worked. Kids aged 5 or so start school together in September, at the start of their first academic year. So, rather than there being a gentle trickle of susceptibles into the system, there is a huge kick once a year. This is enough to stop our damped pendulum settling down. The estimated T for measles (using the expression above, put $R_0 = 20$, infectious period about 12 days, lifetime about 70 years) is about 2.18 years. The forcing rounds this to a whole number of years, and so we end up with our epidemics every two years.

The author of these notes is probably biased (this is JRG's research area), but the field of disease modelling is very rich in terms of both mathematical interest and practical

importance. However, we have some other interesting topics to cover in this course, so we will move on. If you want to read more, then do have a look at the book called 'Modeling Infectious Diseases in Humans and Animals' by Keeling and Rohani.

end of
lecture 9

1.3.5 Excitable systems

In this section we consider *excitable systems* where a biological system can 'fire' following a small impulse. Heart muscle cells are one example of this: these are primed to contract in response to small electrical impulses. At the huge scale we could also consider large population phenomena such as plankton blooms, where relatively small environmental effects can lead to a dramatic burst of plankton growth. Arguably, some epidemic models should be considered also as excitable systems: a small number of initial infecteds can spark off a major epidemic. However for contrast here we will focus on a small physiological system: a neuron (nerve or brain cell) which receives a small input signal and can then fire an onwards signal to other cells.

The biochemical basis for signal propagation by neurons is well understood. It comes down to considering the difference in electrical potential (voltage) between the inside and outside of the cell, and how positively charged sodium and potassium ions are transported across the membrane by pumps. These pumps themselves change their behaviour according to the voltage difference. It is possible to develop a set of equations to represent these quantities (the Hodgkin-Huxley model of the action potential). While this system behaves in the right sort of way, it is not very easy to use it to gain mathematical insight as to why there is a threshold where an initial impulse is large enough to cause a 'spike'.

Here we move to a second layer of models (so this is a really a model of a model³⁵). This type of model was developed independently by two researchers, hence it is the *Fitzhugh-Nagumo* model. Students of dynamical systems will see similarities with the van der Pol oscillator, but this has a subtle but important modification.

Here is the Fitzhugh-Nagumo model:

$$\begin{aligned}\dot{u} &= c \left(v + u - \frac{1}{3}u^3 + z(t) \right) \\ \dot{v} &= -\frac{1}{c} (u - a + bv)\end{aligned}$$

where


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

The term $z(t)$ represents external input, which might be applied at some time and not at others. To start with, we analyse the system with $z(t) = 0$. Note that the u and v are

³⁵one could argue it is actually models all of the way down: the Hodgkin-Huxley model is itself a simple mathematical description of our understanding of the electrochemistry, and of course our understanding is an idealised model of how some imagined neuron might fire, and 'neuron' is catch-all term for a variety of cells, and so it goes on until we are in the land of philosophy, which means we probably should have stopped a while back, probably before reading this footnote.


not quantities that we can easily relate back to the original biochemical model: they are not directly comparable to potassium levels, voltage difference or other physical quantities, but the idea is that this model will illustrate the mathematical processes at work.

The nullclines are themselves fairly simple (u on horizontal axis, v on vertical):




However we must check how they sit relative to each other. Given the $\dot{v} = 0$ nullcline (in purple above) has gradient $-1/b$ it is not too hard to check this is steeper than the downwards section of the $\dot{u} = 0$ nullcline. So there must be a single intersection. The range in a above is rigged to ensure that at $u = 1$ the $\dot{v} = 0$ nullcline is between $-2/3$ and 0 .

For $c \gg 1$, we have fast-slow dynamics again. The u dynamics are fast and so we move quickly near-horizontal away from the $\dot{u} = 0$ nullclines. The vertical adjustment happens more slowly. This means we crawl along just outside the cubic nullcline much of the time, but if we reach the maximum or the minimum, then we ‘fall off’ and zip across until we hit another section of the cubic.



For this plot, $a = 0.7$, $b = 1$ and $c = 5$.



Same parameters, just a zoom-in on the fixed point showing how all trajectories head to it.

Depending on where we start, we might first do a stint moving up the left side of the cubic, then a hop across, and then drift down to the single fixed point. Unlike the van der Pol oscillator, that's it! So long as $z(t) = 0$ we just stay there.


Finding the value of u and v at the fixed point involves solving a cubic of course, but luckily we do not actually need to do this. It suffices to say the fixed point is at (u^*, v^*) and to note that by the arrangement of nullclines, we see that $u^* > 1$. Check the Jacobian at the fixed point to confirm stability (though we are already sure it must be stable):

$$J_{(u^*, v^*)} = \begin{pmatrix} c(1 - u^{*2}) & c \\ -1/c & -b/c \end{pmatrix} = \begin{pmatrix} - & + \\ - & - \end{pmatrix}$$

which is clearly a matrix with negative trace and positive determinant, so indeed the fixed point is stable.


Suppose now that the system has settled to the fixed point, i.e. suppose the neuron is at rest. Now apply some input: $z(t) = -V_0$, where V_0 is a (possibly small) positive constant. We can see that the $\dot{v} = 0$ nullcline is completely unchanged, but the $\dot{u} = 0$ nullcline shifts up by that constant. There are now three cases we must consider. In each of the plots for the cases below, the original cubic nullcline is shown in blue, and the new shifted one in red, with the new fixed point in red. For the phase diagrams on the left, the axes have been left out for clarity. The plot on the right shows v against time.

(i) No spike:




Quite a dull case: basically the fixed point has not moved too much and the system just corrects to the new system (same for $V_0 < 0$).

(ii) Enough impulse for one spike:



This is the excitable behaviour! The trajectory from the old fixed point skims below the new nullcline and thus has to go the whole way across to the other branch of the cubic, climb up and then back to the right and climb down. This corresponds to spike in v . After that, we arrive at the new fixed point, which is much as before.

(iii) Enough impulse for multiple spikes:



Now V_0 is large enough that it is actually a structural change to the system. The new fixed point is not stable: the nullcline is now shifted so far that the fixed point has moved beyond the minimum to the middle portion of the cubic. So, we start off as in case (ii): from the original fixed point and do a loop, but now there is no fixed point to land on, so it goes around again, and again, and again. It will keep going until $z(t)$ is switched off again.

Exercise 27: Find the threshold for V_0 for repeated firing (case (iii)).

So this Fitzhugh-Nagumo model illustrates how a mathematical system can be excitable. We can gain the insight that we need a nullcline which is like a cubic, with multiple branches for a fast-slow solutions to follow. We can also see how a small external input can shift the system so that there is a large single transient, or even so that the system itself has changed and there are cycles. The fast-slow dynamics make it easy to illustrate what is going on, but it turns out that a large value of c is not required³⁶.

end of
lecture 10

³⁶Have a go with the mathematica simulation.

2 Stochastic systems

2.0 Preliminaries

2.0.0 Revision: discrete probabilities and generating functions

We will be considering discrete systems where the population size takes non-negative integer values $(0, 1, 2, 3, \dots)$. This could be representing the number of individuals, cells, or perhaps molecules. We will usually³⁷ write the random variable as N and particular values it takes as n .

$$P(n, t) = \text{Probability}(N = n \text{ at time } t)$$

For clarity, often we will use the simpler notation of p_n to mean $P(n, t)$ when it is unambiguous. As N can only take positive integer values, summing probabilities over all of these values:

$$\sum_{n=0}^{\infty} p_n = 1.$$

We will also be making extensive use of generating functions:

$$\phi(s, t) = \sum_{n=0}^{\infty} s^n p(n, t) = \sum_{n=0}^{\infty} s^n p_n = \langle s^N \rangle$$

Where the angle brackets $\langle \cdot \rangle$ denote the expected value. Given a generating function, we can recover p_n by looking for the coefficient of s^n . Recall some of the useful properties of generating functions, particularly by evaluating at special values of s :

$$\phi(1, t) = \sum_{n=0}^{\infty} p_n = 1$$

$$\phi(0, t) = p_0 = \text{Probability that the population has died out}$$

We can also differentiate with respect to s :

$$\left. \frac{\partial}{\partial s} \right|_t \phi = \phi_s = \sum_{n=0}^{\infty} n s^{n-1} p_n$$

and from now on take partial derivatives with respect to s to mean t treated as constant, and vice-versa. And using this derivative of the generating function we can find the mean by evaluating at $s = 1$:

$$\phi_s(1, t) = \sum n p_n = \langle N \rangle = \text{mean pop.} = \mu$$

³⁷I will try to stick to this throughout these notes just for the sake of clarity (capitals for the random variable, lower case for the number it takes in an equation, usually as an index being summed over), but usually easy to interpret what is meant in any given expression, so not essential to be super careful about all this.

(where μ is just the most commonly used letter for the mean). We do not need to be super careful about putting limits on the sums so long as we say $p_n = 0$ for $n < 0$. Differentiating a second time then evaluating at $s = 1$:

$$\phi_{ss}(1, t) = \sum n(n-1)p_n = \langle N(N-1) \rangle = \langle N^2 \rangle - \langle N \rangle = \langle N^2 \rangle - \mu$$

and recall the variance is given by

$$\text{Var}(N) = \langle (N - \mu)^2 \rangle = \langle N^2 - 2\mu N + \mu^2 \rangle = \langle N^2 \rangle - 2\mu \langle N \rangle + \mu^2 = \langle N^2 \rangle - \mu^2$$

We know μ from the generating function already, so we now have variance:

$$\text{Var}(N) = \phi_{ss}(1, t) + \mu - \mu^2.$$

2.0.1 Why bother?

So far, we have only considered deterministic systems, and also we have focussed on continuous variables for things like population size. We know that we are doing this for convenience of course. The number of individuals in a population should be a non-negative integer. To some extent, all biological processes are subject to random fluctuations. However by taking these simplifications, we are in a realm where we have a great deal of mathematical machinery we can use (differential equations, and dynamical systems in general). It means we can write down models and explore them without much difficulty. We can also understand them clearly enough to gain insights which often extend far beyond the particular model we have in front of us. We might be confident that some of our results (for example the concept of R_0 in disease dynamics and hence the idea of a vaccination threshold) will hold even if we were to add more layers of complexity to the models.

However, what have we missed by ignoring random fluctuations and forcing population numbers to be continuous? Perhaps the most obvious problem is that populations that are decaying and getting very small still cannot ever reach zero, e.g.:

$$\frac{dN}{dt} = -\lambda N \quad \implies \quad N = N_0 e^{-\lambda t}$$

($\lambda > 0$) and even as t becomes very large, still N is positive. Does this matter? Well, yes it does if something changes later and the population size, which should have been zero, can claw its way back from tiny numbers to being large.


A classic case of this was in a model for rabies in foxes proposed by Murray and others. In a published paper: the scenario of a rabies introduced to the south of England was explored by a deterministic diffusion model. As one might expect, rabies spreads out almost radially over a few years. However it doesn't end there: after some years, rabies *reappears* from the same initial location and spreads out again (though less dramatically). This second introduction was not put in there by the modellers, but

just happens in the simulations. This was explained by Mollison in a paper in 1991³⁸: the density of infected foxes at the origin location never actually died down to zero, though it did get to a very small value: 10^{-18} of a fox per square kilometre. Anyone reasonable would describe this as essentially *no* rabid foxes, but mathematically it is still a quantity that can grow back when conditions allow (renewed susceptibility when the fox population has turned over for a few years). This Mollison called the ‘atto-fox’, and this has become a useful term to describe this modelling issue³⁹. Note that this is not a numerical issue from inaccuracy in computation, but one which comes from the very nature of continuous population deterministic models. Does this mean that original model was worthless? I would argue it still has use, but attention to results from simulations should stop after the first wave.

So although we do not always need to use stochastic models, we should be vigilant for atto-foxes and other artifactual effects in our deterministic models. And to watch out for these effects, we must be aware of what they can be. As well as this issue with small numbers not going to zero, we can already guess (from thinking about excitable systems already studied) that in some special situations, small fluctuations can lead to large effects. Further, we also do not always expect ‘average’ behaviour to in fact be representative of typical behaviours. In addition, we will see in the later parts of this section that small fluctuations in two quantities are not always just independent noise: they can actually covary if the dynamics are interlinked.

2.0.2 The first step

Consider first this very simple system: we have states A and B only, and we are always in exactly one of those states. We jump from A to B at probability rate λ . We can make a diagram like this:



Where dots (or nodes) denote states, arrows are the possible ‘jumps’ and the arrows are labelled with the rates. There are two ways to proceed from here. The less formal approach (taken in lectures) is simply to go from the diagram to consider flows of probability, an idea which might be familiar from continuous time Markov Chains from Ib . For completeness, the slightly more formal way is included in these notes [in blue](#). If the more formal approach helps you understand things better than great, if not then then you can safely skip it. In both approaches, we introduce notation $P(A, t)$ to mean probability we are in state A at time t .

³⁸Mollison, D. *Dependence of epidemic and population velocities on basic parameters*. 1991, Mathematical biosciences 107 (2)

³⁹I have also heard this phenomena described as ‘nano-hawks’

More formal What does probability rate λ actually mean? Essentially it is a probability per unit time of making a jump so that for small δt :

$$P(\text{jump by } t + \delta t \mid \text{in } A \text{ at time } t) = \lambda \delta t + \mathcal{O}(\delta t^2)$$

One could justify this by essentially saying this is the definition of probability rate (a bit like a probability density function $f(x) dx$). But perhaps more intuitive is to think of it as being an instantaneous rate λ until it happens. It might have already happened in δt so then the rate is no longer λ . In fact that helps us to see why the correction is order δt^2 .

We can build up a differential equation for the probability by considering the probability of being in A at time $t + \delta t$ conditional on each of the possibilities at time t , and what happens in the time period between of length δt :

$$\begin{aligned} P(A, t + \delta t) &= P(A, t) P(\text{stayed in } A \text{ for } \delta t) + P(B, t) P(\text{went from } B \text{ to } A \text{ in } \delta t) \\ &= P(A, t) [1 - P(\text{left } A \text{ during } \delta t)] + P(B, t) [0] \\ &= P(A, t) [1 - \lambda \delta t] + \mathcal{O}(\delta t)^2 \end{aligned}$$

Then strategic rearrange and divide by δt :

$$\frac{P(A, t + \delta t) - P(A, t)}{\delta t} = -\lambda P(A, t) + \mathcal{O}(\delta t)$$

and taking the limit $\delta t \rightarrow 0$:

$$\frac{d}{dt} P(A, t) = -\lambda P(A, t)$$

Alternatively, working straight from the diagram, it would not be unreasonable to go directly for the same equation immediately, i.e.

$$\frac{d}{dt} P(A, t) = -\lambda P(A, t).$$

We can build a similar differential equation for $P(B, t)$, but we do not need it as in fact we must be in A or B , so $P(A, t) + P(B, t) = 1$

We can solve this if we have an initial condition. Let us suppose we are at A initially so $P(A, 0) = 1$:

$$P(A, t) = e^{-\lambda t}$$

and hence also

$$P(B, t) = 1 - e^{-\lambda t}.$$

This rather unsurprising result lets us quickly check the distribution of times to jump. Let T be the time when we jump from state A to B . Then we can find the cumulative distribution for T by noting we are in B if T has passed already:

$$P(T < t) = P(B, t) = 1 - e^{-\lambda t}$$

We can use this to find the probability density function:

$$f(t) = \frac{d}{dt}(1 - e^{-\lambda t}) = \lambda e^{-\lambda t}$$

which by now we see is the exponential distribution. And as always:

$$\int_0^\infty f(t) dt = 1.$$

We can find the expected jump time

$$\langle T \rangle = \int_0^\infty t f(t) dt = \frac{1}{\lambda}$$

and the variance of the jump time

$$\text{Var}(T) = \sigma^2 = \langle T^2 \rangle - \langle T \rangle^2 = \int_0^\infty t^2 f(t) dt - \frac{1}{\lambda^2} = \frac{1}{\lambda^2}.$$

and we could square root this for the standard deviation σ and we see that it is equal to the mean.

2.1 Discrete population sizes

2.1.1 Single populations

We now study the pure import model. There is a probability rate λ of adding one individual to the population, no matter what the current population size is. We start with zero population. Again, represent states as dots and jumps as arrows:



And writing $P(n, t)$ to mean probability of the population size (N) being n at time t , then the initial condition just says that $P(0, 0) = 1$ and $P(n, 0) = 0$ for $n > 0$.

Again we can set things up by considering where we are at $t + \delta t$ starting from t :

$$P(n, t + \delta t) = P(n, t) \times (1 - \lambda \delta t + \mathcal{O}(\delta t^2)) + P(n - 1, t) \times (\lambda \delta t + \mathcal{O}(\delta t^2)) + \mathcal{O}(\delta t^2).$$

In theory we should be including a term for every possible previous state on the right hand side here, but that last $\mathcal{O}(\delta t^2)$ takes care of all of the others. For example, starting from $n - 2$ we would have to jump *twice* during δt to end up in n . Other states are even more remote, or even impossible to get to n : that last $\mathcal{O}(\delta t^2)$ is a big bin for them all.

As before, rearrange and divide by δt to prepare for a derivative:

$$\frac{P(n, t + \delta t) - P(n, t)}{\delta t} = -\lambda P(n, t) + \lambda P(n - 1, t) + \mathcal{O}(\delta t).$$

Take the limit $\delta t \rightarrow 0$.

Or again, work directly from the diagram:

$$\frac{d}{dt}P(n, t) = \lambda P(n - 1, t) - \lambda P(n, t)$$

and at this point it would be sensible to use the shorter notation $p_n = P(n, t)$:

$$\dot{p}_n = \lambda(p_{n-1} - p_n) \quad (17)$$

and we can say this holds for all $n \geq 0$ by the nice convention of saying $p_n = 0$ for $n < 0$.

Equation (17) is an example of a *master equation*, which is just a fancy way of saying a differential equation for p_n . It is a slightly curious object though as it is (continuous) differential equation in t , but a (discrete) difference equation in n .

In this case, we see that the master equation gives a system of differential equations that we can actually fully solve explicitly.

Exercise 28: Set $p_0 = 1$, $p_n = 0$ for $n > 0$ at $t = 0$. Solve the master equation inductively to get

$$p_n = \frac{(\lambda t)^n}{n!} e^{-\lambda t}$$

Generally we cannot fully solve in this way, in which case it is often a good idea to try a generating function. Recall:

$$\phi(s, t) = \sum_{n=0}^{\infty} s^n p_n = \langle s^N \rangle$$

Then we can differentiate ϕ with respect to t (keeping s fixed), and substitute in using the master equation:

$$\begin{aligned} \frac{\partial \phi}{\partial t} &= \sum s^n \dot{p}_n = \sum s^n \lambda (p_{n-1} - p_n) = \lambda \sum s^n p_{n-1} - \lambda \sum s^n p_n \\ &= \lambda \sum s^{n+1} p_n - \lambda \sum s^n p_n \\ &= \lambda s \phi - \lambda \phi \\ &= (s - 1) \lambda \phi \end{aligned}$$

And we can integrate this (remember that s was treated as a constant):

$$\phi = A(s) e^{(s-1)\lambda t}$$

To resolve $A(s)$, consider $t = 0$:

$$\phi(s, 0) = \sum s^n p_n(0) = p_0(0) + s p_1(0) + s^2 p_2(0) + s^3 p_3(0) + \dots$$

and $p_0(0) = 1$ and the rest of the $p_n(0) = 0$ for $n > 0$, so $\phi(s, 0) = 1$. And from this we see that $A(s) = 1$ for all s . So here we have the full solutions for the generating function

$$\boxed{\phi(s, t) = e^{(s-1)\lambda t}}$$

Great. So what can we do with ϕ ? Well, it encodes all of the p_n , so we could use it to recover the p_n by taking the coefficient of s^n . This we could do by Taylor series (recognition or differentiating n times and evaluating at $s = 0$):

$$\begin{aligned} p_n &= \text{Coef of } s^n \text{ in } \phi = \text{Coef of } s^n \text{ in } e^{-\lambda t} e^{\lambda t s} \\ &= \text{Coef of } s^n \text{ in } e^{-\lambda t} \left(1 + (\lambda t s) + \frac{1}{2!} (\lambda t s)^2 + \dots + \frac{1}{n!} (\lambda t s)^n + \dots \right) \\ &= e^{-\lambda t} \frac{(\lambda t)^n}{n!} \end{aligned}$$

Often we are not very interested in p_n itself, but rather want things like the mean and variance, which can be found directly without having to go via the p_n . For the mean:

$$\mu = \langle N \rangle = \left. \frac{\partial \phi}{\partial s} \right|_{s=1} = \lambda t e^{(s-1)\lambda t} \Big|_{s=1} = \lambda t$$

and for the variance:

$$\langle N(N-1) \rangle = \left. \frac{\partial^2 \phi}{\partial s^2} \right|_{s=1} = (\lambda t)^2 e^{(s-1)\lambda t} \Big|_{s=1} = (\lambda t)^2$$

so

$$\sigma^2 = \text{Var}(N) = \langle N^2 \rangle - \langle N \rangle^2 = \left. \frac{\partial^2 \phi}{\partial s^2} \right|_{s=1} + \mu - \mu^2 = (\lambda t)^2 + (\lambda t) - (\lambda t)^2 = \lambda t.$$

This could be used to say something about the size of small fluctuations. Typical fluctuations are of the order of the standard deviation $\sigma = \sqrt{\lambda t}$. Typical population sizes are of the order of the mean $\mu = \lambda t$, so relative to population size, fluctuations go like $(\lambda t)^{-1/2}$ so fluctuations become less important for this system as time goes on.


One further note on this simple import system: if we were to construct an analogous deterministic system, it would surely be

$$\frac{dN}{dt} = +\lambda$$

and with $N(0) = 0$ the solution would be $N = \lambda t$. This agrees with the mean of the stochastic system. Note: this agreement does not always hold⁴⁰.

⁴⁰You might like to try and identify what is special about this system. It might help to think about for what functions f is it true that $\langle f(N) \rangle = f(\langle N \rangle)$.

Import and death model Here we use the same model for import (probability rate λ of adding a new individual, including when at $n = 0$). Now we add a *per capita* probability rate of death β so the total rate is βn :



By now, we should be able to build the master equation by considering ‘flows’ of probability between states (the approach with δt would be fine but rather tedious). Focussing on state n in the diagram above, we can see two arrows going in and two going out, giving us the four terms:

end of
lecture 12

$$\begin{aligned}\dot{p}_n &= \lambda p_{n-1} - \lambda p_n + \beta(n+1)p_{n+1} - \beta n p_n \\ &= \lambda(p_{n-1} - p_n) + \beta[(n+1)p_{n+1} - n p_n]\end{aligned}$$

Here, we go straight for the equation satisfied by the generating function:

$$\frac{\partial \phi}{\partial t} = \sum s^n \dot{p}_n = \lambda \sum s^n p_{n-1} - \lambda \sum s^n p_n + \beta \sum s^n (n+1) p_{n+1} - \beta \sum s^n n p_n$$

and as always, it is a good idea to split into separate sums, so we can shift the index on each one separately, aiming for p_n in each:

$$\frac{\partial \phi}{\partial t} = \lambda \sum s^{n+1} p_n - \lambda \sum s^n p_n + \beta \sum n s^{n-1} p_n - \beta \sum n s^n p_n$$

The first two sums are much as before ($s\phi$ and ϕ), but for the latter two, we must differentiate ϕ with respect to s to get the bonus n in the coefficients to appear:

$$\frac{\partial \phi}{\partial s} = \phi_s = \sum n s^{n-1} p_n$$

so we can immediately recognise the third sum as ϕ_s . The fourth one is the same thing again but with a spare factor of s :

$$\frac{\partial \phi}{\partial t} = \lambda s \phi - \lambda \phi + \beta \phi_s - \beta s \phi_s$$

and we can factorise⁴¹ out $(s - 1)$ to get

$$\phi_t = (s - 1) [\lambda \phi - \beta \phi_s] \quad (18)$$

We should also specify an initial condition. Again we will take $N = 0$ initially so that $\phi(s, 0) = 1$.

⁴¹This keeps happening and should not be a surprise. We know that $\phi = 1$ for $s = 1$, so $\phi_t = 0$ for $s = 1$.

Exercise 29: Get from the dots, arrows and rates diagram above to equation (18) without lecture notes

We can actually fully solve for ϕ again in this case. Take a strategic form:

$$\phi = e^{(s-1)f(t)}$$

where $f(t)$ is a function yet to be determined and substitute into equation (18):

$$(s-1)f'(t)\phi = (s-1)[\lambda\phi - \beta f(t)\phi]$$

and then it is clear that this form was designed to cancel the $(s-1)\phi$ and we're left with a DE in purely f :

$$f'(t) = \lambda - \beta f(t) \quad \implies \quad f(t) = \frac{\lambda}{\beta} + Ae^{-\beta t}$$

The initial condition that $\phi = 1$ at $t = 0$ corresponds to $f(0) = 0$, and this determines A , which fully determines f and hence we have

$$\phi(s, t) = e^{\frac{\lambda}{\beta}(s-1)(1-e^{-\beta t})}$$

and the double exponential might not be pretty, but it is nice to have fully solved for ϕ .

Exercise 30: Use this ϕ to show that the mean is given by

$$\langle N \rangle = \frac{\lambda}{\beta}(1 - e^{-\beta t})$$

and the variance is actually the same

$$\text{Var}(N) = \frac{\lambda}{\beta}(1 - e^{-\beta t})$$

Suppose we are now interested in what happens to this system at large time, so we are looking for a steady state solution and not worrying about initial condition. We could use our solution and let $t \rightarrow \infty$. However, to illustrate a more generally useful technique we will suppose that we do not actually have the solution for ϕ . We can go back to equation (18) and seek the solution that is independent of time (i.e. set $\phi_t = 0$):

$$0 = (s-1)[\lambda\phi - \beta\phi_s]$$

which is just a DE for ϕ in s , which we can solve:

$$\phi_s = \frac{\lambda}{\beta}\phi \quad \implies \quad \phi = Ae^{\frac{\lambda}{\beta}s}$$

and we cannot use the time initial condition here, but need something for a fixed value of s . The one to go for here is always $\phi = 1$ at $s = 1$. So $A = e^{-\frac{\lambda}{\beta}}$ and we have the steady state solution for ϕ :

$$\phi = e^{\frac{\lambda}{\beta}(s-1)}$$

and we see this is reassuringly consistent with what we would get with $t \rightarrow \infty$ in the general solution above.


We can also quickly see the steady state p_n by pulling out the appropriate coefficient of s^n :

$$p_n = \text{Coef of } s^n \text{ in } \phi = e^{-\frac{\lambda}{\beta}} \left(\frac{\lambda}{\beta} \right)^n \frac{1}{n!}$$

(which we recognise as the Poisson distribution).

Aside on simulations

As always, it helps build intuition for what this is all about by looking at the output of simulations to get a feel for how these things behave. Here's one for this import and death model from this last section with $\lambda = 20$ and $\beta = 1$:



The dotted line is $\lambda/\mu = 20$, what we expect the mean to settle to. Indeed you can see that it heads quickly away from the initial condition and bumbles around near the mean.

You could even say the standard deviation looks to be about 4 or 5, which is consistent too with long term behaviour ($\sigma = \sqrt{20}$).

This programming approach is certainly not examinable, but it is actually very easy to simulate models like these. The *Gillespie algorithm* is both simple to code and fast to run. It relies on some nice properties of the exponential distribution. Suppose current time is t_i and population size is n_i . Then from the model we have some probability rates for jumps to other states. Each of these possible jumps happens at exponentially distributed time, but only the first one is actually the one that happens. The nice property is that the time for the *first* of a bunch of exponential distributions is itself an exponential distribution, with rate just sum of the others. Simulating an exponential distribution is super easy: just pick uniform random real in $(0, 1)$, take minus log of it, and divide by rate.


Then which event is it that happens? Also easy to run: just pick another uniform random and use it to choose event by weighting the unit interval in proportion to the rates of each possible event. Then you have $t_{i+1} = t_i + \Delta t$ and n_{i+1} is whatever it is after the chosen event.

Then rinse and repeat a number of times, but how many? Can either do a set number of events (some max for i), or until time is above some threshold (some max for t), but need to check some event is possible, i.e. total rate has not hit zero ($r_{tot} \neq 0$). That last possibility can't happen in this model, but could in others later, if we are in some stuck state (extinction, below). Here it is as pseudo code for this model: start with $i = 0$, choose some n_0 and probably $t_0 = 0$:

- Loop back to here, got n_i, t_i at some step i
- Compute rates for events $r_b = \lambda, r_d = \beta n_i$. Total rate $r_{tot} = r_b + r_d$.
- Generate a couple of uniform $(0, 1)$ random numbers: $random1, random2$
- Time to next event $\Delta t = -\log(random1)/r_{tot}$
- Choose which event: if $random2 < r_b/r_{tot}$ it is a birth, else a death
- update variables: $t_{i+1} = t_i + \Delta t$ and $n_{i+1} = n_i + 1$ if birth, $n_{i+1} = n_i - 1$ if death
- $i = i + 1$, loop back

Variance and step size

A generalisation of the import and death model allows us to see how the variance of the steady state depends on step size. The death rate is still β per capita, but now we suppose that M individuals are added each time, rather than just one, but still at rate λ . (Take M to be some positive integer of course.) How do we expect the mean and variance in steady state to change?



This works much as above, so let us jump straight to the master equation:

$$\dot{p}_n = \lambda (p_{n-M} - p_n) + \beta [(n+1) p_{n+1} - n p_n]$$

and then find the PDE satisfied by the generating function ϕ .

Exercise 31: Show

$$\frac{\partial \phi}{\partial t} = \lambda (s^M - 1) \phi - \beta (s - 1) \phi_s$$

We could work directly with the generating function for example to find steady state:⁴²

$$\phi_t = 0 \implies \phi_s = \frac{\lambda}{\beta} \left(\frac{s^M - 1}{s - 1} \right) \phi$$

And this would work well for any given M , as we'd just have a nice polynomial in that bracket. But to illustrate another more general approach, suppose we couldn't progress with the generating function directly, then we could instead think of working with the master equation directly to find differential equations (in time) for the moments (expected N to some power):

$$\begin{aligned} \frac{d}{dt} \langle N \rangle &= \sum n \dot{p}_n \\ &= \lambda \sum n p_{n-M} - \lambda \sum n p_n + \beta \sum n(n+1) p_{n+1} - \beta \sum n^2 p_n \\ &= \lambda \sum (n+M) p_n - \lambda \sum n p_n + \beta \sum (n-1)n p_n - \beta \sum n^2 p_n \\ &= \lambda \sum M p_n + \beta \sum (-n) p_n \\ &= \lambda M - \beta \langle N \rangle \end{aligned}$$

where we have used the master equation to sub in for \dot{p}_n , and then expanded to write one sum per p term, then shifted the index on each sum so the p term is p_n , then it tidies from there. So this is a nice⁴³ differential equation in $\langle N \rangle$. For initial condition, we supposed $N = 0$ at $t = 0$, so also $\langle N \rangle = 0$ at $t = 0$, and we can then easily solve:

⁴²Do note that in stochastic models, steady state means that the p_n are constant, not N . The general behaviour of the system has settled into something, not that the population size has settled to a fixed value.

⁴³actually this all seems rather too nice, and you might be suspicious that there was something special here. Indeed there is! The birth and death rates were only linear in n , which means the coefficients in the master equation are only linear. The equation for time derivative for each moment will end up being just in terms of that moment and lower. If birth or death rates were quadratic in n , we would find $\frac{d}{dt} \langle N \rangle$ depends on $\langle N^2 \rangle$ and so on up, which means we no longer have a closed system.

$$\langle N \rangle = \frac{\lambda M}{\beta} (1 - e^{-\beta t})$$

and we can also see $\langle N \rangle \rightarrow \lambda M / \beta$ as $t \rightarrow \infty$.

A similar approach with higher moments will also work:

$$\begin{aligned} \frac{d}{dt} \langle N^2 \rangle &= \sum n^2 \dot{p}_n \\ &= \lambda \left(\sum n^2 p_{n-M} - \sum n^2 p_n \right) + \beta \left(\sum n^2 (n+1) p_{n+1} - \sum n^3 p_n \right) \\ &= \lambda \sum ((n+M)^2 - n^2) p_n + \beta \sum ((n-1)^2 n - n^3) p_n \\ &= \lambda \sum (2nM + M^2) p_n + \beta \sum (-2n^2 + n) p_n \\ &= \lambda (2\langle N \rangle M + M^2) + \beta (-2\langle N^2 \rangle + \langle N \rangle) \\ &= \lambda M^2 + (2\lambda M + \beta) \langle N \rangle - 2\beta \langle N^2 \rangle \end{aligned}$$

and this also gives us a sensible DE as we already have solved to get $\langle N \rangle$ as a function of time. If we write $\langle N^2 \rangle = X$ then we have a DE in this form:

$$\dot{X} = -2\beta X + a + b e^{\beta t}$$

where the coefficients a and b are some constants that can be deduced from above.

Exercise 32: Show that at steady state

$$\langle N^2 \rangle = \frac{\lambda}{2\beta} \left(1 + \frac{2\lambda}{\beta} \right) M^2 + \frac{\lambda}{2\beta} M.$$

Note, you have a choice here. Fast method (which would be fine) is to just use the differential equation at steady state to find the equilibrium value directly. Full method (good for extra practice) is to fully solve the DE to get $\langle N^2 \rangle$ as a function of t , and then let $t \rightarrow \infty$.


Either way... show also that if $M \gg 1$,

$$\text{var}(N) = \frac{\lambda}{2\beta} M^2.$$

Intuition and effect of step size: now think about the effect of changing M . Remember that the model is to add M individuals at rate λ and they die at some per capita rate β . It is intuitive that the population will head towards some balance between introduction and death: if the population is much larger than this balance then death rate

will dominate, if much smaller then the birth rate will dominate. So, this matches with expected N being $\lambda M/\beta$. We can also imagine that once the system is near that, then it will just stay vaguely near, but some random fluctuations coming from the stochastic nature of the model.

Suppose now that we make M ten times as big, and shrink λ by a corresponding factor of 10. Then intuitively we have the same average rate of individuals being added, and indeed we see $\langle N \rangle$ is unchanged (proportional to λM). However, we see that the variance (proportional to λM^2 for large M) will increase, and indeed ten times as much if we started with large M . Hopefully, this will also seem intuitive: we've made the stochastic process more 'lumpy' by having a rarer but larger step. This idea will be useful later!



Small M on left, large M on the right, also changing λ to keep the mean constant. There's clearly more variability with the larger M .

end of
lecture 13

2.1.2 Extinction

The examples above had a constant import rate, so even when $n = 0$ there is a constant immigration into the system. It turns out that this is crucial to get non-trivial steady states. In models with a *closed* population, no immigration is possible, so there is no way to get from $n = 0$ to $n = 1$, and once the population has gone extinct, it stays that way forever. In this section, we will build up some typical properties we might expect from a single population stochastic model, and we see that extinction is inevitable.

Here are three reasonable properties for a closed population model:

1. No births/immigration from $n = 0$
2. Death is possible for all $n > 0$
3. For large n , the death rate is much larger than the birth rate



Setting b_n and d_n to be birth and death rates respectively from state n , we can translate these properties to something in terms of these rates:

1. $b_0 = 0$
2. $d_n > 0$ for all $n > 0$
3. Either $b_n = 0$ for some n , or $d_n/b_n \gg 1$ for large n

We now divide into two cases to consider extinction:

(i) If $b_n = 0$ for some $n > 0$: then the state space is actually finite $(0, 1, 2, 3 \dots n-1, n)$ so we can use a result from IB Markov Chains: if we have a finite state space and all states communicate with $n = 0$, and $n = 0$ is absorbing (i.e. once we are there, we stay there), then as $t \rightarrow \infty$, $p_0 \rightarrow 1$. In other words, extinction is inevitable.

(ii) If $b_n > 0$ for all $n > 0$: then we seem to have an infinite state space, but actually the same result turns out to be true. Perhaps this is almost intuitive if we remember that one of our conditions is that the death rate far outweighs birth rate large for large n , so it should behave as if the state space is effectively finite. But, here's a slightly more mathematical argument (which of course uses that condition of deaths outweighing births). We set up q_n the probability of extinction eventually, given we start at n . Then we can immediately have that $q_0 = 1$ (this is instant game over: we started at zero). For $n > 0$, we set up a recurrence, again using law of total probability: if we're at $n > 0$ then either a birth or a death next. If there's a birth, then we will be at $n+1$ and the probability of ultimate extinction starting from here is q_{n+1} . Similarly for death next. So:

$$q_n = P(\text{birth next}) q_{n+1} + P(\text{death next}) q_{n-1}$$

and what is the probability of birth or death next? We have the rates b_n and d_n (and know neither is zero in this case), so:

$$P(\text{birth next}) = \frac{b_n}{b_n + d_n}, \quad P(\text{death next}) = \frac{d_n}{b_n + d_n}.$$

Subbing these in, and multiplying through by total rate $b_n + d_n$:

$$(b_n + d_n)q_n = b_n q_{n+1} + d_n q_{n-1}$$

Rearrange to split b_n and d_n :

$$\begin{aligned} b_n(q_n - q_{n+1}) &= d_n(q_{n-1} - q_n) \\ (q_n - q_{n+1}) &= \frac{d_n}{b_n}(q_{n-1} - q_n) \end{aligned}$$

Note that we can't immediately solve this recurrence with only one initial condition, but we can argue through that there is only one feasible solution. Notice the bracket terms are just difference in consecutive q_i , and this recurrence equation holds for all $n > 0$, so chase the bracket on the right down:

$$\begin{aligned} (q_n - q_{n+1}) &= \frac{d_n}{b_n}(q_{n-1} - q_n) \\ &= \frac{d_n}{b_n} \frac{d_{n-1}}{b_{n-1}}(q_{n-2} - q_{n-1}) \\ &= \frac{d_n}{b_n} \frac{d_{n-1}}{b_{n-1}} \cdots \frac{d_1}{b_1}(q_0 - q_1) \\ &= \left(\prod_{i=1}^n \frac{d_i}{b_i} \right) (q_0 - q_1) \end{aligned}$$

And again this holds for all n . Notice the LHS is the difference of two probabilities, so must be between -1 and 1 for all n . Consider the product on the RHS and taking larger and larger n . We can make this product as large as we like⁴⁴. Obviously $(q_0 - q_1)$ doesn't depend on n , but that equation holds for all n , so the only possibility we are left with is $q_0 - q_1 = 0$. So $q_1 = q_0 = 1$.

Now by the above equation, $q_{n+1} - q_n = 0$ for all $n > 0$, so finally we have that $q_n = 1$ for all n . Extinction is inevitable, starting from any n . This completes the second case (infinite state space, but actually effectively finite state space).

It can be shown that a similar result holds for multiple populations. Think about a vector for replacing n , where the elements are population numbers for different species. We should still have similar conditions: the vector of zeros being absorbing, deaths in each species being possible if any members are present, and very large populations being squished down by death rate outweighing the birth rate.

Treat this blue paragraph as starred, and ignore if it doesn't help you... It is intuitive that the same sort of argument as above should work, but here's a slightly different way to think about it (not given in lectures): divide up the infinite state space into three: (i) extinction (all populations zero, a single state), (ii) all populations within sensible size (some very large but finite state space), (iii) at least one population ridiculously large (the rest of the infinite state space). Then think of the possible transitions between these three states. State (iii) has death rates pushing us back towards (ii) and by carefully choosing what we mean by 'ridiculously large' we could make conditions on

⁴⁴and actually this tells us more precisely what we mean by death rate outweighing birth rate for large n . We can see that it would be enough to say $d_n \geq R b_n$ for all $n > \hat{n}$, for some \hat{n} and some $R > 1$.

birth/death rates that give finite expected time until we are back in (ii). From state (ii) we can go to (iii) or (i), and as there are finitely many states, we pop out one way or the other in finite time. Obvious in (i) we stay there. So now we have something similar to a three state Markov chain with single absorbing state, and the other two states communicate with it. Hence, we will end up there with probability 1.

However, for single or multiple populations, the expected time to extinction could be extremely large⁴⁵


2.1.3 Multiple populations

Here we are going to look at a stochastic model with two populations: wildebeest (represented by m) and flies (n). This model is a bit unusual in that we are going to have the wildebeest affecting the dynamics of the flies, but not vice versa. This is not typical, but means we can handle the mathematics fairly easily, and be able to see the consequence of the interactions clearly on the flies.

The wildebeest have population size m , and ‘births’ are rate λ_1 (i.e. this is the import model above, so no longer worry about extinction) and deaths at rate $\beta_1 m$ (so β_1 *per capita*). The flies have population size n , ‘birth’ rate $\lambda_2 m$, and death rate $\beta_2 n$. The weird bit here is that the birth rate is proportional to the number of wildebeest. This is not because wildebeest give birth to flies, but interpret it here as a kind of import rate, and the flies are attracted by the wildebeest. The more wildebeest there are, the higher the rate of flies that join the system.

The current state of the system is given by (m, n) , which we can think of as corresponding to grid coordinates. Then we could do our arrows diagram like this:

⁴⁵It might seem alarming that any stochastic model like this leads to total extinction always, and one might start fretting that life on earth could in principle be represented by such a system, and complications like time-dependence from days and seasons could be ironed out. However, there is a big difference between it being a mathematical inevitability, and it being something to actually lose sleep over. Expected time could be orders of magnitude larger than the age of our solar system. The ‘end of life on earth’ is not on my personal worry list. Extinction of individual species, however, is. If this also interests you, please *do* go and think and read more on this. Here’s one starting point: <https://www.durrell.org/wildlife/wildlife/durrell-index/explore/>.



For the master equation, really best to use the simpler notation $p_{m,n}$ rather than $P(m, n, t)$, which would be very cumbersome.

$$\begin{aligned} \dot{p}_{m,n} = & \lambda_1 [p_{m-1,n} - p_{m,n}] + \beta_1 [(m+1)p_{m+1,n} - m p_{m,n}] \\ & \lambda_2 [m p_{m,n-1} - m p_{m,n}] + \beta_2 [(n+1)p_{m,n+1} - n p_{m,n}] \end{aligned}$$

And using a similar method as above (lots of sums and shifting the indices carefully, we can find differential equations for the time evolution of moments. Note that the sums are over both population variables (m and n). (I'm going to reduce the number of steps I take each time, so if you are struggling to follow, try to put an extra step in, following an earlier example.) Let us start with the expected number of wildebeest, $\langle M \rangle$:

$$\begin{aligned} \frac{d}{dt} \langle M \rangle = \sum_{m,n} \dot{p}_{m,n} = & \lambda_1 [\langle M+1 \rangle - \langle M \rangle] + \beta_1 [\langle (M-1)M \rangle - \langle M^2 \rangle] \\ & \lambda_2 [\langle M^2 \rangle - \langle M^2 \rangle] + \beta_2 [\langle MN \rangle - \langle MN \rangle] \end{aligned}$$

And this tidies up very nicely:

$$\frac{d}{dt} \langle M \rangle = \lambda_1 - \beta_1 \langle M \rangle$$

and this is exactly as we'd expect, given that we already studied the import and death model above, and the dynamics of the wildebeest are just that. The flies just don't matter to the wildebeest.

Now move on to the expected number of flies, $\langle N \rangle$:

$$\begin{aligned}\frac{d}{dt}\langle N \rangle &= \lambda_1 \langle N - N \rangle + \beta_1 \langle MN - MN \rangle \\ &\quad \lambda_2 \langle M(N + 1) - MN \rangle + \beta_2 \langle (N - 1)N - N^2 \rangle \\ &= \lambda_2 \langle M \rangle - \beta_2 \langle N \rangle\end{aligned}$$

Which interesting just appears to have the parameters from the birth and death processes that directly affect the flies, but of course the birth rate includes expected number of wildebeest (which is itself a function of time, remember).

Exercise 33: Continue and get the equivalent equations for the second moments. Should get:

$$\begin{aligned}\frac{d}{dt}\langle M^2 \rangle &= \lambda_1 (2\langle M \rangle + 1) + \beta_1 (-2\langle M^2 \rangle + \langle M \rangle) \\ \frac{d}{dt}\langle MN \rangle &= \lambda_2 \langle M^2 \rangle + \lambda_1 \langle N \rangle - (\beta_1 + \beta_2) \langle MN \rangle \\ \frac{d}{dt}\langle N^2 \rangle &= \lambda_2 (2\langle MN \rangle + \langle M \rangle) + \beta_2 (-2\langle N^2 \rangle + \langle N \rangle)\end{aligned}$$

And supposing we are interested in steady state, we could set the time derivatives to zero and solve for any of these moments. For the population means:

$$\langle M \rangle = \frac{\lambda_1}{\beta_1}, \quad \langle N \rangle = \frac{\lambda_2}{\beta_2} \langle M \rangle = \frac{\lambda_1 \lambda_2}{\beta_1 \beta_2}$$

The variance of the wildebeest will be similar to the import and death model above (hence will be λ_1/β_1) and hence not particularly interesting. The variance of the flies might be more interesting, as the wildebeest dynamics affect it. And using the last equation in the exercise above, we have that:

$$\begin{aligned}\langle N^2 \rangle &= \frac{\lambda_2}{\beta_2} \left(\langle MN \rangle + \frac{1}{2} \langle M \rangle \right) + \frac{1}{2} \langle N \rangle \\ &= \frac{\lambda_2}{\beta_2} \langle MN \rangle + \langle N \rangle\end{aligned}$$


And use this to look at $var(N)$:

$$\begin{aligned}var(N) &= \langle N^2 \rangle - \langle N \rangle^2 \\ &= \frac{\lambda_2}{\beta_2} \langle MN \rangle + \langle N \rangle - \langle N \rangle^2\end{aligned}$$

$$\begin{aligned}
 \text{var}(N) &= \langle N \rangle + \frac{\lambda_2}{\beta_2} (\langle MN \rangle - \langle M \rangle \langle N \rangle) \\
 &= \langle N \rangle + \frac{\lambda_2}{\beta_2} \text{cov}(M, N)
 \end{aligned}$$

recognising the second term as the covariance between M and N . We could then describe the flies variance as having two components: intrinsic variability from the first term, and extrinsic variability from the second term. The first term is just the fluctuations which come from the randomness of the flies own birth and death process directly. If we somehow totally fixed the number of wildebeest and looked at the flies dynamics in isolation, the variance would just be the first term. The second term is to do with the number of wildebeest changing and thus moving the flies birth rate around.

Actually, if we look at a typical output of a simulation, it is possible to see that both populations have fluctuations of course, but if the wildebeest numbers take a bit of a detour from the mean, then flies seem to follow suit. If we didn't already know, we could guess from this that the populations are interacting somehow.



One output of a Gillespie simulation of the model described above. The wildebeest M are plotted in red and the flies N in black. The means (as computed above) for each population is marked with dotted horizontal line of appropriate colour. Look for a place where the red curve moves a bit away from the mean, and check what happens to the black curve. It should be that the flies population starts to drift in the same direction (from increased or decreased birth rate) and thus flies curve is very variable: it has its own intrinsic variability, but also it is being shoved around by the wildebeest fluctuations.

We will come back to these ideas again below, thinking about covariances in general. Here, we can see that we would expect M and N to be covarying positively.

end of
lecture 14

2.2 Continuous population sizes


2.2.1 Fokker-Planck for a single variable

So far for stochastic models we have thought of the population size n as an integer. Hence we are dealing with master equations which are discrete in n (but continuous in time). Here, we recognise that these distributions ‘look like’ a continuous distribution, in some sense. This is particularly when the width of our distribution is much larger than our typical step size (usually ± 1). By convention and for clarity, when we think about continuous population sizes, we will write x rather than n , but these are really the same thing. By making an approximation, we replace our master equation (discrete in n) with a Fokker-Planck equation (continuous in x). Watch carefully below for the moment of approximation!

To do all this, we should set up some general notation for a master equation. Let $W(n, r)$ be the ‘jump’ rate from n to $n + r$. In other words, $W(n, r)$ give the rates of moving from n , taking a step size of r (which can be negative). Very often, W is non-zero only for a very small set of r , such as $r \in \pm 1$. So far, the biggest set of r is $r \in \{-1, 1, M\}$.

It is useful to have a couple of examples in hand that we already have fully cracked while working through this, just to see how things might look in practice:

Example 1: Pure import (see above). Take $N = 0$ at $t = 0$.




Then

$$W(n, 1) = \begin{cases} \lambda & \text{if } r = 1 \\ 0 & \text{otherwise} \end{cases}$$

and we already know that $\langle N \rangle = \text{var}(N) = \lambda t$.

Example 2: Import and death (see above). Also take $N = 0$ at $t = 0$.



Then

$$W(n, 1) = \begin{cases} \lambda & \text{if } r = 1 \\ \beta n & \text{if } r = -1 \\ 0 & \text{otherwise} \end{cases}$$

and we already know that $\langle N \rangle = \text{var}(N) = \frac{\lambda}{\beta} (1 - e^{-\beta t})$.

Now we have this general notation, we can write a general master equation:

$$\frac{\partial}{\partial t} P(n, t) = \sum_r [W(n - r, r) P(n - r, t) - W(n, r) P(n, t)]$$

where the sum over r is a finite number of terms (usually very small). We are using the full notation here just for extra safety. The partial derivative wrt time is just to make clear that n is fixed. We are about to have some derivatives wrt x appear, so this might matter.

Looking at the square bracket in the master equation, it is essentially of the form $[f(n-r) - f(n)]$ where $f(n) = W(n, r) P(n, t)$. Change the n to x and consider Taylor expansion about x :

$$f(x-r) - f(x) = \left(f(x) - r f'(x) + \frac{1}{2} r^2 f''(x) + \dots \right) - f(x) = -r f'(x) + \frac{1}{2} r^2 f''(x) + \dots$$

so applying that to the actual square bracket:

$$[W(x-r, r) P(x-r, t) - W(x, r) P(x, t)] = -r \frac{\partial}{\partial x} W(x, r) P(x, t) + \frac{1}{2} r^2 \frac{\partial^2}{\partial x^2} W(x, r) P(x, t) + \dots$$

and if we drop the \dots , then we are making an approximation that those terms are small, i.e. $|r| \ll |x|$. This is where the approximation happens. We have ditched $r^3 \frac{\partial^3}{\partial x^3}$ and higher. So, after this approximation, the master equation has become

$$\begin{aligned} \frac{\partial P}{\partial t} &= \sum_r -r \frac{\partial}{\partial x} W(x, r) P(x, t) + \sum_r \frac{1}{2} r^2 \frac{\partial^2}{\partial x^2} W(x, r) P(x, t) \\ \frac{\partial P}{\partial t} &= - \sum_r r \frac{\partial}{\partial x} W P + \frac{1}{2} \sum_r r^2 \frac{\partial^2}{\partial x^2} W P \\ \frac{\partial P}{\partial t} &= - \frac{\partial}{\partial x} \left(\sum_r r W \right) P + \frac{1}{2} \frac{\partial^2}{\partial x^2} \left(\sum_r r^2 W \right) P \end{aligned}$$

where in the last step, the differential and sum were reordered (not a problem as the sum is over a finite number of terms). Now all the r are contained in sums within the brackets. The bracket terms depend on x but not t . Their form is actually fully determined by the model: these functions of x are just made out of the W . Call them $A(x)$ and $B(x)$ respectively, i.e. define:

$$A(x) = \sum_r r W \quad \text{and} \quad B(x) = \sum_r r^2 W$$

We will see below that A can be interpreted as a ‘mean drift’ or advection, some sort of average movement. B is something to do with spreading, or diffusion (indeed it looks that way in the equation below). So finally, we have the Fokker-Planck equation (FPE):

$\frac{\partial P}{\partial t} = - \frac{\partial}{\partial x} (A P) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (B P) \quad \text{Fokker-Planck Equation (FPE)}$

Example 2 continued:

$$A = \sum_r rW(x, r) = (-1)(\beta x) + (+1)(\lambda) = \lambda - \beta x$$

$$B = \sum_r r^2W(x, r) = (-1)^2(\beta x) + (+1)^2(\lambda) = \lambda + \beta x$$

And just sub these in to get the FPE:

$$\frac{\partial P}{\partial t} = -\frac{\partial}{\partial x}((\lambda - \beta x)P) + \frac{1}{2}\frac{\partial^2}{\partial x^2}((\lambda + \beta x)P)$$

The FPE can be used to see how expected values of x , x^2 or indeed any function of x evolve in time. Here we will do expectation of a general function $f(x)$:

$$\begin{aligned}\frac{d}{dt}\langle f(x) \rangle &= \frac{d}{dt} \int f(x)P(x, t)dx \\ &= \int f(x) \frac{\partial P}{\partial t} dx \quad \text{now sub using FPE} \\ &= \int f(x) \left[\frac{\partial}{\partial x}(AP) + \frac{1}{2} \frac{\partial^2}{\partial x^2}(BP) \right] dx\end{aligned}$$

and at this point, sensible to think about where we would like to get to. Integral wrt x of something times P is the expected value of the something. So we'd like to dig the P out from under the partial x derivatives. The way to do that is integrating by parts (once on first time, twice on second). There are some boundary terms from the parts integration, but all involve P or $\frac{\partial P}{\partial x}$ as $x \rightarrow |\infty|$, and we assume that these tend to zero.

$$\begin{aligned}\frac{d}{dt}\langle f(x) \rangle &= \int f(x) \frac{\partial}{\partial x}(AP) dx + \frac{1}{2} \int f(x) \frac{\partial^2}{\partial x^2}(BP) dx \\ &= - \int \left(\frac{\partial}{\partial x} f(x) \right) (AP) dx + \frac{1}{2} \int \left(\frac{\partial^2}{\partial x^2} f(x) \right) (BP) dx \\ \langle f(x) \rangle' &= - \int A f' P dx + \frac{1}{2} \int B f'' P dx \\ \langle f(x) \rangle' &= \langle A f' \rangle + \frac{1}{2} \langle B f'' \rangle\end{aligned}$$

And now we can put in particular $f(x)$ and learn about the time evolution of the mean and variance. With $f(x) = x$ (so $f' = 1$, $f'' = 0$):

$$\langle x \rangle' = \langle A \rangle$$

and this shows how A therefore gives the drift of the mean. This is not hugely surprising as A is just sum of step size times step probability, so in some sense it captures the 'average' behaviour of the system.

Now to work towards the variance, set $f(x) = x^2$ (so $f' = 2x$, $f'' = 2$):

$$\begin{aligned}\langle x^2 \rangle' &= \langle A 2x \rangle + \frac{1}{2} \langle B 2 \rangle \\ &= 2 \langle Ax \rangle + \langle B \rangle\end{aligned}$$

so the variance time evolution is:

$$\begin{aligned}\frac{d}{dt} \text{var}(x) &= \frac{d}{dt} (\langle x^2 \rangle - \langle x \rangle^2) \\ &= \langle x^2 \rangle' - 2 \langle x \rangle \langle x \rangle' \\ &= 2 \langle Ax \rangle + \langle B \rangle - 2 \langle x \rangle \langle A \rangle \\ &= \langle B \rangle + 2 \text{cov}(A, x)\end{aligned}$$

So we can see that the rate of change involves A and how it covaries with x , but also the B term is now involved, and hence B contributes to how the population size distribution spreads, but not to its mean.

Example 2 continued further:

$$\langle x \rangle' = \langle A \rangle = \lambda - \beta \langle x \rangle$$

$$\begin{aligned}\frac{d}{dt} \text{var}(x) &= \langle B \rangle + 2 \langle Ax \rangle - 2 \langle x \rangle \langle A \rangle \\ &= \lambda + \beta \langle x \rangle + 2 \langle \lambda x - \beta x^2 \rangle - 2(\lambda - \beta \langle x \rangle) \langle x \rangle \\ &= \lambda + \beta \langle x \rangle - 2\beta \text{var}(x)\end{aligned}$$

which can easily be solved for steady states:

$$\langle x \rangle = \frac{\lambda}{\beta}, \quad \text{var}(x) = \frac{\lambda}{\beta}.$$

Interestingly the mean and variance at steady state match what we had for the full master equation, even though Fokker-Planck is an approximation. If we asked what higher terms would be if we continued the approximation earlier, we can see why this might be. The higher terms would involve third and higher derivatives with respect to x . After doing integration by parts, these end up on f . If we're looking for mean and variance, we'd use $f = x$ and $f = x^2$, so *third and higher derivatives would vanish*. So, that gets the first two moments. What happens for higher moments?

Exercise 34: Continue Example 1. Use the FPE to find $\langle x \rangle$, $\langle x^2 \rangle$, $\langle x^3 \rangle$ and $\langle x^4 \rangle$ (explicitly as functions of time).

Go back to the full master equation and find $\langle n \rangle$, $\langle n^2 \rangle$, $\langle n^3 \rangle$ and $\langle n^4 \rangle$.

Checkpoint:

$$\langle x^3 \rangle = (\lambda t)^3 + 3(\lambda t)^2, \quad \langle n^3 \rangle = (\lambda t)^3 + 3(\lambda t)^2 + (\lambda t)$$

You should find the first two moments match, but the next two do not, but actually they are a reasonable match for large t , why is this?

2.2.2 Multivariate Fokker-Planck

All of this can be extended to multivariate systems (for multiple populations). Much of the algebra works through in similar ways, just sometimes with vectors and matrices now. We do however get a nice result at the end of this, i.e. a relatively simple equation to find how populations covary with each other at steady state. We started to think about this with the wildebeest and flies example above (section 2.1.3). Actually we will use that example again to illustrate things in this section.

So now start setting up the language and equations: we write vector \mathbf{x} to give the population sizes. The ‘jump’ is now also a vector (\mathbf{r}) and if an event leads to only one population changing size, then the vector \mathbf{r} will be zeros except for one entry. We allow it to be more general though: for example our ‘populations’ might be stages of the same species e.g. caterpillars and butterflies, in which case we could reasonably have $\mathbf{r} = (-1, +1)$ for the event of one turning into the other. Anyway, there will only be a small (and certainly finite) set of \mathbf{r} to consider.

As before, we let $W(\mathbf{x}, \mathbf{r})$ be the rate at which we make a jump from \mathbf{x} of size \mathbf{r} , i.e. ending up at $\mathbf{x} + \mathbf{r}$.

Then our general master equation⁴⁶ is given by

$$\frac{\partial}{\partial t} P(\mathbf{x}, t) = \sum_{\mathbf{r}} [W(\mathbf{x} - \mathbf{r}, \mathbf{r}) P(\mathbf{x} - \mathbf{r}, t) - W(\mathbf{x}, \mathbf{r}) P(\mathbf{x}, t)]$$

and again the square bracket is expanded using Taylor series, but now for a function of a vector. Here is the relevant expansion for a general F , expanding around \mathbf{x} and hoping \mathbf{r} is small (at least when compared with \mathbf{x}):

$$F(\mathbf{x} - \mathbf{r}) = F(\mathbf{x}) - r_i \frac{\partial F}{\partial x_i} + \frac{1}{2} r_i r_j \frac{\partial^2 F}{\partial x_i \partial x_j} + \dots$$

⁴⁶for 1D we were a bit more careful about using N for when we were thinking about population size as discrete and x when continuous. By now, we know where this is going so just launch straight into \mathbf{x} here

where we are using summation convention. We could also write this in vectors using divergences, but this more explicit form seems like a sensible idea while we are working through the detail below.

Hence our master equation becomes, approximating to the first two terms only (so correct up to r^2 in effect):

$$\frac{\partial P}{\partial t} = \sum_{\mathbf{r}} -r_i \frac{\partial}{\partial x_i} (W P) + \frac{1}{2} \sum_{\mathbf{r}} r_i r_j \frac{\partial^2}{\partial x_i \partial x_j} (W P)$$

where we are still using summation convention. Now reorder sums and integrals to bunch up the sums into $\underline{\mathbf{A}}$ and $\underline{\mathbf{B}}$ terms:

$$\frac{\partial P}{\partial t} = - \frac{\partial}{\partial x_i} \underbrace{\left(\sum_{\mathbf{r}} r_i W \right)}_{A_i(\mathbf{x})} P + \frac{1}{2} \frac{\partial^2}{\partial x_i \partial x_j} \underbrace{\left(\sum_{\mathbf{r}} r_i r_j W \right)}_{B_{ij}(\mathbf{x})} P$$

where we will see that A_i give components of the vector of average movement. The matrix $\underline{\mathbf{B}}$ is trickier, but we can already see that it is a symmetric matrix (same if we swap i and j). So in summary we have our multidimensional FPE:

$$\boxed{\frac{\partial P}{\partial t} = - \frac{\partial}{\partial x_i} (A_i P) + \frac{1}{2} \frac{\partial^2}{\partial x_i \partial x_j} (B_{ij} P) \quad \text{Fokker-Planck Equation (FPE)}}$$

Exercise 35: Find $\underline{\mathbf{A}}$ and $\underline{\mathbf{B}}$ for wildebeest and flies.

Much as we did for one dimension, we can use the FPE to find time evolution of any function of \mathbf{x} , but here we will explicitly look at components of first and second moments.

$$\begin{aligned} \frac{d}{dt} \langle x_m \rangle &= \frac{d}{dt} \int x_m P(\mathbf{x}, t) dx_1 dx_2 \dots dx_n \\ &= \int x_m \frac{\partial P}{\partial t} dV \quad \text{now sub using FPE} \\ &= \int -x_m \frac{\partial}{\partial x_i} (A_i P) + \frac{1}{2} x_m \frac{\partial^2}{\partial x_i \partial x_j} (B_{ij} P) dV \end{aligned}$$

and again we use parts (or divergence theorem if thinking in terms of vectors) to liberate the P from the derivatives. Again, assume that P and its derivatives go to zero at the boundaries:

$$\begin{aligned}\frac{d}{dt}\langle x_m \rangle &= \int + (A_i P) \underbrace{\frac{\partial}{\partial x_i} x_m}_{=\delta_{im}} + \frac{1}{2} (B_{ij} P) \underbrace{\frac{\partial^2}{\partial x_i \partial x_j} x_m}_{=0} dV \\ &= \int (A_i P) \delta_{im} dV = \int A_m P dV = \langle A_m \rangle\end{aligned}$$

Hence in vectors we could write:

$$\langle \mathbf{x} \rangle' = \langle \mathbf{A} \rangle$$

and hence we can see that \mathbf{A} represents the ‘mean drift’, i.e. the time derivative of the mean population sizes. Now continue with same approach for $\langle x_m x_n \rangle$, but we can reuse the working above as it is almost the same initially, and jump in after the integration by parts:

$$\begin{aligned}\langle x_m x_n \rangle' &= \int + (A_i P) \underbrace{\frac{\partial}{\partial x_i} x_m x_n}_{=\delta_{im} x_n + \delta_{in} x_m} + \frac{1}{2} (B_{ij} P) \underbrace{\frac{\partial^2}{\partial x_i \partial x_j} x_m x_n}_{=\delta_{im} \delta_{jn} + \delta_{in} \delta_{jm}} dV \\ &= \int (A_m x_n + A_n x_m) P + \frac{1}{2} (B_{mn} + B_{nm}) P dV \\ &= \int (A_m x_n + A_n x_m) P + B_{mn} P dV\end{aligned}$$

using that $B_{mn} = B_{nm}$. Hence

$$\langle x_m x_n \rangle' = \langle A_m x_n \rangle + \langle A_n x_m \rangle + \langle B_{mn} \rangle.$$

These second moments can be used to consider the covariances

$$C_{mn} = \text{cov}(x_m, x_n) = \langle x_m x_n \rangle - \langle x_m \rangle \langle x_n \rangle,$$


so $C_{ii} = \text{var}(x_i)$, and $\underline{\underline{\mathbf{C}}}$ is a symmetric matrix. Using the above:

$$\begin{aligned}C'_{mn} &= \langle x_m x_n \rangle' - \langle x_m \rangle' \langle x_n \rangle - \langle x_m \rangle \langle x_n \rangle' \\ &= \langle A_m x_n \rangle + \langle A_n x_m \rangle + \langle B_{mn} \rangle - \langle A_m \rangle \langle x_n \rangle - \langle x_m \rangle \langle A_n \rangle \\ &= \langle A_m x_n \rangle - \langle A_m \rangle \langle x_n \rangle + \langle A_n x_m \rangle - \langle x_m \rangle \langle A_n \rangle + \langle B_{mn} \rangle \\ C'_{mn} &= \text{cov}(A_m, x_n) + \text{cov}(A_n, x_m) + \langle B_{mn} \rangle\end{aligned}$$


and all this holds for general \mathbf{A} and $\underline{\underline{\mathbf{B}}}$.

Finding behaviour at steady state:

Here is another run of the wildebeest and flies, each plotted against time:



a different way to visualise this would be plotting flies against wildebeest, i.e. exactly the two-dimensional space that the system is moving around in:



and as the tracks go over themselves repeatedly, we can add the blue dots at fixed time intervals⁴⁷. And by now we are comfortable with thinking of this as steady state: the system is not stuck at fixed \mathbf{x} but rather the probability of being a in a given state has settled.

We can see that the steady state mean (marked in black) is in the middle of this random splurge, but we can also see that the spread around it has some interesting shape: it is spread out diagonally. This corresponds to saying the wildebeest and flies covary positively.

We would like to find this steady state covariance (and variances) in general (not just wildebeest and flies model). We have an equation above for C' , but it depends on $A_m(\mathbf{x})$ and $B_{mn}(\mathbf{x})$ which in turn depend on \mathbf{x} .

Here's the final idea that we need to do this. Approximate near steady state (summation convention below):

- Approximate A_m as being linear in \mathbf{x} :

$$A_m(\mathbf{x}) = \lambda_m + a_{mk} x_k$$

- Approximate B_{mn} as being constant in \mathbf{x} :

$$B_{mn}(\mathbf{x}) = b_{mn}$$

and in both cases this is usually done by taking approximation (Taylor series) near steady state. All of λ_m , a_{mk} and b_{mn} now are constants. Actually $\underline{\underline{a}}$ can be thought of as the Jacobian near the steady state mean.

Then (still with summation convention):

$$\text{cov}(A_m, x_n) = a_{mk} \text{cov}(x_k, x_n) = a_{mk} C_{kn}$$

$$\langle B_{mn} \rangle = b_{mn}$$

and hence:

$$C'_{mn} = a_{mk} C_{kn} + a_{nk} C_{km} + b_{mn}.$$

This looks nice in matrix notation:

$$\underline{\underline{C'}} = \underline{\underline{a}} \underline{\underline{C}} + (\underline{\underline{a}} \underline{\underline{C}})^T + \underline{\underline{b}}$$

⁴⁷should be careful to use set time period, not fixed number of steps, or will not be sampling uniformly over time

and as $\underline{\underline{C}}$ is symmetric, $(\underline{\underline{a}}\underline{\underline{C}})^T = \underline{\underline{C}}\underline{\underline{a}}^T$. So at steady state, the covariace matrix $\underline{\underline{C}}$ must satisfy:

$$\underline{\underline{a}}\underline{\underline{C}} + \underline{\underline{C}}\underline{\underline{a}}^T + \underline{\underline{b}} = 0 \quad \text{Lyapunov equation}$$

or in components (with summation convention):

$$a_{ik}C_{kj} + a_{jk}C_{ki} + b_{ij} = 0.$$

Exercise 36: Continue with the wildebeest and flies: approximate around steady state, i.e.:


$$x_1 = \frac{\lambda_1}{\beta_1}, \quad x_2 = \frac{\lambda_1\lambda_2}{\beta_1\beta_2}$$

(i.e. where $\mathbf{A} = 0$). Find $\underline{\underline{a}}$, $\underline{\underline{b}}$. By taking components of the Lyapunov equation show that near steady state:

$$C_{11} = \frac{\lambda_1}{\beta_1}, \quad C_{12} = \frac{\lambda_1\lambda_2}{\beta_1(\beta_1 + \beta_2)}, \quad C_{22} = \frac{\lambda_1\lambda_2}{\beta_1\beta_2} \left(1 + \frac{\lambda_2}{\beta_1 + \beta_2}\right).$$

These give the variance of x_1 , covariance between x_1 and x_2 , and the variance of x_2 respectively.

And just for completeness with the example we have been following, $\underline{\underline{C}}$ can be used to find the confidence ellipse (or error ellipse), and shown here for equivalent of one standard deviation, and also the 95% CI. And this looks pretty good at capturing the spread and directionality of fluctuations from the mean at steady state.




Steady-state vs quasi steady-state

(Treat as starred.) For one small logical jump we can have a big extension to our covariance result. We already know that if zero (for all species) is fixed (no way to jump out from that state), then the only true steady state for the full system is total extinction in the long term. However, if we have a solution \mathbf{x} for where the mean is constant in time $\mathbf{A}(\mathbf{x}) = 0$ and it is well away from zero, then we call a spread around that as a 'quasi steady-state'. It is not truly a steady state, but it might look like one for a very very long time, so our above results hold well.

We could in principle be a little more formal about it either by thinking about eigenvectors of the Markov transition matrix (the one with second largest eigenvalue is the quasi-steady-state: the largest corresponds to total extinction). Or we could write about large t conditioning on non-extinction and call that the quasi steady state, but no need to go further for this course: just know that it is often OK to use these results even for systems where extinction is possible.


Example

To illustrate what we mean by quasi steady-state, we can use a one-dimensional system as an example for clarity. Consider again the simple import and death model (example 2). Set $\lambda = 10$, $\beta = 1$ for these plots. So we know that the steady state solution has mean and variance 10. Here's an output for large time for the full model in blue, and the normal distribution with the mean and variance from the FPE at steady state in black:




This matches fairly well.


If we now modify the import and death model so there is NO import if $n = 0$, then we in principle know that if we wait long enough, then $P_0 = 1$ and $P_n = 0$ for $n \neq 0$. However, look at what happens for modest time, here $t = 10$ (modified import and death model in red):



It just looks as if there is no modification, hence the steady-state approximation from before is pretty good for early times for this modified system. Indeed if we increase time, here to $t = 100$:



We can see the shape is right, except some of the probability density has 'leaked' out to extinction at $N = 0$. And go further to $t = 1000$:



It still has basically right shape, even though more area has been lost to P_0 . The rest of the curve still keeps its mean and variance.

The FPE for both the original model and this modified model will be almost the same (just A and B will differ at $x = 0$), and the approximation for the variance near steady state from the Lyapunov equation will be the same (from approximating A as linear in x and B as constant near $x = 10$). Unsurprisingly, we've lost the information about extinction being possible. However, we do have the right mean and variance for some large range of times, or even later times if we condition on non-extinction: call this quasi steady-state.

3 Systems with spatial structure

3.0 Preliminaries

Life is spatial. Whether we are thinking about the dynamics within a single living cell, or the interaction of microbes, or pandemics which encompass the globe, the relative physical position of things matters. Back in section 1, we ignored space. One can think of this as an implicit assumption that things are ‘well mixed’, i.e. things may be spatial, but our populations or other biological quantities are so well mixed or connected across space that they might as well be treated as all being in the same place. This is clearly at best only approximately true, but often the complication of adding space does not really tell us anything new (and model parsimony is essential in mathematical biology: only include things if they are useful for the purpose in hand). But, for this final section of the course, we consider approaches to consider the maths of spatial systems. We reach a rather surprising and delightful result at the end with Turing Instabilities.

To start this section, we derive a rather general system for spatial dynamics. Here, use \mathbf{x} as the variable for spatial location (often our space will be \mathbb{R} , \mathbb{R}^2 or \mathbb{R}^3). As always, t is time. And $C(\mathbf{x}, t)$ is the quantity of interest. Here, we will rather generically call it the concentration of ‘stuff’. And for example, ‘stuff’ could be cytokines (chemical signals), bacteria, insects or humans. We will consider a small region in our space: let our (fixed) volume be \mathcal{V} which has surface \mathcal{S} . Now consider what goes in and out, and changes within the volume. The rate of change of total amount of ‘stuff’ in \mathcal{V} is determined by the creation and loss inside \mathcal{V} and also the gain and loss across \mathcal{S} . We can set this up mathematically by writing $F(\mathbf{x}, t)$ (scalar) is the births/deaths at position \mathbf{x} at time t and $\mathbf{J}(\mathbf{x}, t)$ (vector) is the flux of stuff. So then we have

$$\underbrace{\frac{d}{dt} \int_{\mathcal{V}} C(\mathbf{x}, t) dV}_{\text{rate of change of stuff in } \mathcal{V}} = \underbrace{\int_{\mathcal{V}} F(\mathbf{x}, t) dV}_{\text{creation/loss in } \mathcal{V}} - \underbrace{\int_{\mathcal{S}} \mathbf{J}(\mathbf{x}, t) \cdot \mathbf{n} dS}_{\text{flow out through } \mathcal{S}}.$$

Note the minus sign here and the decision to use the flow *out* of \mathcal{S} . This makes sense with the usual practice below of taking the normal of a surface as being outward from the volume. Now writing it more simply and using the divergence theorem and working towards each of the terms being a volume integral over \mathcal{V} .

$$\begin{aligned} \frac{d}{dt} \int_{\mathcal{V}} C dV &= \int_{\mathcal{V}} F dV - \int_{\mathcal{S}} \mathbf{J} \cdot \mathbf{n} dS \\ \int_{\mathcal{V}} \frac{\partial C}{\partial t} dV &= \int_{\mathcal{V}} F dV - \int_{\mathcal{V}} \nabla \cdot \mathbf{J} dV \end{aligned}$$

Strategic rearrange:

$$\int_{\mathcal{V}} \frac{\partial C}{\partial t} dV + \int_{\mathcal{V}} \nabla \cdot \mathbf{J} dV - \int_{\mathcal{V}} F dV = 0$$

$$\int_{\mathcal{V}} \left(\frac{\partial C}{\partial t} + \nabla \cdot \mathbf{J} - F \right) dV = 0$$

and we're there. This holds for any choice of \mathcal{V} , hence the integrand must be zero for all \mathbf{x} :

$$\frac{\partial C}{\partial t} = F - \nabla \cdot \mathbf{J}$$

The term $F = F(\mathbf{x}, t)$ represents all of the non-spatial processes, effectively everything in section 1 could go here. What can we say about this flux $\mathbf{J}(\mathbf{x}, t)$? Essentially there are two types of flux terms:

(i) Advection or active motion: Here, the 'stuff' moves with velocity $\mathbf{u}(\mathbf{x}, t)$. The flux of 'stuff' at a particular location and time is given by:

$$\mathbf{J}(\mathbf{x}, t) = \mathbf{u}(\mathbf{x}, t) C(\mathbf{x}, t)$$

This could be passive motion, for example the background medium is moving, carrying our stuff with it. Or it could be something more active such as bacteria swimming.

(ii) Diffusion: Here our 'stuff' is moving as a sum of small movements of separate particles (or individuals in a population):

$$\mathbf{J}(\mathbf{x}, t) = -D(\mathbf{x}, t) \nabla C(\mathbf{x}, t)$$

In fact, this is Fick's law⁴⁸. Do remember the minus sign here: remember it as net flux will be *down* the population gradient (random movement makes stuff spread out). The $D(\mathbf{x}, t)$ is diffusivity. It is a coefficient that adjusts the rate of movement. Very often this is a constant. It could depend on time or place though, and we will investigate below where it depends on time and place implicitly through actually being a function of C .

Combining these into one form we have

$$\mathbf{J} = \mathbf{u}C - D\nabla C$$

and putting this into our equation from earlier we have our transport equation

$$\boxed{\frac{\partial C}{\partial t} + \nabla \cdot (\mathbf{u}C - D\nabla C) = F(\mathbf{x}, t)}.$$

Various cases of this will be the subject of the next few sections.

⁴⁸actually this can be connected back to the previous section: think of n as a position rather than population size, and then our jumps are small movements in space. Zooming out to larger length scales, the Fokker-Planck equation represented the dynamics. This term is just like our ∇P .

3.1 Diffusion and growth

3.1.1 Linear diffusion in finite domain

For this subsection (and the next) we will take the diffusivity D to be constant. To start with, we are going to study the diffusion alone, so keep $F = 0$ and $\mathbf{u} = 0$. We will also work in one dimension only, so our transport equation reduces to the classic diffusion equation that you will have seen in other courses:

$$\frac{\partial C}{\partial t} = D \frac{\partial^2 C}{\partial x^2}.$$

We consider a finite domain of length L : suppose $x \in [0, L]$. For boundary conditions at the end, for example suppose that C takes some fixed constant values at each end, i.e.

$$C(0, t) = C_0, \quad C(L, t) = C_1.$$

To imagine a particular example, consider a population of bacteria in a (narrow) tube. The bacteria can move about with diffusion. Each end of the tube is connected to a (large) population with constant density (maybe different populations each end, so maybe different densities). There is no birth/death of bacteria in the tube. The ‘large’ populations at the ends mean they stay constant concentration and the bugs diffusing in and out of the tube really don’t affect the outside density. The ‘narrow’ tube means we don’t care about any variations in concentration across the cross-section of the tube: we can treat the system as 1D.

We use the tactic here of first finding the steady solution first, which helps when we think about the full solution (for given initial conditions) as we can just consider the difference from the steady solution (linear system!).

Steady solution Here we just seek a solution which is constant in time, but not necessarily in space. So, this must just satisfy

$$0 = D \frac{\partial^2 C}{\partial x^2}$$

and hence just be linear in x :

$$C(x, t) = C^*(x) = c_0 + (c_1 - c_0)x/L$$

and we’re done.

It is instructive to get a sense of what this makes the flux J at so-called ‘steady state’.

$$J = -D \frac{\partial C}{\partial x} = -\frac{D}{L}(c_1 - c_0)$$

and this is constant across all $x \in [0, L]$. This means (for $c_1 \neq c_0$) there's a net flux throughout. Remember individual bacteria can be going either direction, but the total flow is from the higher concentration to the lower. This is an example of 'conducting' boundary conditions. Demanding that C is constant at the ends but the two ends are different means that there has to be some net flow to hold this in balance.

We still call this 'steady' as the overall density C is constant in time.

General solution We actually need more boundary conditions. We have conditions on the ends $x = 0, L$, and we should also have some initial conditions such as the concentration being known at $t = 0$, i.e. $C(x, 0)$. For general time, we can now just focus on the difference from the steady solution, so we set

$$C(x, t) = C^*(x) + \hat{C}(x, t).$$

By linearity (or subbing and cancelling if we aren't awake enough), we have that \hat{C} solves the same diffusion equation:

$$\frac{\partial \hat{C}}{\partial t} = D \frac{\partial^2 \hat{C}}{\partial x^2}$$

Then work our boundary conditions to be in terms of \hat{C} . The ends condition becomes $\hat{C}(0, t) = \hat{C}(L, t) = 0$, which is nice and simple.

To build a general solution, try first for a separable solution (one which is made of a product of functions in purely x and t). Set

$$\hat{C}(x, y) = F(x)G(t)$$

sub in to the diffusion equation:

$$F(x)G'(t) = DF''(x)G(t)$$

Divide by $F(x)G(t)$ to make each side a function of a single variable:

$$\frac{G'(t)}{G(t)} = D \frac{F''(x)}{F(x)}.$$

So we have a function of t always equals some function of x , so both must be constant and equal. Because we have simple boundary conditions in x , sort that out first. We must have $F(0) = F(L) = 0$ so only possible solutions are

$$F(x) = \sin\left(\frac{n\pi x}{L}\right) \quad \text{for } n = 1, 2, \dots$$

and following through for each n separately, we must have

$$G(t) = e^{-\lambda_n t} \quad \text{with} \quad \lambda_n = D \frac{n^2 \pi^2}{L^2}.$$

From the IB Methods course, we know to build a linear combination of these to make a general solution:

$$\hat{C}(x, t) = \sum_{n=1}^{\infty} a_n e^{-\lambda_n t} \sin\left(\frac{n\pi x}{L}\right)$$

and just add back in the steady solution to get it back in terms of C :

$$C(x, t) = c_0 + (c_1 - c_0) \frac{x}{L} + \sum_{n=1}^{\infty} a_n e^{-\lambda_n t} \sin\left(\frac{n\pi x}{L}\right)$$

where the coefficients a_n are determined by the initial conditions.


And we can see no matter what the initial conditions are, each of the terms in the sum will decay exponentially until only $C^*(x)$ remains. So, things do indeed settle to the steady solutions.

There's a bit more to be seen here though: the exponents λ_n give the decay rates and these are proportional to n^2 and hence get large (and negative) quickly. So suppose we wait until the first exponential is fairly small, e.g. 0.1. Then the next exponential is 10^{-4} and the one after is 10^{-9} and so on, so really (so long as $a_1 \neq 0$) it's all about that first one:

$$C = C^*(x) + a_1 \underbrace{e^{-\frac{D\pi^2 t}{L^2}}}_{=0.1} \sin\left(\frac{\pi x}{L}\right) + a_2 \underbrace{e^{-4\frac{D\pi^2 t}{L^2}}}_{=0.0001} \sin\left(\frac{2\pi x}{L}\right) + \dots$$

The deviation from the steady solution is mainly proportional to $\sin(\pi x/L)$: we call this the 'gravest mode'. This is not a massive surprise: diffusion wipes out sharp patterns quickly, so the more wiggly $\sin(2\pi x/L)$, $\sin(3\pi x/L)$ etc. will decay much faster. We are left with the largest scale pattern that fits the conditions⁴⁹.

⁴⁹all of this is assuming we do not happen to have some special initial condition that leads to $a_1 = 0$. Of course this would be some special case: 'nearly all' initial conditions contain some component of each of the modes, so usually assume all $a_n \neq 0$



This shows $C(x, t)$ for successive timepoints for an example initial condition, starting with red at $t = 0$ (the most wiggly one) and then moving through the rainbow. We can see that the 'wiggles' decay away quickly, leaving effectively a single loop of $\sin(x\pi/L)$ which then gently settles in.

Exercise 37: Change the boundary conditions so that we still have a fixed concentration at one end, but instead we have no flux at the other end, i.e.:

$$\frac{\partial C}{\partial x} = 0 \quad \text{at} \quad x = 0, \quad C = C_1 \quad \text{at} \quad x = L.$$

(This is equivalent to closing off one end of the tube.) What is the gravest mode now?

Add linear growth to finite domain. Finally, still with linear diffusion in a finite domain, we get a very interesting result if we include linear growth. Suppose that our bacteria concentration grows with per capita rate λ (and take $\lambda > 0$). In other words, we have $F = \lambda C$ in the transport equation so

$$\frac{\partial C}{\partial t} = D \frac{\partial^2 C}{\partial x^2} + \lambda C$$

and take $C = 0$ at $x = 0, L$ now, so there's no population outside of the tubes. Bacteria can diffuse out, but none are coming in. If it was just diffusion ($\lambda = 0$), all the bacteria would leak out of the tube, so we'd head to an empty tube. If it was just growth ($D = 0$, $\lambda > 0$), the bacterial population would grow exponentially without bound. So, with both, which one wins out?

We can solve this immediately with one trick: basically factorise out the growth. Set $C(x, t) = e^{\lambda t} \tilde{C}(x, t)$ and sub this in for an equation in \tilde{C} :

$$\lambda e^{\lambda t} \tilde{C} + e^{\lambda t} \frac{\partial \tilde{C}}{\partial t} = D e^{\lambda t} \frac{\partial^2 \tilde{C}}{\partial x^2} + \lambda e^{\lambda t} \tilde{C}$$

And we see this works as things cancel out to leave

$$\frac{\partial \tilde{C}}{\partial t} = D \frac{\partial^2 \tilde{C}}{\partial x^2}$$

which is exactly what we had with zero growth, so we can solve exactly as before.

Exercise 38: Show (for general a_n , i.e. assume none are zero) that the the population will die out for $L < L_c$ and grow for $L > L_c$ for some critical L_c (which you should find!).

If the tube is too short, the bacterial population dies out, even though it has positive growth rate. This is intuitive in some ways: if the tube is short then diffusion wins out as the bacteria keep falling out of the ends faster than they can grow inside. If the tube is long enough, the growth inside will more than compensate for the loss at the ends. (Similarly, if you've done the exercise, you will be able to see that population dies out if D is too large or λ too small, so this all fits.)

This is non-intuitive in other ways: *the stability of the system depends on the geometry* (the length of the tube). This surprising result is a nice precursor to the Turing instabilities that we will see at the end of this chapter.

end of
lecture 17

3.1.2 Linear diffusion in infinite domain

Linear diffusion (constant D) in an infinite domain (real line) has been covered in other earlier courses already, more than once probably, so the results will not be terribly surprising. Here we will use this recap to work through some language and approaches that will be more generally useful (in nonlinear diffusion, next subsection). We also need it as a basis of comparison with what happens when we add the local dynamics (see travelling waves in next section), or more than one variable (spatial instabilities, final section).

We are still using the linear diffusion equation

$$\frac{\partial C}{\partial t} = D \frac{\partial^2 C}{\partial x^2}$$

and now considering $x \in \mathbb{R}$. To start with, we will consider systems where the concentration of stuff is 'localised to the origin'. This means that the initial conditions are that the initial concentration is all somewhere close from origin (or indeed all at origin).

Mathematically we should interpret that as C goes to zero far from origin, and does so sufficiently fast that we may C , $\partial C/\partial x$, and whatever other derivatives as tending to zero for large x (positive or negative).

The total amount of ‘stuff’ (population, chemical or whatever it was) can be found just by integrating:

$$M = \int_{-\infty}^{\infty} C(x, t) dx$$

where we have labelled this total as M . By differentiating with respect to time subbing using the diffusion equation, we can check this is constant:

$$\frac{dM}{dt} = \frac{d}{dt} \int_{-\infty}^{\infty} C(x, t) dx = \int_{-\infty}^{\infty} \frac{\partial C}{\partial t} dx = \int_{-\infty}^{\infty} D \frac{\partial^2 C}{\partial x^2} dx = \left[D \frac{\partial C}{\partial x} \right]_{-\infty}^{\infty} = 0$$

where the last move is using our ‘localised to origin’ condition to say integrand tends to zero at ends.

We will solve our diffusion equation, subject to total stuff being constant, using similarity solutions. There’s two natural ways to do this, and of course they’re actually not fundamentally different, but useful to see both. In both, we will set up a new variable ξ as a *space-like* variable, which means it can be used as a substitute for x (in fact it is proportional to x here).

Approach 1: non-dimensionalise We will seek a solution of the form $C(x, t) = \eta f(\xi)$, where η has no dependence on x . Here the idea is to think about dimensions to set up the only possible ξ and η , and then we’ll have lovely simple differential equations and boundary conditions to solve. But, there’s a bit of work to set it up. We will have dimensions of length (L) and time (T). We also need dimensions for concentration ‘stuff’: just call that C again. We will use square brackets to determine taking the dimension of a variable, e.g. $[x] = L$. We want a non-dimensional space-like variable i.e. $[\xi] = 1$, and we want $[\eta] = C$ to make the overall solution have right dimension.

Here are all of our available quantities to make our ξ and η , and their dimensions:

$$\{ [t] = T, [x] = L, [D] = L^2 T^{-1}, [M] = CL \}$$

The first two are our obvious free variables. The next two are the constants in our PDE and in the integral (which says total stuff is constant). For the dimensions of a diffusion constant, it is probably easiest just to look at the PDE. Same for M and the integral (integrate C wrt space so it has dimensions CL).

We can immediately simplify by noting we don’t want the dimension T for anything, so eliminate⁵⁰ it:

$$\{ [x] = L, [Dt] = L^2, [M] = CL \}$$

⁵⁰This is all basically Gaussian elimination. We could encode the dimensions above into a matrix where columns represent quantities that we can use (t, x, D, M) and rows are the powers of the dimensions (T, L and C respectively here). Then by writing $\xi = t^{u_1} x^{u_2} D^{u_3} M^{u_4}$ and seeking dimensionless ξ

Then working towards ξ we would eliminate C (which just wipes out M) and make it proportional to x , and balance the other to make it dimensionless, hence we are led to

$$\xi = \frac{x}{\sqrt{Dt}}.$$

For η , we want no x -dependence but we do want dimension of C , so we must have

$$\eta = \frac{M}{\sqrt{Dt}}.$$

And so to summarise where we are, we are now looking for a solution to

$$\frac{\partial C}{\partial t} = D \frac{\partial^2 C}{\partial x^2}, \quad \int_{-\infty}^{\infty} C(x, t) dx = M$$

with

$$C(x, t) = \eta f(\xi), \quad \text{with } \xi = \frac{x}{\sqrt{Dt}} \text{ and } \eta = \frac{M}{\sqrt{Dt}}.$$

The next move is to update our PDE and integral to our new variables, and hope something nice happens. There's a danger of algebraic mess while making this change, so it is a very good idea to prepare the first derivatives beforehand in a tidy way, and then pull them in as needed:

$$\frac{\partial \eta}{\partial t} = -\frac{1}{2} \frac{\eta}{t}, \quad \frac{\partial \xi}{\partial t} = -\frac{1}{2} \frac{\xi}{t}, \quad \frac{\partial \eta}{\partial x} = 0, \quad \frac{\partial \xi}{\partial x} = \frac{1}{\sqrt{Dt}} = \frac{\xi}{x}.$$

Then working first with the PDE (diffusion equation):

$$LHS = \frac{\partial C}{\partial t} = \frac{\partial}{\partial t} (\eta f(\xi)) = \frac{\partial \eta}{\partial t} f + \frac{\partial \xi}{\partial t} \eta f' = -\frac{1}{2t} (\eta f + \eta \xi f') = -\frac{\eta}{2t} (\xi f)'$$

where prime denotes differentiation wrt ξ . Note we have tidied as far as we can, and also have one step of integration at the ready there. Meanwhile the RHS is not at all bad:

$$\frac{\partial}{\partial x} = \frac{1}{\sqrt{Dt}} \frac{\partial}{\partial \xi} \quad \text{and} \quad \frac{\partial^2}{\partial x^2} = \frac{\partial}{\partial x} \left(\frac{1}{\sqrt{Dt}} \frac{\partial}{\partial \xi} \right) = \frac{1}{\sqrt{Dt}} \frac{\partial}{\partial x} \frac{\partial}{\partial \xi}$$

using that partial differentiation with respect to x can commute past functions of t (as that it is what is kept fixed), so

$$\frac{\partial^2}{\partial x^2} = \frac{1}{\sqrt{Dt}} \left(\frac{1}{\sqrt{Dt}} \frac{\partial}{\partial \xi} \right) \frac{\partial}{\partial \xi} = \frac{1}{Dt} \frac{\partial^2}{\partial \xi^2}$$

we are solving for the u_i such that

$$\begin{pmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 2 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

and demand that $u_2 = 1$.

and hence we can now prepare the right-hand side:

$$RHS = D \frac{\partial^2 C}{\partial x^2} = D \frac{1}{Dt} \frac{\partial^2}{\partial \xi^2} (\eta f(\xi)) = \frac{\eta}{t} f''$$

Equating the two sides and tidying, our PDE diffusion equation becomes this equation for f . :

$$f'' + \frac{1}{2}(\xi f)' = 0.$$

This is all ready to integrate:

$$f' + \frac{1}{2}\xi f = k$$

where k is the constant of integration, which can be resolved by using the 'localised to origin' condition. Far away from the origin, we have C and its derivatives in space are zero, which must mean we have $f(\xi)$ and derivatives tend to zero far away. Hence LHS tends to zero, so $k = 0$. So now we have

$$f' = -\frac{1}{2}\xi f$$

which we can integrate again to get

$$f = A e^{-\frac{1}{4}\xi^2}$$

with a new constant of integration A . As well as working with the diffusion equation, we should also work with the integral (that says total stuff is constant), and this will resolve A :


$$M = \int_{-\infty}^{\infty} C(x, t) dx = \int_{-\infty}^{\infty} \underbrace{\frac{M}{\sqrt{Dt}}}_{\eta} f(\xi) \underbrace{\sqrt{Dt} d\xi}_{dx}.$$

Hence

$$\int_{-\infty}^{\infty} f d\xi = 1$$

which is rather nice. We can then sort out the constant $A = (4\pi)^{-\frac{1}{2}}$. Putting it all together we have

$$C(x, t) = \frac{M}{\sqrt{4\pi Dt}} e^{-\frac{x^2}{4Dt}}$$



This shows $C(x, t)$ for successive timepoints, starting with red at t very small and then moving through the rainbow. It is a Gaussian where the height drops width grows both proportional to square root time.

Running time backwards towards zero, this will tend to a Dirac delta function (or rather $M \delta(x)$): the area stays constant but C tends to zero everywhere except $x = 0$. In fact, we could use this as a Green's function now to handle any general initial condition. But, that's enough for now.

Approach 2: just sort powers of t Sorting out the dimensions worked nicely in giving us a form which would tidy up. Really though, getting things right with x and t were the only important bits mathematically, the rest just made it nice in M and D . This leads to a second approach to reach the same thing: just sort out the powers of t . Start with a general exponents and then see what is required to make the similarity solution actually work. Set

$$C(x, t) = t^\alpha G(\xi), \quad \text{with } \xi = \frac{x}{t^\beta}$$

and we go ahead with this form, and see what α or β would work. First the diffusion equation, aiming to write this in terms of ξ and t only (no x):

$$LHS = \frac{\partial C}{\partial t} = \alpha t^{\alpha-1} G + t^\alpha \left(-\beta \frac{x}{t^{\beta+1}} \right) G' = \alpha t^{\alpha-1} G - \beta t^{\alpha-1} \xi G' = t^{\alpha-1} \underbrace{(\alpha G - \beta \xi G')}_{\text{fn of } \xi, \text{ no } t}$$

$$RHS = D \frac{\partial^2 C}{\partial x^2} = D t^\alpha t^{-2\beta} G'' = D t^{\alpha-2\beta} G''$$

Equating, we must have the powers of t match up if this is to be a valid solution.

$$\alpha - 1 = \alpha - 2\beta$$

hence $\beta = \frac{1}{2}$.

Next, the integral should also behave correctly with time:

$$M = \int_{-\infty}^{\infty} C dx = \int_{-\infty}^{\infty} t^{\alpha} G(\xi) t^{\beta} d\xi = t^{\alpha+\beta} \underbrace{\int_{-\infty}^{\infty} G d\xi}_{\text{no } t}$$

this must not depend on t , so $\alpha + \beta = 0$. Combining with earlier, $\alpha = -\frac{1}{2}$ and $\beta = \frac{1}{2}$.

Reassuringly, this matches up with the time dependence of our ξ and η from the approach using dimensions. Things will proceed almost as before to resolve G , except there will be more stray factors of D and M lurking around, but it will come out the same as before.

Exercise 39: Try this powers approach for n dimensions. Solve

$$M = \int_{-\infty}^{\infty} C(x, t) dV, \quad \frac{\partial C}{\partial t} = D \nabla^2 C$$

using

$$C(x, t) = t^{\alpha} G(\xi), \quad \text{with } \xi = \frac{|x|}{t^{\beta}}.$$

What α and β make this work?

In addition to thinking about solutions which are ‘localised to the origin’, for some applications in mathematical biology we might be interested in solutions which tend to different constants far from the origin. In principle we could use the solution above as a Green’s function to resolve these, but we can also just solve directly, but now with different boundary conditions:

$$\frac{\partial C}{\partial t} = D \frac{\partial^2 C}{\partial x^2} \quad \text{with} \quad \begin{cases} C \rightarrow +1 \text{ as } x \rightarrow +\infty \\ C \rightarrow -1 \text{ as } x \rightarrow -\infty \end{cases}$$

and again assume that derivatives of C with respect to x tend to zero as fast as we need below. Note we cannot talk about integral of C over the range (it’s tending to a non-zero constant far away, so we certainly have a problem with the integral). Here, we will take approach 2 from above, and just sort out powers of t . Indeed the condition above to make the PDE depend on time in a consistent manner is exactly the same, so we already have $\beta = \frac{1}{2}$, so we are already at

$$C(x, t) = t^{\alpha} G(\xi), \quad \text{with } \xi = \frac{x}{t^{\frac{1}{2}}}.$$

Demanding that C tends to a non-zero constant far away tells us that we must now have $\alpha = 0$ (and now our G must tend to -1 as ξ tends to $-\infty$, and $+1$ as ξ tends to ∞).

Reusing our working from above, our PDE must have:

$$LHS = \frac{\partial C}{\partial t} = t^{\alpha-1}(\alpha G - \beta \xi G') = t^{-1}(-\frac{1}{2}\xi G')$$

$$RHS = D \frac{\partial^2 C}{\partial x^2} = D t^{\alpha-2\beta} G'' = D t^{-1} G''$$

Hence the diffusion equation has become

$$-\frac{1}{2}\xi G' = D G''$$

Which we can go ahead and integrate once:

$$G' = A e^{-\frac{1}{4D}\xi^2}$$

for some constant A . Integrating again:

$$G = B + A \int_0^\xi e^{-\frac{1}{4D}\hat{\xi}^2} d\hat{\xi}$$

And if we use the *error function*⁵¹ we can write down the solution in fairly nice form:

$$G = B + A \frac{\sqrt{\pi}}{2} \sqrt{4D} \operatorname{Erf} \left(\frac{\xi}{\sqrt{4D}} \right)$$

We also want G tending to ± 1 as x tends to $\pm\infty$, so that resolves A and B ($B = 0$ and A is whatever is needed to make factor 1 at the front of Erf). Moving back from our similarity solution in ξ to a final result in x :


$$C(x, t) = \operatorname{Erf} \left(\frac{x}{\sqrt{4Dt}} \right)$$

⁵¹Error function is defined as

$$\operatorname{Erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt.$$

You might recognise this from probability/stats courses. It's an odd function which goes from -1 to $+1$ as x goes from $-\infty$ to $+\infty$ and zero at zero. For small x we can approximate $\operatorname{Erf}(x)$:

$$\operatorname{Erf}(x) \approx \frac{2}{\sqrt{\pi}} x \quad \text{for } |x| \ll 1$$




This shows $C(x, t)$ for successive timepoints, starting with red at t very small and then moving through the rainbow.

Example: FRAP A simple application of this solution is in analysing FRAP: fluorescence recovery after photobleaching. This is a technique to estimate the diffusion coefficient D . The details don't matter for the course, but in case you are interested: suppose there are some proteins on some surface (like membrane proteins on the surface of a cell) and these are tagged with fluorescence. This fluorescence can be *bleached* by using a laser. Then there is some small patch which no longer glows. As the membrane proteins are moving around, and the bleached patch was small, eventually the fluorescence is recovered after some time. By observing how fluorescence recovers, these experiments can give estimates for D .

Set it up mathematically: let $C(x, t)$ be the fluorescence, have it starting at one everywhere and suppose bleaching takes it to zero. so the patch is width $2w$ (from $-w$ to $+w$). We want something that at time zero is one for $|x| > w$ and zero inside that window $|x| < w$. We are using the same diffusion equation as above, and the trick is to note that it is linear we can superpose solutions. We want to step down from 1 to 0 at $x = -w$ and back up from 0 to 1 at $x = w$. So, have error functions centred on $x = \pm w$ and appropriate coefficients:

$$C(x, t) = 1 - \frac{1}{2} \operatorname{Erf} \left(\frac{x + w}{\sqrt{4Dt}} \right) + \frac{1}{2} \operatorname{Erf} \left(\frac{x - w}{\sqrt{4Dt}} \right)$$



This shows $C(x, t)$ for successive timepoints, starting with red at t very small and then moving through the rainbow.

We know what width w is, and we can observe fluorescence as a function of time, so we should be able to estimate D from this.

Exercise 40: Suppose that fluorescence is half recovered after about 1 second. The window of bleaching is $w = 1\mu m$ (one micron, which is $= 10^{-6}$ meters). Estimate D . (Hint, use $x = 0$).

end of
lecture 18

Add simple growth to linear diffusion in infinite domain. One final note before moving on to more funky forms of diffusion: what happens when we introduce growth to diffusion in an infinite domain, when things are localised to the origin? In other words, use

$$\frac{\partial C}{\partial t} = D \frac{\partial^2 C}{\partial x^2} + \lambda C$$

with initial condition

$$M = \int_{-\infty}^{\infty} C(0, t) dx$$

where this is not going to be constant any more. Recall our trick from linear diffusion in a finite domain to basically factor out the growth by a strategic change of variable. Set $C(x, t) = e^{\lambda t} \tilde{C}(x, t)$ and then, as before, we will have

$$\frac{\partial \tilde{C}}{\partial t} = D \frac{\partial^2 \tilde{C}}{\partial x^2}.$$

Which we can solve immediately as before, then return our factor of $e^{\lambda t}$ to see that

$$C(x, t) = \frac{M}{\sqrt{4\pi Dt}} e^{\lambda t - \frac{x^2}{4Dt}}$$

Hence this behaves as expected. If $\lambda \leq 0$ then for each value of x , $C(x, t) \rightarrow 0$ as $t \rightarrow \infty$. If $\lambda > 0$ then for each value of x , $C(x, t) \rightarrow \infty$ as $t \rightarrow \infty$. In other words, the growth λ does exactly what you expect it to do. It was different in a finite domain as stuff diffusing out of the sides could prevent total growth, even for $\lambda > 0$. Here, with an infinite domain, diffusion cannot stop the growth.

3.1.3 Nonlinear diffusion

Up until here, the diffusion coefficient D has been a constant, though we did set up the original system to allow for it being a function of space and time. In practice the diffusion rate D could depend on local concentration C . For example bacteria may 'sense' the presence of other bacteria and either move more to avoid competition. Or, more interestingly, they could choose to move less in order to cluster together (and move more when local concentration is low). However, their movement will still be random and undirected, just that there is more or less of it.

Consider the example of $D = kC$ where k is some positive constant (case where more bacteria means more movement). We will still work in one dimension only, and consider solutions which are localised to the origin again. Our transport equation is given by

$$\frac{\partial C}{\partial t} = \frac{\partial}{\partial x} \left(D \frac{\partial C}{\partial x} \right) = \frac{\partial}{\partial x} \left(kC \frac{\partial C}{\partial x} \right) = k \frac{\partial}{\partial x} \left(C \frac{\partial C}{\partial x} \right).$$

We still have constant total:

$$\int_{-\infty}^{\infty} C(x, t) dx = M.$$

Now we follow the approach of non-dimensionalising. As before a solution of the form $C(x, t) = \eta F(\xi)$, where η has no dependence on x . Again, we want a non-dimensional space-like variable i.e. $[\xi] = 1$, and we want $[\eta] = C$. Rather than having the diffusion coefficient in our set of available quantities, we now just have the constant k and we can deduce its dimensions from looking at the transport equation again, or by using $D = kC$ and recalling the dimensions of D from before. So, now here are all of our available quantities to make our ξ and η , and their dimensions:

$$\{ [t] = T, [x] = L, [k] = L^2 T^{-1} C^{-1}, [M] = CL \}$$

Again the natural first move is to eliminate T as we don't need it for ξ or η :

$$\{ [x] = L, [kt] = L^2 C^{-1}, [M] = CL \}$$

Working towards ξ , eliminate C :

$$\{ [x] = L, [Mkt] = L^3 \}$$

then make ξ proportional to x and eliminate L :

$$\xi = \frac{x}{(Mkt)^{\frac{1}{3}}}.$$

Then for η , return to

$$\{ [x] = L, [kt] = L^2 C^{-1}, [M] = CL \}$$

and we don't use x , so just

$$\{ [kt] = L^2 C^{-1}, [M] = CL \}$$

eliminate L :

$$\{ [M^2/kt] = C^3 \}$$

and make $[\eta] = C$:

$$\eta = [M^2/kt]^{\frac{1}{3}} = \frac{M}{(Mkt)^{\frac{1}{3}}}$$

where that last rearrangement is the natural one to bring out the shared factor in ξ and η .

Exercise 41: Find suitable ξ, η for $D = kC^p$ (where p is some constant)

This time we will update the integral first:

$$M = \int_{-\infty}^{\infty} C(x, t) dx = \int_{-\infty}^{\infty} \eta F(\xi) dx = \int_{-\infty}^{\infty} \underbrace{\frac{M}{(Mkt)^{\frac{1}{3}}}}_{\eta} F(\xi) \underbrace{(Mkt)^{\frac{1}{3}} dx}_{dx} = M \int_{-\infty}^{\infty} F(\xi) d\xi,$$

hence

$$\int_{-\infty}^{\infty} F d\xi = 1.$$

For the PDE, and it is a good idea to prepare all the necessary partial derivatives to change variable:

$$\frac{\partial \xi}{\partial t} = -\frac{1}{3} \frac{\xi}{t}, \quad \frac{\partial \eta}{\partial t} = -\frac{1}{3} \frac{\eta}{t}, \quad \frac{\partial \xi}{\partial x} = \frac{1}{(Mkt)^{\frac{1}{3}}}, \quad \frac{\partial \eta}{\partial x} = 0.$$

And a little more preparation this time, using chain rule:

$$\frac{\partial}{\partial t} = -\frac{1}{3} \frac{\xi}{t} \frac{\partial}{\partial \xi} - \frac{1}{3} \frac{\eta}{t} \frac{\partial}{\partial \eta} = -\frac{1}{3t} \left(\xi \frac{\partial}{\partial \xi} + \eta \frac{\partial}{\partial \eta} \right)$$

and

$$\frac{\partial}{\partial x} = \frac{1}{(Mkt)^{\frac{1}{3}}} \frac{\partial}{\partial \xi}$$

(And don't worry too much about the prefactors involving t on the right hand side, hope for them to cancel somewhere, and remember the $\partial/\partial x$ can commute past them.)

So LHS of the PDE becomes

$$LHS = \frac{\partial C}{\partial t} = -\frac{1}{3t} \left(\xi \frac{\partial}{\partial \xi} + \eta \frac{\partial}{\partial \eta} \right) \eta F(\xi) = -\frac{1}{3t} (\eta \xi F' + \eta F) = -\frac{\eta}{3t} (\xi F)'$$

And with RHS we have to be careful about multiple occurrences of C now:

$$RHS = k \frac{\partial}{\partial x} \left(C \frac{\partial C}{\partial x} \right) = \frac{k \eta^2}{(Mkt)^{\frac{2}{3}}} \frac{\partial}{\partial \xi} \left(F \frac{\partial F}{\partial \xi} \right) = \frac{k M \eta}{(Mkt)} (FF')' = \frac{\eta}{t} (FF')'$$

Equate, tidy, integrate (which is not hard as we tidied things into derivatives as we went):

$$FF' + \frac{1}{3} \xi F = \text{some constant}$$

and as before, we see this constant must be zero by noting that the LHS goes to zero far from the origin. Hence

$$F(F' + \frac{1}{3} \xi) = 0$$

Be very careful not to just divide out F at this point: we need it still to get the full solution. Factorising, either $F = 0$ or $(F' + \frac{1}{3} \xi) = 0$. The first one just zero everywhere. The second one gives a parabola $F = A - \frac{1}{6} \xi^2$. However, neither of these seems to work with our boundary conditions and integral $\int F = 1$. If we use $F = 0$, then we certainly have things tending to zero far from the origin, but the integral will never be satisfied. If we use the other one, then we get something non-trivial at least, but neither the integral or the condition far from origin have any hope of being satisfied.

The answer here is to splice together solutions, and this works⁵². We can cut between them only where $F = 0$, so that $F(F' + \frac{1}{3} \xi) = 0$ is (almost!) satisfied. Hence, here is the only hopeful form:

$$F(\xi) = \begin{cases} A - \frac{1}{6} \xi^2 & \text{for } |\xi| < \xi_0 \text{ where } \xi_0 = \sqrt{6A} \\ 0 & \text{otherwise.} \end{cases}$$

⁵²Your inner pure mathmo will be growling here... strictly this doesn't quite work as F' is not defined at the splice point. Your inner applied mathmo will be saying 'wait, were we always allowed to do this and I never knew?'. Both can be resolved by noting the very special form here with the two factors: $F(F' + \frac{1}{3} \xi) = 0$. Strictly we should probably say something like we smooth things out very close to the corner where we splice between solutions (to make F' sensible), and that solution will evolve close to the 'solution' presented here. But, for math bio, it's fine to go with these spliced solutions here.

This satisfies the PDE. The localised boundary condition is certainly satisfied ($F \rightarrow 0$ as $\xi \rightarrow \pm\infty$). We still have a free choice of A and this is resolved by demanding the integral condition ($\int F = 1$):

$$1 = \int_{-\infty}^{\infty} F(\xi) d\xi = \int_{-\sqrt{6A}}^{\sqrt{6A}} A - \frac{1}{6}\xi^2 d\xi = \sqrt{\frac{32}{3}} A^{\frac{3}{2}}$$

hence

$$A = \left(\frac{3}{32}\right)^{\frac{2}{3}} \quad \text{hence} \quad \xi_0 = \sqrt{6A} = \left(\frac{9}{2}\right)^{\frac{1}{3}}$$

and we are done. Just write the final solution back in terms of the original variables:


$$C(x, t) = \eta F(\xi) = \begin{cases} \frac{M}{(Mkt)^{\frac{1}{3}}} \left(\frac{1}{6} \left(\frac{9}{2}\right)^{\frac{2}{3}} - \frac{1}{6} \frac{x^2}{(Mkt)^{\frac{2}{3}}} \right) & \text{for } |x| < \left(\frac{9}{2}Mkt\right)^{\frac{1}{3}} \\ 0 & \text{otherwise.} \end{cases}$$

This can be tidied up somewhat by setting $x_0 = \left(\frac{9}{2}Mkt\right)^{\frac{1}{3}}$. Then,

$$C(x, t) = \begin{cases} \frac{1}{6kt} (x_0^2 - x^2) & \text{for } |x| < x_0 \\ 0 & \text{otherwise.} \end{cases}$$

Exercise 42: Starting from $D = kC$, get to here without notes.

So this is just a chunk of a parabola. The (half) width is x_0 which is proportional to $t^{\frac{1}{3}}$ and the maximum height (at $x = 0$) is given by $x_0^2/(6kt)$ and so is proportional to $t^{-\frac{1}{3}}$. This is all very reassuring: the area stays constant. As time increases, the concentration spreads outwards, but holds the same shape (as hoped for, given we were trying for a similarity solution).



This shows $C(x, t)$ for successive timepoints, starting with red at t very small and then moving through the rainbow.

The sharp (or nearly sharp) edges might seem odd to start with, but recall the original model: $D = kC$. So, our substance/population essentially cannot diffuse at very low concentration, so it piles up at the edge of its reaches, until there are enough to start moving, explaining why this is such a weird boundary. This corner is a consequence of our form of diffusion, together with condition that things are localised to the origin.

end of
lecture 19

3.2 Travelling waves in reaction-diffusion systems

3.2.0 General F

From now on we focus on constant diffusivity D and work in one-dimension, unless otherwise stated.

Recall the general transport equation had a term $F(x, t)$ to represent the ‘creation/loss’ of our stuff. So far we have looked at the behaviour of systems with diffusion, but the F has been restricted to just be linear growth of stuff i.e. $F(x, t) = \lambda C$. In this section, we bring back the possibility of a more general $F(C)$ (so local dynamics depends on concentration of stuff, not time or space on its own), and we start to explore more the interaction between local dynamics (F) and spatial dynamics (diffusion terms).

$$\frac{\partial C}{\partial t} = D \frac{\partial^2 C}{\partial x^2} + F(C)$$

Another way to describe this is as a ‘reaction-diffusion’ equation: F is the reaction (change from non-spatial processes, or local processes) and obviously $D \frac{\partial^2 C}{\partial x^2}$ is the diffusion term.

This is, in general, a non-linear PDE in C (unless F happens to be proportional to C). First, see what happens if we revert to our usual tactic of solve and perturb. Think about a purely spatially homogeneous dynamics first. Spatially homogeneous means $C(x, t) = C(t)$. Then subbing this in, our PDE becomes

$$\frac{dC}{dt} = F(C)$$

which is exactly the sort of thing we studied near the start of the course. Let us suppose we find C_0 is a fixed point of this system, i.e. $F(C_0) = 0$. Then consider a small spatial perturbation $C(x, t) = C_0 + \epsilon(x, t)$. Subbing this in, we get a PDE for ϵ :

$$\frac{\partial \epsilon}{\partial t} = D \frac{\partial^2 \epsilon}{\partial x^2} + F(C_0 + \epsilon(x, t)) ,$$

and as we take ϵ to be small, we can Taylor Series expand:

$$F(C_0 + \epsilon(x, t)) = \underbrace{F(C_0)}_{=0} + \underbrace{\epsilon(x, t) F'(C_0)}_{\lambda} + \underbrace{\mathcal{O}(\epsilon^2)}_{\text{neglect}}$$

and $F(C_0) = 0$ as we were working around a fixed point. $F'(C_0)$ is some number, call it λ . Ignore higher order terms. Then we have

$$\frac{\partial \epsilon}{\partial t} = D \frac{\partial^2 \epsilon}{\partial x^2} + \lambda \epsilon$$

which is exactly what we had near the end of subsection 3.1.2 when we considered linear growth. So we already know the outcome. If $F'(C_0) < 0$ then perturbations about $C = C_0$ will decay and if $F'(C_0) > 0$ then perturbations will grow. Note that these are also the same conditions for the non-spatial system, just classic section 1 style results.

What we ask next in this section is this: OK, we’ve done the linearised dynamics near homogeneous solution, what happens when the solution is further away from the equilibrium and the linear approximation isn’t valid? Can we get at the full dynamics in any way? This turns out to be rather tricky in general, but we will see some possible approaches below.

3.2.1 Fisher’s equation

We start by studying a classic equation, motivated by a classic model (first published in 1937 I think) which describes how a beneficial mutation can spread in a population. Imagine a spatial population and somewhere a mutation arises which turns out to be good for those individuals who carry that mutation. In population genetics, ‘good’ mean they are fitter, which means they have more offspring.

First we will build the model in a non-spatial setting. Let $N(t)$ be the total population and let $p(t)$ be the proportion of the population which currently has the ‘good’ mutation ($N > 0$, $0 \leq p \leq 1$). Suppose that at time t the non-mutant population is growing at rate λ and the mutant population at rate $\lambda + s$. So $s > 0$ is the boost to the growth rate conferred by the mutation. Once the mutation has first appeared, there is no further mutations happening, i.e. mutants give birth to mutants and non-mutants give birth to non-mutants only.

Now build a picture of how things change in some short time δt (a bit like an approach you may have seen to derive the rocket equation in *la Dynamics & Relativity*). Just for clarity in the table and below, write $N = N(t)$ and $p = p(t)$.

time	Number of non-mutants (growth rate λ)	Number of mutants (growth rate $\lambda + s$)
t	$N(1 - p)$	Np
$t + \delta t$	$N(1 - p)(1 + \lambda \delta t) + \mathcal{O}(\delta t^2)$	$Np(1 + (\lambda + s)\delta t) + \mathcal{O}(\delta t^2)$

To find $N(t + \delta t)$, add the number of non-mutants and mutants at time $t + \delta t$:

$$\begin{aligned} N(t + \delta t) &= N(1 - p)(1 + \lambda \delta t) + Np(1 + (\lambda + s)\delta t) + \mathcal{O}(\delta t^2) \\ &= N(1 + \lambda \delta t + sp \delta t) + \mathcal{O}(\delta t^2) \end{aligned}$$

And similarly to find $p(t + \delta t)$, divide number of mutants at time $t + \delta t$ by $N(t + \delta t)$:

$$\begin{aligned} p(t + \delta t) &= \frac{\text{mutants}(t + \delta t)}{N(t + \delta t)} = \frac{N(1 - p)(1 + \lambda \delta t)}{N(1 + \lambda \delta t + sp \delta t)} + \mathcal{O}(\delta t^2) \\ &= p \left[1 + \frac{(s - sp)\delta t}{1 + (\lambda + sp)\delta t} \right] + \mathcal{O}(\delta t^2) \\ &= p[1 + s(1 - p)\delta t] + \mathcal{O}(\delta t^2) \end{aligned}$$


Hence (now writing $p(t)$ properly) we have

$$p(t + \delta t) = p(t) + sp(t)(1 - p(t))\delta t + \mathcal{O}(\delta t^2).$$

So as usual, rearrange, take limit $t \rightarrow 0$ to get a differential equation for p :

$$\frac{dp}{dt} = sp(1 - p)$$

and we recognise this as the logistic equation, again. This is maybe not a massive surprise: a beneficial mutation, once introduced, should take over. In other words, p should grow from small but positive up towards 1.




Solution to the logistic equation, starting with $0 < p(0) \ll 1$ and $s > 0$. Once introduced, the beneficial mutation will take over.

Now we extend this system to be a spatial model, just by adding linear diffusion. (Or more formally I suppose, use the transport equation where we put $F = sp(1-p)$ for the local process. We then reach Fisher's equation:

$$\frac{\partial p}{\partial t} = \frac{\partial^2 p}{\partial x^2} + p(1-p) \quad \text{Fisher's equation}$$

Suppose we start with the beneficial mutation absent almost everywhere, so we have $p(x, t) = 0$ except for some small range of x , where we perturb p from zero. We find numerically that the solution looks like it is heading to a steady wave solution:



Initial condition in red (small Gaussian perturbation away from zero) and then moving through the rainbow in time. The solution seems to quickly settle to a steady wave progressing away from the initial perturbation in each direction.

It does not seem to much matter what initial conditions are used, so long as $p \in [0, 1]$ and disturbances confined to some finite region of x , we seem to head to the same sort of travelling wave solution. Can we find these travelling waves?

We seek a steady wave solution by setting $p(x, t) = f(\xi)$ where $\xi = x - ct$. Without loss of generality, we will take $c > 0$ (we can see the system is symmetric in x , so might as well choose one direction to look for, and we know the other direction will be the same).

Note carefully that at this point, we are just trying some constant c : we have no guarantee that this will work for this c , or indeed for any value of c . In addition, suppose we find some solutions, we don't immediately know from this which travelling wave solution it is that our general initial condition seems to be headed to. Still, the simulations are enough that this looks like it should work, so we carry on in hope.

Substituting this form into Fisher's equation, we get a differential equation for f :

$$-cf' = f'' + f(1 - f).$$

There's more than one way to proceed now, and we look at three apparently-different approaches in this course:

Approach 1 to Fisher's equation: phase-plane analysis Essentially, we are looking for a solution to the DE for f that runs from near $f = 1$ to near $f = 0$. We could think of the second order DE as a coupled first order dynamical system for (f, f') , and use all of our usual tools to analyse it. Set $f' = g$, then:

$$\begin{aligned} f' &= g \\ g' &= -cg - f(1 - f) \end{aligned}$$



Exercise 43:

- Check fixed points are at $(0, 0)$ and $(1, 0)$
- Find the Jacobian at each of the fixed points
- Show that $(1, 0)$ is a saddle (for $c > 0$)
- Show that $(0, 0)$ is a stable node for $c \geq 2$ and a stable focus for $c < 2$



(A stable node has both eigenvalues real and negative, and stable focus has a complex conjugate pair of eigenvalues both with negative real part.)

The nullclines are straightforward: $f' = 0$ corresponds $g = 0$, i.e. the x-axis. And $g' = 0$ gives $g = -c^{-1}f(1 - f)$ which is the parabola in dashed blue in the diagrams below. Using the results of the fixed point types together with a nullcline analysis, it is not too hard to draw the phase-diagram⁵³. Below we will consider $c = 1, 2$ and 3 and consider the corresponding solutions that join $(1, 0)$ with $(0, 0)$. These trajectories between the fixed points correspond with solutions for f which run from near constant $f = 1$ to $f = 0$, as we want.


⁵³and it's even easier if you use mathematica...



This is $c = 1$. These do correspond to valid solutions of the Fisher equation, but as it goes into $f < 0$, they can't correspond to the travelling wave solution that we were looking for.



This is $c = 2$. This turns out to be the one we are after. The origin is a stable node.



This is $c = 3$. These also correspond to valid solutions of the Fisher equation, and are in right range, but turn out to not be the solution that we head to for general initial conditions. Note that the red trajectory is near the nullcline, but not exactly on it (this was where I had problems drawing in lectures)

So with all this we can see that for each $c > 0$, there exists a corresponding steady wave solution. If we want $f \in [0, 1]$ for all of x , then we can rule out the oscillatory solutions (as $f < 0$ for some x). That just leaves us with all of $c \geq 2$ to consider.

Approach 2 to Fisher's equation: linearise near the leading edge Still thinking of a wave moving to the right, we have $f \approx 0$ for large positive ξ . So the dynamics out there ought to be close to that of a system linearised around $f = 0$. Linearising just means ejecting the f^2 term here:

$$-cf' = f'' + f(1 - f).$$

i.e.


$$f'' + cf' + f = 0$$

This is just a linear second order system with constant coefficients, so we can go ahead and solve with⁵⁴ $f = e^{-\lambda\xi}$ for some λ and sub in:

$$\lambda^2 - c\lambda + 1 = 0$$

Normally we would now solve for λ (two solutions) and write down the complimentary function and so on, but it all depends on c , so really focus on the relationship between λ (which determines the shape of the solution) and c (the speed):

⁵⁴note minus in exponent, just so that we want real positive λ to have solutions that decay exponentially towards zero for large ξ .



Wave speed c as a function of the exponent of the shape of the wave λ . The red dot at $(1, 2)$ marks the minimum.

end of
lecture 20

We can see that for real positive λ , we have c takes values in $[2, \infty)$.

Thinking back to the original system we have the Fisher equation in p

$$\frac{\partial p}{\partial t} = \frac{\partial^2 p}{\partial x^2} + p$$

has solution $p = e^{-\lambda(x-ct)}$ if $c = \lambda + 1/\lambda$. We have yet to say why the system settles to one particular such speed and shape. Here, treat this next bit as starred, but solutions of the full nonlinear system will settle to the *slowest* wave speed of the *linearised system*. Here is a sketch argument for this, but treat as starred:

****Arguments for slowest linear speed** First, the nonlinear system speed must be one of the linear speeds (consider the leading edge as above, if we have a steady nonlinear solution, its front must satisfy the linear system). Second, we will show that any nonlinear speed must be slower or equal to any linear speed. Combining these, the nonlinear speed can only be the slowest linear speed.

This can be done more formally (see work of Kolmogorov and others...) but here's an intuitive picture. Our initial conditions are 'localised' to some region, so we can say at least that $p(x, 0) = 0$ for $x > x_0$ for some x_0 . Everywhere, $p(x, 0) \leq 1$. So, we can certainly bound our initial condition above by a curve of the form $\hat{p} = Ae^{-\lambda x}$ by choosing A large enough⁵⁵. Use this \hat{p} as an initial condition for the linearised system. So now we're running two separate systems:

$$p(x, t) \text{ solves } \frac{\partial p}{\partial t} = \frac{\partial^2 p}{\partial x^2} + p - p^2$$

⁵⁵ $A = e^{\lambda x_0}$ will do: then $\hat{p} > 1$ for $x < x_0$ and $\hat{p} > 0$ everywhere

$$\hat{p}(x, t) \text{ solves } \frac{\partial \hat{p}}{\partial t} = \frac{\partial^2 \hat{p}}{\partial x^2} + \hat{p}$$

And we have rigged it so that $p(x, 0) < \hat{p}(x, 0)$. Consider both p and \hat{p} running forward in time now, each according to their own systems. Can p push past \hat{p} ? No, surely no. The extra $-p^2$ pulls the p down. And, even without that, the diffusion tend to pull any protrusions back. To be slightly more formal, consider $g(x, t) = \hat{p}(x, t) - p(x, t)$. We have $g(x, 0) > 0$ for all x .

$$\begin{aligned} \frac{\partial}{\partial t} g(x, t) &= \frac{\partial \hat{p}}{\partial t} - \frac{\partial p}{\partial t} \\ &= \left(\frac{\partial^2 \hat{p}}{\partial x^2} + \hat{p} \right) - \left(\frac{\partial^2 p}{\partial x^2} + p - p^2 \right) \\ &= \frac{\partial^2 g}{\partial x^2} + g + p^2 \\ &> \frac{\partial^2 g}{\partial x^2} + g \end{aligned}$$

So the g system is just diffusion and growth but with an extra p^2 growth, which is non-negative. So if g starts non-negative everywhere, it will remain non-negative for all time.

So we have that \hat{p} bounds p for all time. And \hat{p} is just one of the nonlinear solutions (we could choose the λ and it will trundle along with the appropriate speed $c(\lambda)$ from above). Whatever p does, it cannot get past that big rolling road-block of the linear solution. Hence, if it settles to a wave form with constant wave speed, that speed cannot be faster than $c(\lambda) = \lambda + 1/\lambda$. This all works for any $\lambda > 0$.

If we accept this minimum speed argument, we now know that if the Fisher equation settles to a constant speed, it must have speed $c = 2$.

Approach 3 to Fisher's equation: mechanical analogy Really, this is firmly into the land of overkill with respect to the Fisher equation: we don't learn anything we don't already know from the other two approaches. However, it is included here for completion as it turns out to be useful with other systems where the earlier approaches don't work so nicely.


The idea is to take our differential equation in f and reinterpret it as a mechanical system. We take ξ to be 'time' and f to be position⁵⁶.

$$\underbrace{f''}_{\text{acceleration}} = \underbrace{-cf'}_{\text{force from friction}} - \underbrace{V'}_{\text{Force from potential } V}$$

and take mass to be one to avoid that hassle. The first force is some sort of resistance to motion: proportional to speed and in opposite direction. The second force is from

⁵⁶this is the confusing bit as we really want x to be space, but just think of this as a separate system now for a moment. More serious is that ξ as time will feel backwards.


some potential $V(f)$, so here we need $V'(f) = -f(1 - f)$ hence $V(f) = \frac{1}{2}f^2 - \frac{1}{3}f^3$. Perhaps the easiest way to imagine a potential is as a wire where $V(f)$ gives its height as a function of its position f .



$V(f) = \frac{1}{2}f^2 - \frac{1}{3}f^3$. Think of a bead on a wire of this shape, with a bit of friction.

We are seeking a solution which goes from one to zero as time (ξ) increases. Imagine the bead starts close to that maximum at $f = 1$ and then slides towards the minimum at $f = 0$. If friction coefficient c is low (but not zero), it will overshoot and oscillate around $f = 0$. If friction is large enough, there is no overshoot. We call this overdamping/underdamping in mechanics language.

It would be entirely reasonable to guess at this point that critical damping is at $c = 2$, without bothering to do any further work in this system.



Solving system with $f(0) = 0.99$ and $c = 0.2$ (left) and $c = 3.5$ (right)

When working with this analogy, it is useful to remember that stability will appear to be the wrong way around. In the world of the bead, $f = 1$ is unstable and $f = 0$ is stable. In the world of the Fisher equation, it is the other way around. The explanation is simply that time is reversed: $\xi = x - ct$ is 'time' in bead world, and we have taken $c > 0$.

3.2.2 Bistable systems

For Fisher's equation above, there was one unstable state ($p = 0$) and one stable state ($p = 1$). So maybe it was not too surprising that we would find a travelling wave solution. Thinking in terms of the beneficial mutation: it is not surprising that it will spread and fix everywhere (aka mutant invasion?). It is less obvious what happens when there's *two* stable states: will one will win out if we're at one state for part of space, and the other for another part of space?

We call such a system *bistable*. We actually have a bistable system at the ready from earlier in the course: in section one we studied competition models using the example of wolbachia and mosquitos. After rescaling and so on we had

$$\begin{aligned}\dot{x} &= x \left[x_0 - \frac{y}{x+y} - (x+y) \right] \\ \dot{y} &= y \left[y_0 - \frac{x}{x+y} - (x+y) \right]\end{aligned}$$

Where x represent uninfected mosquitos, y is infected mosquitos. We have that $0 < y_0 < x_0 < 1$. We know there are two stable fixed points $(x_0, 0)$ and $(0, y_0)$ which represent the situation when no mosquitos are infected and when they all are, respectively.

It turns out that we can turn this apparently two-dimensional system into a one dimensional system by working in terms of *proportion*⁵⁷ of mosquitoes that are infected: set $u = y/(x+y)$.

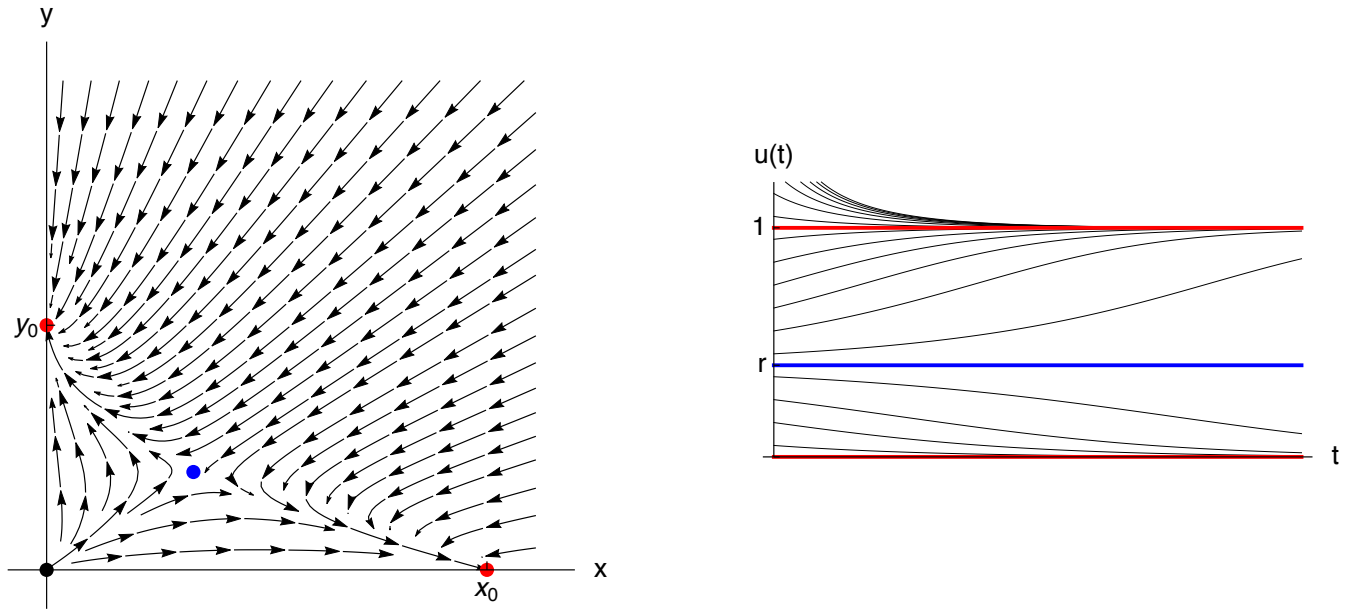
If we compute \dot{u} using quotient rule and substituting the above equations, we find that the dynamics of u depends only on u : any other x, y dependency disappears.

Exercise 44: Show that

$$\dot{u} = -u(u-r)(u-1)$$

where $r = x_0 - y_0$, so we know $r \in (0, 1)$.

⁵⁷We have a small problem at the origin, but that situation is that there are no mosquitoes of any kind, so the model is kind of pointless.



Left: the phase plane in x and y . Right: now u as a function of t .

Dyn Sys note: we're only looking at one dimension now, but we can see from the phase diagram that things collapse fairly quickly to the unstable manifold of that saddle point, so this is not unreasonable. We also can know the basins of attraction of the two stable fixed points⁵⁸: it's just the line $y = rx$.

In fact, this is a nicely generic one-dimensional bistable system: just a cubic giving stable points at 0 and 1 and unstable at $r \in (0, 1)$. We can extend this to a spatial system just by adding diffusion:

$$\frac{\partial u}{\partial t} = -u(u-r)(u-1) + \frac{\partial^2 u}{\partial x^2}$$

And now what happens? Suppose we have initial conditions where we are mostly at zero, but have a patch with $u \approx 1$. This is like we introduce our mosquitos with wolbachia some patch. Will the patch spread outwards or will it shrink? Both 0 and 1 are stable, so which one wins? Things close to 1 should push towards 1, and things close to 0 should push towards 0. But, diffusion hates sharp corners: we're not going to end up with a sharp boundary of all wolbachia and no wolbachia, it must blur out.

It turns out to depend on r , and we can show this by using the techniques we developed above to look for travelling waves, then the mechanics approach is nice here. So, seek a travelling wave solution:

$$u(x, t) = f(\xi) \quad \text{where} \quad \xi = x - ct$$

Sub this in:

$$-cf' = f'' - f(f-r)(f-1)$$

⁵⁸I left this suitably vague in the hand-drawn figures in section 1, but actually they're just straight lines, it turns out

Rearrange to think of as a mechanical system:

$$\underbrace{f''}_{\text{acceleration}} = \underbrace{-cf'}_{\text{friction}} + \underbrace{f(f-r)(f-1)}_{-V'(f)}$$

where $V(f)$ is the potential (e.g. height of bead on a wire). We want

$$-V'(f) = f(f-r)(f-1)$$

So integrate up (choose constant, it does not matter):


$$V(f) = -\frac{1}{4}f^4 + \frac{1}{3}(1+r)f^3 - \frac{1}{2}rf^2$$

Now consider the shape of this potential. There are maxima at 0 and 1 (the two stable fixed points of u) and a minimum at r (the unstable fixed point of u). But, which maximum is higher? It depends on r .

Compare the heights of the maxima:

$$V(0) = 0, \quad V(1) = -\frac{1}{4} + \frac{1}{3}(1+r) - \frac{1}{2}r = \frac{1}{6}\left(\frac{1}{2} - r\right)$$


so it depends: if $r < \frac{1}{2}$ then higher max is at $f = 1$ and if $r > \frac{1}{2}$ it is at $f = 0$. (Note, come back to equal case later.)



Now imagine the dynamics of our bead sliding with friction. Take $c > 0$ so that the friction resists motion. If we start almost stationary near the lower maximum, we clearly cannot slide to the higher maximum. If we start near the higher maximum, then it is *possible*⁵⁹ that we go to the lower max exactly. If friction is too high, we fall short and end up in the minimum. If friction is too low, we will overshoot. So, there is one magic value of c in between where we grind to a halt *exactly* at the lower maximum.

⁵⁹almost like when sports commentators call something a ‘theoretical possibility’: it is a possibility, just not one we think is all likely. However, here we are specifically hunting for the special outcome of going from one of 0 and 1 to the other, not the typical dynamics.


That will give a curve for $f(\xi)$ which goes from 0 to 1 or vice versa. We do not really much care about the shape of this curve, we just want to know which one can happen. Translate back to $u(x, t)$ and we have the wave shape that moves to the right (with the special value of c).




Once the wave has passed through a region, it leaves $u = 1$ for $r < 1/2$ and $u = 0$ for $r > 1/2$. So, if we want wolbachia to take hold in a population, we certainly need $r < 1/2$.

As a bonus, work this back to the original variables to see if it makes sense. We had $r = x_0 - y_0$ where x_0 and y_0 were the carrying capacities of uninfected and infected mosquito populations alone. Then, to say $r < 1/2$ is to say $y_0 > x_0 - 1/2$. In other words, having wolbachia does not reduce the carrying capacity of the population too much.

We know $r < 1/2$ is necessary for a successful introduction of wolbachia. Is it sufficient? Turns out no, not enough in bistable case. If the initial conditions are such that the initial introduction is too small or too narrow, it will not become established.



Here are three realisations with same system ($r=0.3$) and differing initial conditions. The initial conditions are in red (the blocky shape) and subsequent time points go through the rainbow.



In the top run, the initial window is too narrow: diffusion of nearby uninfected mosquitos will pull that initial spike down.

In the middle run, the window of introduction is just wide enough, and you can see several curves close together as the distribution hangs on the brink of collapse, but it just gets through, then settles to steady travelling wave each way.

In the bottom run, the initial window is very wide, and the solution quickly reshapes to be the travelling wave each way.

Exercise 45: Suppose insecticide kills mosquitos at some constant per capita rate (and it is the same for both infected and uninfected mosquitos). How will this change the required introduction width?

end of
lecture 21

3.3 Spatial instabilities

So far in this section, we have had one variable at a time (concentration of bacteria, proportion of mutants, infected mosquitos). Here, we now extend to two variables. This is not just a more complicated version of the above, but it turns out that with two quantities, we have some rather interesting new outcomes.

3.3.1 Chemotaxis

Chemotaxis sounds like something scary, but it just means the movement of something in response to a chemical. Here, think of small organisms (bacteria or amoebas, say) which move around slightly randomly, but generally prefer to go up (or down) the gradient of the chemical. If go up the gradient, then we call the chemical a chemoattractant (it is like the bacteria are looking for the highest concentration, as they are attracted by the chemical). The chemical could be some nutrients that the bacteria are hunting.

Or perhaps the chemical were excreted by the organisms themselves. For example, if amoebas excrete a chemical and are also attracted by it, they should all end up clustering. If they are repelled, then one might expect that to cause the population to spread out.

Let us suppose we are dealing with a population of bacteria denoted by $n(x, t)$ and some chemical $c(x, t)$. We can still use the transport equation for each of n and c . We finally use the active motion component, and say bacteria move speed v . There is still diffusion in there: think of this as a biased random walk⁶⁰.

$$\begin{aligned}\frac{\partial n}{\partial t} + \nabla \cdot (n\mathbf{v} - D_n \nabla n) &= F(n, c) \\ \frac{\partial c}{\partial t} + \nabla \cdot (-D_c \nabla c) &= G(n, c)\end{aligned}$$

where D_n and D_c are the diffusion constants of the bacteria and chemical respectively. The right hand side (F and G) represent all the local dynamics: all the non-spatial processes, e.g more or less anything from section 1 of this course.

The active movement velocity \mathbf{v} is for the chemotaxis. Let us take $\mathbf{v} = \chi \nabla c$, which means if $\chi > 0$ then the bacteria are swimming up the gradient of the chemical. We could in general have $\chi(c)$ so that the chemotaxis rate also depends on the chemical concentration, similar to the non-linear diffusion earlier, but here we will just take it to be constant.

For the local dynamics:

$$\begin{aligned}F(n, c) &= \gamma - \delta n \\ G(n, c) &= \alpha n - \beta c\end{aligned}$$

where all the constants ($\alpha, \beta, \gamma, \delta$) are positive. This represents a simple import (γ) and death (δ) model for the bacteria, and the chemical being excreted (α) by the bacteria and decaying (β). So now our system is:

$$\begin{aligned}\frac{\partial n}{\partial t} + \nabla \cdot (n\chi \nabla c - D_n \nabla n) &= \gamma - \delta n \\ \frac{\partial c}{\partial t} + \nabla \cdot (-D_c \nabla c) &= \alpha n - \beta c\end{aligned}$$

and note we have not actually specified what dimension of space we are working in, but it turns out below that this does not matter.

These tactics are the general to take with spatial systems:

- Deal with the spatially homogeneous system first

⁶⁰or think back to Fokker-Plank: there are analogous drift and spread terms here

- Consider a small spatial perturbation around a fixed point
- Suppose the perturbation has the form $\cos(kx)$ (or $\cos(\mathbf{k} \cdot \mathbf{x})$ if more dimensions)

So working through for this example:

Spatially homogeneous first Here we treat everything as if it were constant in space. We throw out the diffusion and chemotaxis term and set $n(x, t) = n(t)$ and $c(x, t) = c(t)$. Then we just have

$$\begin{aligned}\dot{n} &= \gamma - \delta n \\ \dot{c} &= \alpha n - \beta c\end{aligned}$$

Solve for the fixed points, and find there is always one:

$$n^* = \frac{\gamma}{\delta}, \quad c^* = \frac{\alpha\gamma}{\beta\delta}$$

Find the Jacobian at the fixed point:

$$J_{(n^*, c^*)} = \begin{pmatrix} -\delta & 0 \\ \alpha & -\beta \end{pmatrix}$$

The Jacobian is triangular, so we could just read off the eigenvalues and see they are both negative so the fixed point is stable. For completeness though, note that the trace $tr = -(\delta + \beta)$ and determinant $det = \delta\beta$. We will spot these again below in the spatial analysis.

Consider spatial perturbations Start with a strategic rearrangement (put the diffusion and chemotaxis on the right):

$$\begin{aligned}\frac{\partial n}{\partial t} &= \gamma - \delta n + D_n \nabla^2 n - \nabla \cdot (n \chi \nabla c) \\ \frac{\partial c}{\partial t} &= \alpha n - \beta c + D_c \nabla^2 c\end{aligned}$$

Then perturb around the fixed point:

$$\begin{aligned}n(x, t) &= n^* + u(x, t) \\ c(x, t) &= c^* + v(x, t)\end{aligned}$$

where both u and v are small, so we will linearise in these. The only tricky term is the chemotaxis one, but note that grad of a constant is zero so we lose ∇c^* immediately. Also, we are linearising so throw out $u \nabla v$ when it occurs in the chemotaxis term:

$$\nabla \cdot (n \chi \nabla c) = \nabla \cdot ((n^* + u) \chi \nabla (c^* + v)) = \nabla \cdot ((n^* + u) \chi \nabla (v)) = \nabla \cdot (n^* \chi \nabla (v)) = \frac{\gamma}{\delta} \chi \nabla^2 v$$

so it is not so bad in the end. Continue subbing in our perturbation, linearising:

$$\begin{aligned} u_t &= -\delta u & + D_n \nabla^2 u - \frac{\gamma}{\delta} \chi \nabla^2 v \\ v_t &= \alpha u - \beta v & + D_c \nabla^2 v \end{aligned}$$

We should consider taking general time initial conditions for u and v , but we could consider them as a sum of Fourier modes⁶¹. In fact, in this linear system we find that each mode evolves on its own, so we can just consider a general mode.

If we are doing one-dimensional space, we could just look at perturbations with spatial form $\cos(kx)$. In general though, we could have vector k and x so we are looking at $\cos(\mathbf{k} \cdot \mathbf{x})$ and it will all work out the same. Some text books use $e^{i\mathbf{k} \cdot \mathbf{x}}$ instead: I personally prefer $\cos(kx)$, which makes it easier to think about the nature of the perturbation in terms of bacteria and chemicals, but it is all mathematically equivalent, so feel free to use whichever you like. Whichever approach, the key thing is that ∇^2 becomes $-k^2$ (or perhaps we should say $-|\mathbf{k}|^2$ for the vectors).

Set

$$u(x, t) = \hat{u}(t) \cos kx, \quad \text{then} \quad \nabla^2 u = -k^2 \hat{u} \cos kx$$

and similarly

$$v(x, t) = \hat{v}(t) \cos kx, \quad \text{then} \quad \nabla^2 v = -k^2 \hat{v} \cos kx$$

Substitute these in, cancel the $\cos kx$ everywhere, then our spatial system with perturbation of the form $\cos kx$ has become

$$\begin{aligned} \hat{u}_t &= -\delta \hat{u} & -k^2 D_n \hat{u} + k^2 \frac{\chi \gamma}{\delta} \hat{v} \\ \hat{v}_t &= \alpha \hat{u} - \beta \hat{v} & -k^2 D_c \hat{v} \end{aligned}$$

There are lots of approaches continuing on from here, but I think the simplest is to reuse as much of our existing machinery as possible. We can read off the coefficients here to get a *modified Jacobian*: some terms are as before, and some appear with the spatial process. Say J_{mod} is the Jacobian in the $\hat{u} \hat{v}$ system (i.e. under $\cos kx$ perturbation). Writing it strategically

$$J_{mod} = \underbrace{\begin{pmatrix} -\delta & 0 \\ \alpha & -\beta \end{pmatrix}}_{\text{Local dynamics}} - k^2 \underbrace{\begin{pmatrix} D_n & 0 \\ 0 & D_c \end{pmatrix}}_{\text{Diffusion}} + k^2 \underbrace{\begin{pmatrix} 0 & \frac{\chi \gamma}{\delta} \\ 0 & 0 \end{pmatrix}}_{\text{Chemotaxis}}$$

The local dynamics are represented by the old Jacobian (which is exactly as we analysed above when considering the homogenous system). The diffusion is $-k^2$ times a diagonal matrix, where the entries are the diffusion constants. The chemotaxis is represented also by a matrix with k^2 factor, and here is a single off diagonal entry (the bacteria are undergoing chemotaxis due to the chemical).

⁶¹if this is confusing, look back at your Ib Methods notes: superposition of solutions, completeness of eigenfunctions and so on, or just go with the flow here.

As usual, we check for stability: it is stable if trace is negative and determinant is positive. We already know that without the spatial components (which we can see is $k = 0$) then indeed these conditions are met. With the spatial components, the trace only becomes even more negative

$$Tr(J_{mod}) = \underbrace{-(\delta + \beta)}_{Tr(J) < 0} - k^2(D_n + D_c) < 0$$


So, it comes down then to the determinant. If that is positive for a particular value of k , then the system is stable for that k . If it is positive for *all* (real) values of k , then we can say the system is stable to small spatial perturbations of any kind. If the determinant is negative for *any* k , then we have a spatial instability. So, let us go ahead and investigate the determinant, remembering that we're interesting in what happens for the full real range of k (wlog $k > 0$).

$$\begin{aligned} Det(J_{mod}) &= \left| \begin{pmatrix} -\delta - k^2 D_n & k^2 \frac{\chi \gamma}{\delta} \\ \alpha & -\beta - k^2 D_c \end{pmatrix} \right| \\ &= (k^2 D_n + \delta)(k^2 D_c + \beta) - k^2 \chi \frac{\alpha \gamma}{\delta} \\ &= \underbrace{D_n D_c}_{\text{diffusions}} k^4 + \underbrace{\left(\delta D_c + \beta D_n - \chi \frac{\alpha \gamma}{\delta} \right)}_{\text{interactions between local and spatial}} k^2 + \underbrace{\delta \beta}_{\text{old det}} \end{aligned}$$

and so we have a quadratic in k^2 . The coefficient of k^4 is the product of the diffusion coefficients, and so is positive. The constant is the old determinant. This should not be a surprise: if $k = 0$, then perturbing with spatial form $\cos kx$ is just the same as old homogeneous perturbation (old trace is also recovered above). We were working with a fixed point that is stable to homogeneous perturbations, so this det is positive. The middle term (coefficient of k^2) is the fun one. Recognising the parameters, we can see each term is a product of local (α , β and δ) and spatial (D_n , D_c and χ) parameters.

Can this $Det(J_{mod})$ be negative for some real k ? Well, yes. Just make the chemotaxis coefficient χ big enough and Det goes negative and hence we have an instability for that k . We can rearrange to get a condition for negative det (easier to reverse to second to last line and rearrange for χ):

$$\frac{\alpha \gamma}{\delta} \chi > \frac{(k^2 D_n + \delta)(k^2 D_c + \beta)}{k^2}$$



If we imagine starting with χ small, then there is no real k for which there is instability. If we increase χ , the first k_c that goes unstable is the minimum of the RHS. We call

this the *critical wavenumber*: the value of k when the instability first appears. With a general perturbation, it is this mode ($\cos kx$) that will grow first.

Exercise 46: Show that we will have instability for $\chi > \frac{\delta}{\alpha\gamma} (\sqrt{\delta D_c} + \sqrt{\beta D_n})^2$. Find k_c and check its dimensions.

end of
lecture 22

Think back to the original system: bacteria excrete chemical and also head up the chemical gradient by chemotaxis. It's hardly a surprise that it is unstable if there is enough chemotaxis: the bacteria are going to head to cluster together by calling to each other with the chemical. It is fairly interesting though that the instability is first at an intermediate value of k (i.e. not k very small or very large), and we will see more of this below.

The fact that there is instability here is down to the chemotaxis. If $\chi = 0$ there is no spatial instability. So, how about in general without chemotaxis? Consider general reaction-diffusion systems. Suppose we have some system where the local dynamics is stable (around some fixed point). If the only spatial process is diffusion, then alone is also going to support stable dynamics as it just smoothes stuff out. Can the combination cause an instability?

In one dimension, we already know the answer. In section 3.1.2 we had a diffusion equation with linear growth, and found that the growth term is beaten by the diffusive spread in an infinite domain. In 3.2.0 we found that a more general local term still behaves like a linear term for small perturbations, so nothing much changes. So in one dimension, there's no possibility of a spatial instability if the local dynamics are stable and all we have is diffusion.

In two or more dimensions, well, read on...

3.3.2 Turing instabilities

We will now see that in two (or more) dimensions, it is possible to have a spatial instability in a system which is stable to homogenous diffusion and the only spatial process is linear diffusion. This result was first shown in a paper by Alan Turing published in 1952. This result is not just mathematically interesting, but also is exciting as it gives a mechanism for patterns to emerge from a homogenous system. This is important in developmental biology: how do things grow? For anything from coat patterns (spots and stripes) to anatomy of any growing creature, there must be some mechanism to create spatial structure starting from something uniform. The beautiful thing is that we do not need to write down some extremely complicated system, but we can set it up from two chemicals (morphogens here) in the type of mathematical language we have used above.

So, to the general theory: start with two chemical concentrations, u and v , which are both functions of space and time, and write down a general system where they have

some local dynamics (given by functions f and g) and their spatial dynamics are just linear diffusion (with rates D_1 and D_2):

$$\begin{aligned}\frac{\partial u}{\partial t} &= f(u, v) + D_1 \nabla^2 u \\ \frac{\partial v}{\partial t} &= g(u, v) + D_2 \nabla^2 v.\end{aligned}$$

We will first analyse the *homogenous* system and suppose we find a stable fixed point, then we look at what happens to spatial perturbations about this fixed point. In other words, this tracks through much as the chemotaxis example above, but we work with a general system here with just diffusion.

First, the *homogenous* system then. Formally, this is considering the case where u and v are constant in space but may vary in time, so we are setting $u(x, t) = u(t)$ and $v(x, t) = v(t)$. More simply in practice, just cross out the diffusion terms and we're back to Section 1. We are seeking fixed points, i.e. we want u^* and v^* such that $f(u^*, v^*) = g(u^*, v^*) = 0$. In particular, we want *stable* fixed points. We write the Jacobian at this fixed point as:

$$J = \begin{pmatrix} f_u & f_v \\ g_u & g_v \end{pmatrix}$$

where this is slightly lazy notation as we really mean these all evaluated at the fixed point, i.e.

$$f_u = \left. \frac{\partial f}{\partial u} \right|_{(u^*, v^*)}, \quad f_v = \left. \frac{\partial f}{\partial v} \right|_{(u^*, v^*)} \quad \text{etc.}$$

For a Turing instability, we want the homogenous system to be stable (otherwise it's rather boring). So we require:

$$\begin{aligned}Tr(J) &= f_u + g_v < 0 \\ Det(J) &= f_u g_v - f_v g_u > 0\end{aligned}$$

for stability. For given functions f and g , this usually boils down to some nice parameter range. So now suppose that we have a fixed point, and that we are in the parameter range where it is stable from now on.

Now, consider spatial perturbations about this homogenous fixed point of this form:

$$\begin{aligned}u(x, t) &= u^* + \hat{u}(t) \cos kx \\ v(x, t) &= v^* + \hat{v}(t) \cos kx\end{aligned}$$

where \hat{u} and \hat{v} are functions of time (but not space) and we are interested in the dynamics when they are small. So, speeding through this (as it goes much like the

chemotaxis example above), we know that the \cos will cancel everywhere. This will just leave a factor $-k^2$ with the diffusion, so to linear order we have:

$$\begin{aligned}\hat{u}' &= f_u \hat{u} + f_v \hat{v} - k^2 D_1 \hat{u} \\ \hat{v}' &= g_u \hat{u} + g_v \hat{v} - k^2 D_2 \hat{v}\end{aligned}$$

where prime denotes differentiation wrt t . Hence, reusing the notation of the modified Jacobian, we can read off:

$$J_{mod} = \underbrace{\begin{pmatrix} f_u & f_v \\ g_u & g_v \end{pmatrix}}_{\text{original } J} - k^2 \underbrace{\begin{pmatrix} D_1 & 0 \\ 0 & D_2 \end{pmatrix}}_{\text{from diffusion}}.$$

To see if the spatial perturbation is stable or unstable, we can now just work with this modified Jacobian. Calculate the trace and det, and try to pull out terms which are the trace and det from the original Jacobian:

$$Tr(J_{mod}) = \underbrace{(f_u + g_v)}_{\text{old } Tr} - k^2(D_1 + D_2)$$

and as the trace of the homogenous J was negative (as we decided we were going to look at stable fixed points only) this trace of J_{mod} is certainly negative. So, that's dealt with.

The real fun begins with the determinant:

$$Det(J_{mod}) = \underbrace{D_1 D_2}_{>0} k^4 - \underbrace{(D_1 g_v + D_2 f_u)}_{\text{the magic!}} k^2 + \underbrace{(f_u g_v - f_v g_u)}_{\text{old } Det}$$

Looking at this as a quadratic in k^2 , there are three terms. The coefficient of k^4 is just the product of the diffusion coefficients, so this is always positive. At the other end, the term which is constant in k is just the original determinant. This is not a surprise as $k = 0$ just recovers the the original Jacobian. The term in between is the really interesting one. Notice the k^4 term contains just constants from diffusion, nothing from the local dynamics. The constant term is the other way: just the local dynamics, and nothing from the diffusion process. The term in the middle somehow represents the interaction between the diffusion and the local dynamics: there are both D_1, D_2 and f, g in there.


Is this perturbation stable? It is if both the trace is negative (which it is) and the det is positive (which is in question). So, is it possible that $det < 0$ for some value of k (or some range of values of k)?

Rewrite as follows⁶²:

$$Det(J_{mod}) = Ak^4 - Bk^2 + C$$

⁶²Note the minus sign before B . Can do it either way, but this makes more sense to me in the end.

where we have already that A, C are positive. For this to be negative, we certainly need $B > 0$ (sometimes this is referred to as the ‘necessary’ condition). Suppose that is true for now. Think about this graphically, plotting determinant against k^2 . It is a parabola, with a minimum at $k^2 = B/2$ which is positive. We can now see that it depends whether the minimum is above or below zero:



Parabolae with $A, B, C > 0$: one in blue does not have any real roots and never negative (in real k^2). One in gold has real roots, so there is a range of $k^2 > 0$ where $Det < 0$.

In fact, within $B > 0$ we can now see we also need there to be real roots⁶³ for $Det < 0$ in some range of (real) k , hence we need positive discriminant: $B^2 - 4AC > 0$.

Pulling this together, we have spatial instability for some k iff

$$B > 0 \quad \text{AND} \quad B^2 - 4AC > 0$$

(and this pair of conditions together are now necessary and sufficient). Of course, as we know the signs of A and C we can get this down to being equivalent to a single condition:

$$B > 2\sqrt{AC}.$$

This might look simpler, but sometimes in a given example it is easier to use the two conditions.

Putting back in the full expressions for A, B and C we can then write down a condition in full:

$$D_1 g_v + D_2 f_u > 2\sqrt{D_1 D_2 (f_u g_v - f_v g_u)}$$

condition for Turing Instability

⁶³Plural if we're being careful: equal roots just means it touches $Det = 0$, but we actually need it negative to be sure of instability so use strict inequalities here.

This is a nice core result here. However, please do not put any effort into memorising the expression here: the condition for a given example can usually be derived more directly than plugging things into here.

Imagine now that there is some parameter (within f or g) that we are varying, and for some range this condition is not met, and for some other range, it is met. The transition as we move between them is for no instability to having instability for some range of k . The value of k that first is unstable as we transition is the *critical wave number*, k_c . Algebraically, it is when our quadratic in k^2 has a double root. In other words the double root of $Ak^4 - Bk^2 + C = 0$ in the case when $B^2 - 4AC = 0$. One can write this value down in many ways e.g.

$$k_c^2 = \sqrt{\frac{C}{A}} \quad \text{or} \quad k_c^2 = \frac{B}{2A}$$

and they will all be equivalent so just choose whichever one is convenient. It is usually possible to eliminate one parameter.

Finally, we can think about Turing Instabilities in a few different ways to try to build more intuition about what they are and how they occur:

Thinking about wavelength Our perturbation was of the form $\cos kx$ so the wavelength is $2\pi/k$.

- If $k \ll 1$: these are really long waves. On the scale that diffusion operates on, these look to be almost constant. So diffusion does not do very much, and hence the dynamics are approximately the local dynamics alone. Indeed, we can check above and see that for small k , $J_{mod} \approx J$. And in our quadratic for the determinant, it is indeed approximately the old determinant. If we have conditioned on the local dynamics being stable, then we should not find instability for $k \ll 1$.
- If $k \gg 1$: these are really short waves, with wiggles with high curvature. Diffusion is fast to eliminate wiggles, so in this range expect diffusion to dominate, and hence things to be stable, no matter the local dynamics.

Hence we can start to understand why instability can only occur for some intermediate range of k . Either regime of diffusion or local dynamics dominating will result in stability. The only possibility is somewhere in between.

Thinking about the algebra This is a nice way to grasp why we really have to have a system of two or more dimensions (u and v) and to discover that we must have differing diffusion rates. We can see in general that in general our modified Jacobian will be the old Jacobian minus k^2 times some diagonal matrix with positive elements:

$$J_{mod} = J - k^2 D$$

where D is a matrix with D_1, D_2 etc down the diagonal.

For a Turing instability, we want the local dynamics to be stable, but not the spatial dynamics (for some values of k) – indeed perhaps this is the definition of a Turing instability. In terms of matrices then, we want all eigenvalues of J to have negative real part, but not true for J_{mod} .

Brief aside: If a matrix M has eigenvector \mathbf{v} with eigenvalue λ , then matrix $M - \mu I$ (where I is the identity matrix) has the same eigenvector but now the eigenvalue is $\lambda - \mu$.⁶⁴

So if the diagonal matrix D was in fact a multiple of I , then all of our new eigenvalues would be the old ones minus some real positive number, and hence still stable. Hence we cannot have $D_1 = D_2$ and still get a Turing instability. Or in more dimensions, we can't have the diffusion coefficients all be equal. In one dimension? Well, it is consistent: we really can't have a Turing instability.

Back to two dimensions: might be confusing that Turing instabilities are possible at all now. It still feels like subtracting stuff off the diagonal should make the eigenvalues have more negative real part. Maybe this tends to be true more often than not, but it certainly is possible to subtract things off the diagonal and grow the real part of *one* of the eigenvalues, which is all we need.

All of this is mainly to build intuition about what is going on here, but it is also slightly useful in practice: we often have $D_1 = 1$ and $D_2 = d$ or similar. Then we know $d = 1$ cannot be in the range where a Turing instability is possible. This often a useful sanity check in working through a particular example.

Activator-Inhibitor dynamics Turing instabilities are often introduced with the mechanistic explanation being in terms of 'activator-inhibitor' dynamics. Think of u and v as being some chemicals and here one is an 'activator' so its presence generally increases one/both of them and the other decreases one/both. This doesn't work perfectly for all examples, but it is still a useful way to think about it sometimes. We can connect it with the algebra above:

$$B > 0 \quad \text{corresponds to} \quad D_2 f_u + D_1 g_v > 0$$

$$Tr < 0 \quad \text{corresponds to} \quad f_u + g_v < 0$$

So we have two different linear combinations of f_u and g_v with positive coefficients, and one combination is positive and the other is negative. So we must have exactly one of f_u and g_v positive and one negative. Without loss of generality, say $f_u < 0 < g_v$ and then we would naturally think of f as an inhibitor: $f_u < 0$ means that it dampens itself back down (1-D dynamics on its own are stable). Similarly g is the activator. Looking at the coefficients above, we can see in this case we need $D_1 > D_2$ ⁶⁵ Hence the inhibitor

⁶⁴Easy to check: $(M - \mu I)\mathbf{v} = M\mathbf{v} - \mu I\mathbf{v} = \lambda\mathbf{v} - \mu\mathbf{v} = (\lambda - \mu)\mathbf{v}$.

⁶⁵Either stare at the two inequalities for a while, or do something more organised like consider $\underbrace{(D_2 f_u + D_1 g_v)}_{>0} - D_2 \times \underbrace{(f_u + g_v)}_{<0} = (D_1 - D_2)g_v$.

diffuses more than the activator. We also must have $|g_v| < |f_u|$ in this case, so the inhibitor is ‘stronger’ in some sense.

This can make intuitive sense: the activator works very locally to stimulate the perturbation. The inhibitor spreads out and pushes things a bit further away in the opposite direction. So if the perturbation is flat, it gets pushed back. If the perturbation is wavy with about the right wavelength, then it may grow.

This is useful for building intuition as it helps underline why there must be (at least) two quantities. It also shows in another way that Turing instabilities are about how the local dynamics interact with diffusion in just the right way at approximately the right wavelengths. But, be cautious not to push this too far: for setting up which is the activator and which is the inhibitor, we have not thought about how the two quantities interact with each other (which are encapsulated, at least to first order, by f_v and g_u).

end of
lecture 22

Worked example: Finally to finish this section, and the course, here is an example question on Turing instabilities:

Consider this system:

$$\begin{aligned}\frac{\partial u}{\partial t} &= -u + u^2v + \nabla^2 u \\ \frac{\partial v}{\partial t} &= b - u^2v + d\nabla^2 v.\end{aligned}$$

Find condition for this system to have a Turing instability. Find the critical wavenumber k_c on the boundary of instability.

To illustrate to tackle these problems in practice, the solution (with annotation in red) is in my handwriting on the remaining pages.

Worked Example

$$Q: \frac{\partial u}{\partial t} = -u + u^2 v + \nabla^2 u \quad b, d > 0$$

$$\frac{\partial v}{\partial t} = +b - u^2 v + d \nabla^2 v$$

Find condition for this system to have Turing instability.

Find k_c on boundary.

Solution: Thoughts: expect $d=1$ not in range.
.. check later.

Homogeneous: FP $u^* = b, v^* = b^{-1}$

$$J|_{FP} = \begin{pmatrix} -1+2uv & u^2 \\ -2uv & -u^2 \end{pmatrix}|_{FP} = \begin{pmatrix} 1 & b^2 \\ -2 & -b^2 \end{pmatrix} \quad \begin{array}{l} \text{tr} = 1 - b^2 \\ \text{det} = b^2 > 0 \end{array}$$

need $b > 1$ for FP to be stable

know can invoke $b > 1$ below when needed.

Spatial: perturb by $\cos kx$.

$$J_{\text{mod}} = \begin{pmatrix} 1 & b^2 \\ -2 & -b^2 \end{pmatrix} - k^2 \begin{pmatrix} 1 & 0 \\ 0 & d \end{pmatrix} \quad \text{tr} < 0$$

can write more lines here.

$$\det = \underbrace{d}_{A > 0} k^4 - \underbrace{(d - b^2)}_B k^2 + \underbrace{b^2}_{C > 0}$$


" $B > 0$ " need $d - b^2 > 0$ i.e. $\frac{d}{b^2} > 1$ \star_1 L24.2

" $B^2 - 4AC > 0$ " $(d - b^2)^2 - 4db^2 > 0$

$$d^2 - 6db^2 + b^4 > 0$$

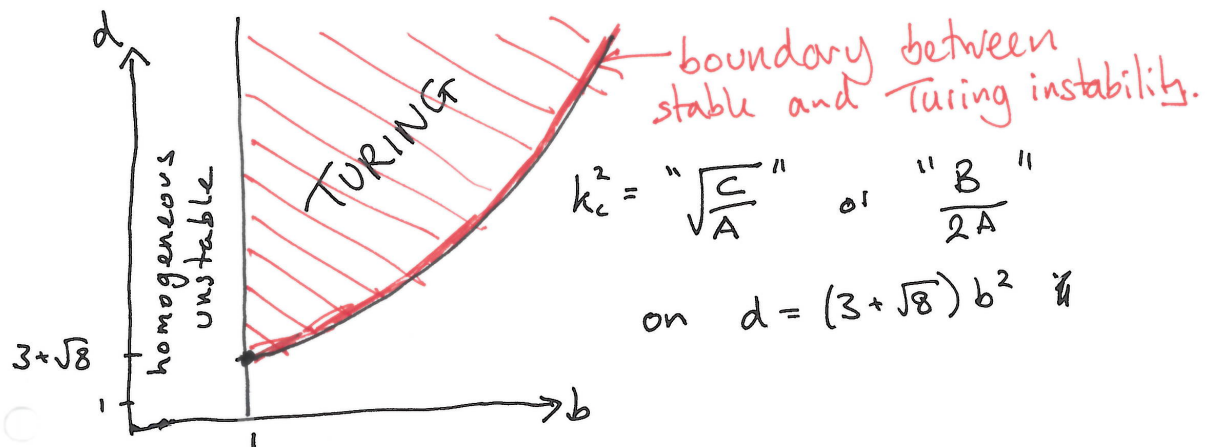
$$\left(\frac{d}{b^2}\right)^2 - 6\left(\frac{d}{b^2}\right) + 1 > 0 \quad \star_2$$

$X^2 - 6X + 1 > 0$ has roots $3 \pm \sqrt{8}$



combine \star_1, \star_2 :

$$\boxed{\frac{d}{b^2} > 3 + \sqrt{8} \quad (b > 1)}$$



ie $k_c^2 = \sqrt{\frac{b^2}{d}} = \sqrt{\frac{1}{3 + \sqrt{8}}} = \frac{1}{1 + \sqrt{2}} = \sqrt{2} - 1$.

so $k_c = \sqrt{\sqrt{2} - 1}$.