

2 Discrete Time

Sometimes it is more natural to model populations and other biological quantities using discrete, rather than continuous, time. A species may give birth only in one particular season. Or you may be interested in how some mutation evolves from one generation to the next. Alternatively, it may be that you really do want to think of time as continuous but you're solving an equation numerically and want to know the pitfalls that could arise when making time discrete.

Whatever the motivation, in this section we will study the behaviour of systems of a single variable x_n where the index $n \in \mathbb{Z}$ plays the role of time. In the simplest case, the variable evolves through some function

$$x_{n+1} = f(x_n) . \tag{2.1}$$

We will learn that one dimensional systems of this kind can be significantly richer than their differential equation counterparts.

2.1 Linear Examples

We start our discussion by looking at a couple of simple examples. These won't bring any substantially new dynamics to the table beyond what we've seen in continuous systems. But they will allow us to get used to some features of discrete time dynamics.

2.1.1 Hello Poppy

Poppies are annual flowers, living for just one year. Their seeds sit in the ground over winter. Some fraction of them germinate the following year, some the year after that, and some not at all. We would like to model this mathematically.

We'll introduce the following variable

- Let x_n be the number of plants in season n .
- Let γ be the number of seeds produced by each plant.
- Let σ be the probability that a seed germinates after one year.
- Let τ be the probability that, having failed to germinate the first year, a seed is successful the next.

Now we're in business and can write down the equation that describes the poppy population: it is

$$x_n = \sigma\gamma x_{n-1} + \tau(1 - \sigma)\gamma x_{n-2} . \quad (2.2)$$

What kind of solutions should we be looking for? We could look for a steady state, but generically the only one is $x^* = 0$. And, because the equation is linear, any small perturbation around this just brings us back to (2.2).

Instead, motivated by the form of the equation (2.2), we look for solutions of the form

$$x_n = p^n \quad (2.3)$$

for some p . If we substitute this into (2.2), we get a quadratic

$$p^2 = \sigma\gamma p + \tau(1 - \sigma)\gamma . \quad (2.4)$$

This has roots

$$p_{\pm} = \frac{\sigma\gamma}{2} \pm \frac{1}{2} \sqrt{\sigma^2\gamma^2 + 4\tau(1 - \sigma)\gamma} . \quad (2.5)$$

We have $p_- < 0 < p_+$ and $|p_+| > |p_-|$. The general solution takes the form

$$x_n = Ap_-^n + Bp_+^n . \quad (2.6)$$

Because $|p_+| > |p_-|$, the second term will dominate at large n . The question is: does the population grow or shrink?

This depends on the size of p_+ . If $p_+ > 1$ then the population grows over time; if $p_+ < 1$ then it shrinks. The critical case is

$$p_+ = 1 \quad \implies \quad \gamma[\sigma + \tau(1 - \sigma)] = 1 . \quad (2.7)$$

This makes sense: the quantity $[\sigma + \tau(1 - \sigma)]$ is the probability that a seed germinates, either in the first year or the second. So the quantity $\gamma[\sigma + \tau(1 - \sigma)]$ is the average number of offspring that a given plant produces. If this number is greater than one, then $p_+ > 1$ and the poppies flourish. If this number is less than one, then $p_+ < 1$ and it's goodbye poppy.

2.1.2 Breathe Again

In Section 1.2.3, we introduced a model of breathing in which the volume of the breath, V , depends on the concentration of CO_2 in your blood. Because breaths are things you can count, it makes sense to construct such model using discrete time.

We previously introduced the breathing model to illustrate delay equations. We'll keep this feature, with the volume of the n^{th} breath V_n determined by the concentration of CO₂ k breaths previously,

$$V_{n+1} = \alpha C_{n-k} . \quad (2.8)$$

We will then model the change in the CO₂ level by the equation

$$C_{n+1} - C_n = M - \beta V_{n+1} . \quad (2.9)$$

Here $M, \alpha, \beta > 0$. Note that this model is not a straightforward discretisation of our previous differential equation (1.52). Indeed, here we have a rather simple linear system, contrasting with the more complicated non-linear delay differential equation that we previously studied.

We can eliminate the volume V_n entirely, and focus just on the CO₂ concentration,

$$C_{n+1} = M + C_n - \alpha\beta C_{n-k} . \quad (2.10)$$

There is a steady-state solution given by $C_n = C^*$ with

$$C^* = \frac{M}{\alpha\beta} . \quad (2.11)$$

Now we can look at perturbations away from this steady state. We will be particularly interested in how the qualitative behaviour of the solutions depends on the delay parameter k .

$k = 0$: To kick things off, let's analyse (2.10) when $k = 0$ so there is no delay and the volume of breath depends on the present concentration of CO₂. We perturb around the fixed point and write $C_n = C^* + \epsilon_n$ where, as the name suggests $\epsilon_n \ll 1$. Then we have

$$\epsilon_{n+1} = (1 - \alpha\beta)\epsilon_n . \quad (2.12)$$

We see that the steady state is stable provided that $\alpha\beta < 2$ and is unstable for $\alpha\beta > 2$.

$k = 1$: Now what happens if we introduce the smallest possible delay $k = 1$? Perturbing around the fixed point, the equation (2.10) becomes

$$\epsilon_{n+1} = \epsilon_n - \alpha\beta\epsilon_{n-1} . \quad (2.13)$$

This time we're going to look for solutions of the form $\epsilon_n = p^n$ for some p . The equation above then becomes a quadratic:

$$p^2 - p + \alpha\beta = 0 \quad \implies \quad p = p_{\pm} = \frac{1}{2} \left(-1 \pm \sqrt{1 - 4\alpha\beta} \right) . \quad (2.14)$$

For $\alpha\beta < \frac{1}{4}$, both p_+ and p_- are real and, moreover $|p_{\pm}| < 1$. Then we have the general solution

$$\epsilon_n = Ap_+^n + Bp_-^n \quad (2.15)$$

which decays as $\epsilon \rightarrow 0$ as $n \rightarrow 0$. We see that, once again, the fixed point C^* is stable for $\alpha\beta$ small enough.

For $\alpha\beta > \frac{1}{4}$, the roots p_{\pm} become complex. This means that the system now oscillates about the fixed point but doesn't otherwise change our approach. We still want to know if $|p_{\pm}|$ is less than one, and hence the fixed point is stable, or greater than one and hence unstable. We have

$$|p_{\pm}|^2 = \frac{1}{4} + \left(\alpha\beta - \frac{1}{4}\right) = \alpha\beta . \quad (2.16)$$

So we learn that, complex oscillations aside, the system is stable if $\alpha\beta < 1$ and unstable if $\alpha\beta > 1$. The upshot of this is that the delay reduces the range of $\alpha\beta$ over which the system is stable.

We can look more closely what happens at $\alpha\beta = 1$ and the system becomes unstable. Here we have

$$p_{\pm} = \frac{1}{2} \left(-1 \pm \sqrt{3}i\right) = e^{\pm\pi i/3} . \quad (2.17)$$

We see that the system has 6-fold periodicity at this point, with $p_{\pm}^6 = 1$.

2.2 The Logistic Map

The fun with discrete maps really gets going when we look at non-linear maps. A great deal of all that's interesting and surprising about these maps can be found lurking inside the deceptively simple example

$$x_{n+1} = f(x_n) = rx_n(1 - x_n) . \quad (2.18)$$

This is the *logistic map*. Understanding the mysteries of the logistic map will occupy us for the rest of this section.

We will take our parameter to lie in the range $x_n \in [0, 1]$. The logistic map keep us within this range provided that the parameter r is bounded by

$$0 \leq r \leq 4 . \quad (2.19)$$

If $r = 4$, the maximum value of the logistic map (at $x_n = 1/2$) gives $x_{n+1} = 1$. Much of our interest will be focussed on how the dynamics of the logistic map changes as we vary r .




Figure 29. Two plots of x_n . On the left, we have taken $r = 0.9$ and $x_0 = 0.5$. We see that the map quickly tends to the origin. On the right, we have taken $r = 2.9$, close to the end of the window of stability, and $x_0 = 0.1$. The map now tends to x^* , oscillating about it as it goes.

The logistic map is a discrete version of the logistic equation that we studied in section 1.1. However, after nondimensionalisation, the logistic equation has *no* free parameters. That's not true of the logistic map (2.18), which crucially depends on the parameter r which can't be absorbed into rescaling time because that's now a discrete variable. Indeed, if you discretise the logistic equation and rescale then you will end up with the logistic map (2.18) with $r = 1 + \epsilon^2$ where $\epsilon \ll 1$ is related to the small time interval that you choose. We will soon see that, for this value of r , the logistic map does indeed recover the qualitative behaviour of the logistic equation. But, for other values of r , wildly different things can happen.

2.2.1 The Fixed Points

The logistic map has two fixed points: one at $x = 0$ and the other at

$$x^* = 1 - \frac{1}{r} . \quad (2.20)$$

This is only a fixed point for $r > 1$.

What is the stability of these two fixed points? First we can look near the origin where we write $x_n = \epsilon_n \ll 1$. We then have

$$x_{n+1} = r\epsilon_n(1 - \epsilon_n) = r\epsilon_n + \mathcal{O}(\epsilon_n^2) . \quad (2.21)$$

So the origin is a stable fixed point for $r < 1$, as each successive iteration takes us closer to it. It is unstable for $r > 1$.




Figure 30. Cobweb diagrams, made with the same parameters as Figure 29. On the left, $r < 1$ and on the right $1 < r < 3$.

Meanwhile, for the non-trivial fixed point (2.20), we write $x_n = x^* + \epsilon_n$ and look at

$$x_{n+1} = x^* + f'(x^*)\epsilon_n + \mathcal{O}(\epsilon_n^2) \quad \text{with} \quad f'(x^*) = 2 - r. \quad (2.22)$$

This fixed point is stable for $1 < r < 3$ and is unstable for $r > 3$. Moreover, in the stable regime $2 < r < 3$, we see that $f'(x^*)$ is negative and so successive terms will jump either side of x^* , while honing in to the fixed point.

We can see this analysis bearing out by simply plotting successive iterations of the logistic map. This is shown in Figure 29 for $r < 1$ and for $2 < r < 3$, where we see that the results converge to the origin and to x^* respectively.

Cobweb Diagrams

There's a nice graphical way to see how the map behaves. We plot $y = f(x)$ together with the line $y = x$. Start at some value of x and $y = 0$ and move up until you hit the graph $f(x)$. Then it's simple to convince yourself that successive iterations of the map are implemented by bouncing at right angles between the graph and the line.

For $r < 1$, the line sits above the graph $f(x)$ and the bouncing takes you down to the origin. This is shown in on the left of Figure 30. For $r > 1$, the line leaves the origin below the graph $f(x)$ and intersects it again at x^* . For $1 < r < 3$, the bouncing zooms in to the fixed point as shown on the right of Figure 30.

2.2.2 Bifurcation

We've understood the behaviour of the logistic map for $r < 3$. But what happens for $r > 3$ when the fixed point x^* is unstable? We can easily answer this by looking at

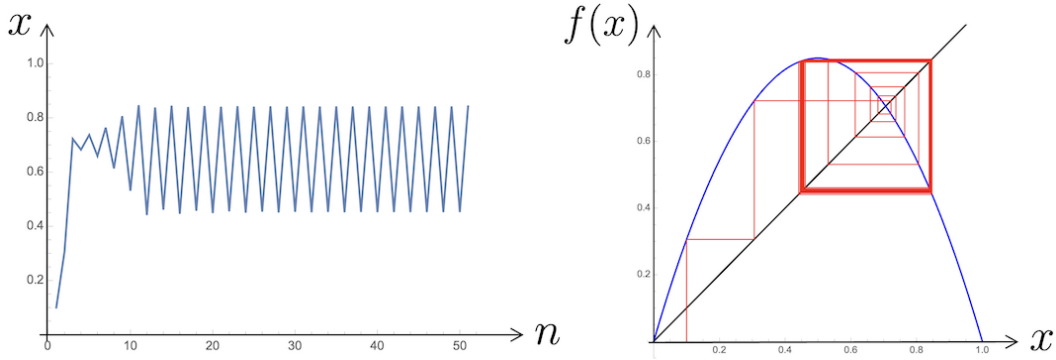


Figure 31. The logistic map for $r = 3.4$ rapidly reaches a stable 2-cycle, bouncing between two different points.

some numerics. Both the behaviour of x_n and the cobweb diagram are shown in Figure 31 for $r = 3.4$. We see that the map starts by honing in on the fixed point x^* , but quickly realises that this is unstable and settles down to a periodic pattern, bouncing between two different points. On the cobweb diagram, the trajectory repeatedly traces out the rectangle as shown in the figure.

We will call a trajectory that bounces between p different points a p -cycle. What we have on our hands in Figure 31 is a 2-cycle.

We can understand this behaviour analytically. We look at the map

$$\begin{aligned}
 f^2(x) &= f(f(x)) \\
 &= rf(x)(1 - f(x)) \\
 &= r^2x(1 - x)(x - rx(1 - x)) .
 \end{aligned} \tag{2.23}$$

This has fixed points

$$x = f^2(x) \implies x(1 - r + rx)(1 + r - r(1 + r)x + r^2x^2) = 0. \tag{2.24}$$

The first two factors give us our previous fixed points, $x = 0$ and $x = x^*$. Now, however, we see that f_2 has two further fixed points, where the second factor vanishes. These are given by

$$x_{\pm} = \frac{1}{2r} \left(1 + r \pm \sqrt{(r - 3)(r + 1)} \right) . \tag{2.25}$$

We see that these fixed points are only real when $r > 3$, which is what we wanted. These are the two points that the system bounces between, as evident in the numerical solution of Figure 31.

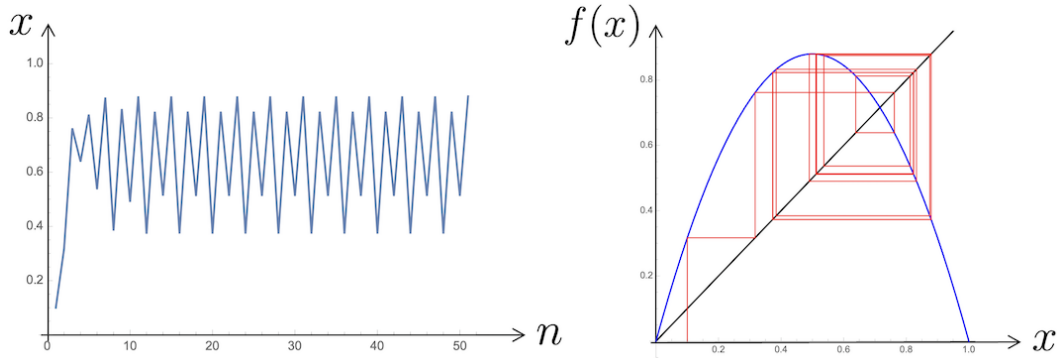


Figure 32. The logistic map for $r = 3.52$ reaches a stable 4-cycle, bouncing between four different points.

Next we should ask: are these fixed points of $f^2(x)$ stable? Here, the standard perturbation analysis prompts us to look at the derivative $f^{2'}(x_{\pm})$ and see if its modulus is bigger than one (in which case the fixed points are unstable) or less than one (in which case they are stable). We have

$$\frac{df^2}{dx} = \frac{d}{dx}f(f(x)) = f'(f(x))f'(x) . \quad (2.26)$$

If we put $x = x_+$ then $f(x_+) = x_-$ and so, rather cutely, we have

$$\left. \frac{df^2}{dx} \right|_{x_+} = f'(x_-)f'(x_+) = r^2(1 - 2x_-)(1 - 2x_+) = -r^2 + 2r + 4 . \quad (2.27)$$

You can check that, for r slightly greater than 3, the modulus of the right-hand side is less than one and so the fixed points of $f^2(x)$ are stable. This is the behaviour that we see in the figure. But the fixed points turn unstable when

$$-r^2 + 2r + 4 < -1 \implies r > 1 + \sqrt{6} \approx 3.449 \quad (2.28)$$

where we've taken the positive root of the quadratic.

The net result is that the 2-cycle behaviour seen in Figure 31 only holds for the regime $3 < r \lesssim 3.45$. Which begs the question: what happens for greater values of r ?

Again, we can plot things numerically to get an idea. This is shown in Figure 32 for $r = 3.52$ where we see that now the value of x bounces periodically between four fixed points. Again, we could get an analytic handle on this by studying the fixed points of $f^4(x)$ although this will now be a polynomial of order 8 and somewhat harder to analyse. (Actually, after factoring out the fixed points of $f(x)$ and $f^2(x)$, you're left with a quartic to deal with.)




Figure 33. The long term behaviour of the logistic map as a function of r . We see the single stable fixed point for $1 < r < 3$ turning into the stable 2-cycle, then the stable 4-cycle, then the stable 2^n -cycle. For $r > r_\infty$, there is (typically, but not always) no stable cycles. But occasionally, out of the murk, windows of order appear.


This story now repeats. The 4-cycle seen in Figure 32 does not persist forever but, it turns out, becomes unstable for $r \gtrsim 3.55$, at which point we find ourselves with a stable 8-cycle. And so on. As r increases, we find ourselves with the period of the cycle constantly doubling in size.

2.2.3 And Then... Chaos

The period of cycles keep doubling and the range over which this happens gets shorter and shorter. Until we get to approximately

$$r > r_\infty = 3.5699 . \quad (2.29)$$

At this point, all hell breaks loose. For most (but not all!) values of r above this value, there is no discernible pattern in the long term behaviour, which bounces around seemingly at random. The plot for $r = 3.7$ is shown in the figure to the right.






Figure 34. Plots of $y = x$, $y = f(x)$ and $y = f^3(x)$. On the left, $r = 3.8$ and on the right, $r = 3.835$. In the latter case, there are solutions to $x = f^3(x)$, reflecting the emergence of a stable 3-cycle.

We can get more insight into this by staring (for a long time!) at the iconic plot shown in Figure 33. This depicts the long-time behaviour of the logistic map for different values of r . For $1 < r < 3$, there is the single fixed point that we found above. At $r \approx 3.45$, this bifurcates into the period 2-cycle and this later bifurcates into the period 4-cycle and so on. For $r > r_\infty$, we see that a seeming continuum of values of x appear. This is the regime of chaos.

But, perhaps most surprisingly, the chaos does not persist for all values of $r > r_\infty$. Occasionally, we find windows of order. These are the white stripes that are apparent in Figure 33 where we again find cycles of some order.

The most prominent of these is the white stripe around $r = 3.88$. Towards the left of this stripe, you can see three points in the diagram. This reflects the fact that, in a small window, there is a stable 3-cycle. This is something new, as all previous p -cycles had p a power of 2.

We could try to repeat the story above to understand how this 3-cycle comes about, looking for fixed points of $f^3(x)$. But, as for all higher powers, this is tricky because it's a higher order polynomial and we don't have analytic expressions for its values. But we can plot the functions $y = x$, $y = f(x)$ and $y = f^3(x)$ to see where the fixed points lie. These functions are plotted in Figure 34 for $r = 3.8$ (on the left) where there is no 3-cycle, and for $r = 3.835$ (on the right) which is in the window with the stable 3-cycle. We can see that, as r changes, the dips in $f^3(x)$ come down sufficiently to intersect the straight line, revealing the fixed points. In fact, the straight line intersects the curve $f^3(x)$ twice each time: these are denoted with red and white dots. The stable 3-cycle

corresponds to intersections where $f^3(x)$ is shallower and these are coloured red. The white dots are then an accompanying unstable 3-cycle.

The period doubling that we saw previously also takes place within these windows of order. For example, the 3-cycle turns into a 6-cycle, which then turns into a 12-cycle, and so on.

2.2.4 The Logistic Map in Ecology

The logistic map is certainly very pretty. But, so far, we haven't really explained why it might be useful in modelling population dynamics.

The potential utility of the logistic map, with its panoply of different behaviours, was advertised in a famous paper by the mathematical ecologist Bob May⁸. The first compelling evidence that period bifurcations and the ensuing chaos are at play in the world of population dynamics was presented in a study from the 1990s on flour beetles⁹. The population at discrete time n consists of larvae L_n , pupae P_n and adults A_n and can be modelled as

$$\begin{aligned} L_{n+1} &= bA_n e^{-c_1 A_n - c_2 L_n} \\ P_{n+1} &= r_L L_n \\ A_{n+1} &= P_n e^{-c_3 A_n} + r_A A_n . \end{aligned} \tag{2.30}$$

Here the constants b , r_L and r_A capture reproductive and death rates while the constants c_1 , c_2 and c_3 are more gruesome, describing the cannibalistic tendencies of flour beetles. This set of equations exhibits many of the features of the logistic map, including a series of period bifurcations before descending into chaos as the parameters are varied. The rates r_L and r_A were artificially varied in the experiment by simply removing individuals from the population. There is then a feedback effect where the constants c_i , with $i = 1, 2, 3$, depend on r_L and r_A . In an experiment that took place over many years, the authors observed both period doubling and chaotic behaviour in the population.

The current consensus is that chaos is possible, but not common, in nature. However, this has been challenged recently in work that suggests chaos may not be so rare after

⁸The paper is “[Simple mathematical models with very complicated dynamics](#)”, Nature 261 (1976) and makes for an easy and fun read.

⁹The original paper is R.F. Constantino et al, “[Chaotic Dynamics in an Insect Population](#)”, Science, vol 275 (1997).

n	0	1	2	3	4	5	6
r_n	3	3.44948974	3.54409035	3.56440726	3.56875941	3.56969160	3.56989125
δ_n		4.751	4.656	4.668	4.668	4.669	4.669

Table 1. Numerical values for r_n and δ_n

all¹⁰. It’s hard to know for sure because it is challenging to collect the long time series needed to estimate Lyapunov exponents and other tell tale signs of chaos.

There is something a little unnerving in finding chaos in population dynamics. One might, quite reasonably, think that wild swings in population size are due to some extreme event, say weather or famine. Moreover, if we can learn to control such events then we could restore order to the universe. But chaotic systems exhibit wild swings for no underlying reason other than the inherent dynamics itself. Chaos is a control freak’s worst nightmare.

2.3 Universality

There’s something magical lurking in the discussion above. As we increase r , we get a series of period doublings. We can ask: for what values of r does the period double?

We calculated the first two of these above. The single fixed point becomes unstable and bifurcates into a 2-cycle at $r = r_0$ with

$$r_0 = 3 . \tag{2.31}$$

The 2-cycle then bifurcates further into a 4-cycle at

$$r_1 = 1 + \sqrt{6} \approx 3.44948974 . \tag{2.32}$$

The value of r_n for which the 2^n cycle bifurcates into a 2^{n+1} -cycle can be calculated numerically and is shown in Table 2.3.

With these values in hand, we see that there is a pattern. We define the ratio of differences

$$\delta_n = \frac{r_n - r_{n-1}}{r_{n+1} - r_n} . \tag{2.33}$$

¹⁰See T. Rogers, B. Johnson, and S. Munch, “[Chaos is not rare in natural ecosystems](#)”, Nature Ecology & Evolution, 6, (2022).




Figure 35. The heights of successive bifurcations, measured relative to the line $x = 1/2$, are d_n .

These too are given in Table 2.3 where they are clearly converging to

$$\delta = \lim_{n \rightarrow \infty} \delta_n = 4.669 \dots \quad (2.34)$$

The magic is that this number appears in many other maps. For example, if you consider the map

$$f(x) = r \sin(\pi x) \quad (2.35)$$

then, as you vary $r \in [0, 1]$, you will again see period doubling at a rate that converges towards the same value of δ . What we're seeing here is that, hidden within the logistic map, is a new mathematical constant, δ . This is known as the *Feigenbaum constant*. Or, more precisely, it is one of two Feigenbaum constants.

The other Feigenbaum constant comes from noting that the heights of successive bifurcations get smaller in Figure 33. We would like to find a way to characterise their height. We do this, by measuring the height relative to the line $x = 1/2$ (corresponding to the maximum of the function $f(x)$). We call successive heights d_n , as shown in the Figure 35. (As an aside, the kind of bifurcating diagrams shown in Figure 35 are affectionately known as fig tree diagrams, not because they look particularly like fig trees but because this is a direct translation of the German word “Feigenbaum”.)

The second Feigenbaum constant comes from noting that the ratio of heights also converges to

$$\alpha = \lim_{n \rightarrow \infty} \frac{d_n}{d_{n+1}} = -2.5029 \dots \quad (2.36)$$

Here the minus sign reflects the fact that, as shown above, the heights are measured alternately above and below the line $x = 1/2$. The value of α is another universal constant, in the sense that the same number emerges for many different maps.




Figure 36. Plots of $y = x$, $f(x)$ and $f^2(x)$ for $r = 2.9$ (on the left) and for $r = 3.2$ (on the right).

The Feigenbaum constants α and δ arise for many maps, but not for all maps. The maps $f(x)$ in question should be smooth and “unimodal”, which means that they go up, then down with a single maximum. Moreover, the maximum should be quadratic. Any map with these features will exhibit period doubling with the Feigenbaum constants α and δ .

Of course, that begs the question: what about other maps? Here too there is an interesting story. Suppose, for example, that we consider maps that have a quartic, rather than quadratic, maximum. Then you’ll find the same kind of period doubling but with different constants, $\alpha \approx -1.7$ and $\delta \approx 7.3$.

The Feigenbaum constants are reminiscent of other mathematical constants, known as *critical exponents*, that arise in the theory of phase transitions. You can read more about them in the lectures on [Statistical Field Theory](#). Indeed, we will borrow some ideas of “renormalisation” from the theory of phase transitions below when we describe how to calculate the Feigenbaum constants.

2.3.1 Zooming in on Bifurcations

Let’s first get a sense for why this universality might be happening. We can go to the very first bifurcation, in which the stable fixed point changes to the 2-cycle.

We can get some understanding of what’s going on by staring at Figure 36 where we plot $y = x$, $y = f(x)$ and $y = f^2(x)$. On the left, these are plotted for $r = 2.9$ where the only solutions to $x = f^2(x)$ are also solutions to $x = f(x)$. On the right, we plot these same functions for $r = 3.2$. Now we see that there are additional solutions to $x = f^2(x)$ that are *not* solutions to $x = f(x)$. This transition happens at $r = 3$, and this is where the single stable fixed point of $f(x)$ becomes unstable and is replaced by the stable 2-cycle, corresponding to the new fixed points of $f^2(x)$.





Figure 37. Plots of $y = x$, $f(x)$, $f^2(x)$ and, $f^4(x)$. This is plotted for $r = 3.5$, after the 2-cycle becomes unstable and there are new fixed points of $f^4(x)$. These can be shown in the plot on the left and, more clearly, in the plot on the right where we've zoomed in to the relevant piece of the graph.

So far, so good. Now what happens when we increase r further so that this 2-cycle turns into a 4-cycle? This time, the fixed points of $f^2(x)$ are becoming unstable. We can see this happening in the same graphical manner as before, this time also plotting $f^4(x)$ to see how its fixed points emerge at some value of r . This is shown on the left of Figure 37 where it's all a little cluttered. But we can zoom in to the relevant part, as shown on the right of 37.

Now comes the key observation. Take the zoomed in plot in Figure 37. Flip it upside down and reflect it. It looks like this:



But that looks very much like the right-hand plot of Figure 36!

Now we can repeat this, looking at how the fixed points of $f^4(x)$ become unstable as new fixed points of $f^8(x)$ appear, and then how these become unstable as new fixed points of $f^{16}(x)$ appear, and so on. At each stage, we zoom in and flip and what we're left with is always a figure that looks like the one above. The existence of the two universal Feigenbaum constants suggests that, as we do this procedure over and over again, we might converge on some universal function. Our task is to try to understand some properties of this function and to extract the Feigenbaum constants from it.

2.3.2 Renormalisation

We want to put the ideas above on a firmer footing. How do we implement the iterated “taking $f^2(x)$, zooming and flipping” procedure described above in more concrete terms?

There are a number of steps that we need to take. First, we will describe the class of functions that we care about. It’s convenient to rescale things slightly. We will consider regular, unimodal functions such that

- $x_{n+1} = f(x_n)$ with $x \in [-1, +1]$.
- The map is symmetric about the maximum at $x = 0$, with $f(0) = 1$
- The maximum is quadratic.

For example, after some rescaling the logistic map can be rewritten as

$$x_{n+1} = 1 - rx_n^2 . \tag{2.37}$$

This map exhibits all the universal features described above, including the two Feigenbaum constants α and δ . But everything that we say below holds for any map obeying the three criteria above.

Part of our iteration in going from $f(x)$ to $f^2(x)$ involves increasing the parameter r , so that we get to the point where $f^{2^n}(x)$ develops fixed points. For example, Figure 36 is plotted with $r = 2.9$ while Figure 37 is plotted with $r = 3.5$. Let’s first look more carefully at how we do this.

A map $x \rightarrow f(x)$ has a fixed point $x^* = f(x^*)$. This fixed point is stable if $f'(x^*) < 1$ and is unstable if $f'(x^*) > 1$. The map fixed point is said to be *superstable* if

$$f'(x^*) = 0 . \tag{2.38}$$

This is the most stable that a fixed point can be. In this case, the convergence towards the fixed point is typically exponential rather than power-law.

It’s clear that a fixed point of our class of maps is superstable if x^* coincides with the maximum of the map x_{\max} which, by construction, we’ve taken to be $x_{\max} = 0$. Similarly, for the higher maps $f^2(x)$ with two new fixed points x_+ and x_- , we saw in (2.27) that $df^2/dx(x_+) = f'(x_+)f'(x_-)$, so this fixed point is superstable if either x_+ or x_- coincides with x_{\max} .

We will denote the value of r at which the fixed point of $f^{n-1}(x)$ is superstable as R_n . These are shown in Figure 38. Note that $r_n < R_n < r_{n+1}$. It turns out that the superstable points R_n converge in the same manner as the bifurcation points r_n ,

$$\delta = \lim_{n \rightarrow \infty} \frac{R_n - R_{n-1}}{R_{n+1} - R_n} = 4.669 \dots \quad (2.39)$$

In what follows, we will phrase everything in terms of maps evaluated at the superstable points R_n .

The Renormalisation Map

Now we can start to put our iteration process in place. We start with a map $f(x; R_0)$ which has a superstable fixed point. We then want to turn this into the appropriately zoomed and flipped map $f^2(x; R_1)$ which, importantly, also has a superstable fixed point. Moreover, we want the map $f^2(x)$ to fall into our general class of maps, obeying the various criteria listed above.

To achieve, this suppose that

$$f(1) = -a . \quad (2.40)$$

Then, because $f(0) = 1$, we have $f^2(0) = f(1) = -a$. So to keep ourselves within the class of maps with $f(0) = 1$, we should rescale f^2 . This is the zooming described above. We should also rescale x to ensure that the domain remains in $x \in [-1, +1]$, with a minus sign to give us the necessary reflection. The upshot is that the zooming and flipping procedure is described by the following action on a map

$$f(x; R_0) \mapsto -\frac{1}{a} f^2(-ax; R_1) . \quad (2.41)$$

Both the original map $f(x)$ and the new map $f^2(x)$ are evaluated at the appropriate point $r = R_n$ where they have a superstable fixed point. We say that the map $f(x; R_0)$ has been *renormalised*, terminology stolen from quantum field theory.

The renormalisation map (2.41) can itself be viewed as a dynamical system, but not one acting on a single variable x but now acting on a class of functions $f(x)$. Said differently, this is a dynamical system with an infinite number of degrees of freedom. Nonetheless, we could push on and think of it like any other dynamical system. We could, for example, reiterate the process n times

$$f(x, R_0) \mapsto \left(-\frac{1}{a}\right)^n f^{2^n}((-a)^n x; R_n) . \quad (2.42)$$




Figure 38. The bifurcations happen at points $r = r_n$; the fixed points are superstable at $r = R_n$.

Then we can ask: does this process converge? In other words, is there some universal function defined by

$$g(x) = \lim_{n \rightarrow \infty} \left(-\frac{1}{a} \right)^2 f^{2n}((-a)^n x; R_n) ? \quad (2.43)$$

It's not at all obvious that such a function $g(x)$ exists. But the universality observed in the Feigenbaum constants suggests that, under the right circumstances, it might. But what are these circumstances?

2.3.3 The Feigenbaum constant α

To get a sense for when universal function $g(x)$ in (2.43) might exist, we need to think more carefully about the meaning of that rescaling factor a . It is rescaling the x coordinate by a , but we expect from our previous discussion that this is what the Feigenbaum constant α is doing. This means that we should identify

$$\alpha = -\frac{1}{a} . \quad (2.44)$$

Moreover, it suggests that the limit (2.43) should only exist if we take a very specific value of α . We just need to compute this value.

Now we're on the home straight. If there's a universal function $g(x)$ that is the limit of the renormalisation map then, when we put it in the renormalisation map, nothing

should happen. In other words, $g(x)$ should be a fixed point of the renormalisation map and satisfy

$$g(x) = \alpha g^2\left(\frac{x}{\alpha}\right) . \quad (2.45)$$

This slightly strange, self-referential equation defines both the function $g(x)$ and the constant α . Indeed, as we've chosen $g(0) = 1$, we have

$$1 = \alpha g(1) \implies \alpha = \frac{1}{g(1)} . \quad (2.46)$$

To make progress, we can simply Taylor expand $g(x)$ around the origin. We know that it is a symmetric function with a quadratic maximum, so we can Taylor expand

$$g(x) = 1 + \sum_{n=1}^N c_{2n} x^{2n} . \quad (2.47)$$

Substitute this into (2.45) and compare various terms. The constant terms give

$$1 = \alpha(1 + c_2 + c_4 + \dots) \quad (2.48)$$

which just reiterates the result (2.46). The x^2 terms give

$$\alpha = 2c_2 + 4c_4 + \dots \dots \quad (2.49)$$

The x^{2n} terms are polynomials of degree $2n$ in the variables α and c_{2n} . The upshot is that we have $N + 1$ equations in $N + 1$ variables which we can solve numerically. (Because these are higher order polynomials, there are several solutions and you have to make sure that you get the right one.) The higher the value of N , the better the accuracy. For low values we have:

$$\begin{aligned} N = 2 &\implies c_2 \approx -1.52 , c_4 \approx 0.13 , \alpha \approx -2.53 \\ N = 3 &\implies c_2 \approx -1.52 , c_4 \approx 0.073 , c_6 \approx 0.046 , \alpha \approx -2.479 . \end{aligned} \quad (2.50)$$

By the time you get to $N = 6$, you have $\alpha \approx -2.502897$, which is correct to 1 part in 10^6 .

You can see that our ansatz (2.47) assumes that the maximum is quadratic. You could repeat the calculation setting $c_2 = 0$ so that the maximum is quartic. Then you get the different universal constant $\alpha \approx -1.7$ appropriate for such quartic maps.

2.3.4 The Feigenbaum Constant δ

Computing the other Feigenbaum constant δ is a little more involved. We will be somewhat heuristic in what follows, but still put together enough of the story to allow us to compute δ .

To start, we define the renormalisation map to act on any function $\phi(x)$ as

$$\text{Ren}[\phi(x)] = \alpha\phi^2\left(\frac{x}{\alpha}\right) . \quad (2.51)$$

This is the same as our previous map (2.41), except we're not changing the variable r in any way: just iterating the map and rescaling. We take the constant α to be the Feigenbaum constant.

Now consider the one-parameter family of functions $f(x; r)$ that defines our original map. We know that there are special values of $r = R_n$ where this map has a superstable 2^n -cycle. The claim of universality is that they converge as

$$R_n = R_\infty - \frac{A}{\delta^n} \quad (2.52)$$

for some (non-universal) constant A and with δ the Feigenbaum constant (2.34). This is what we would like to show. We will first need to develop some machinery to do this.

Expanding About the Universal Function

The renormalisation map $\text{Ren}[\phi(x)]$ has a “fixed point”, or more precisely a “fixed function”, which is our universal function $g(x)$,

$$\text{Ren}[g(x)] = \alpha g^2\left(\frac{x}{\alpha}\right) = g(x) . \quad (2.53)$$

Usually, when presented with a dynamical system with a fixed point, our first inclination is to linearise around the fixed point to see what happens in its vicinity. The same is true here. We look at functions $\phi(x)$ that are close to the universal function, with

$$\phi(x) = g(x) + \epsilon\theta(x) \quad (2.54)$$

with $\epsilon \ll 1$ and $\theta(x)$ some other arbitrary function. Now we act with the renormalisation map of $f(x; r)$. The map isn't linear so we have to tread slowly. We have

$$\begin{aligned} \text{Ren}[\phi(x)] &= \alpha\phi\left(\phi\left(\frac{x}{\alpha}\right)\right) \\ &= \alpha\phi\left(g\left(\frac{x}{\alpha}\right) + \epsilon\theta\left(\frac{x}{\alpha}\right)\right) \\ &= \alpha\left[\phi\left(g\left(\frac{x}{\alpha}\right)\right) + \epsilon\phi'\left(g\left(\frac{x}{\alpha}\right)\right)\theta\left(\frac{x}{\alpha}\right)\right] \\ &= \alpha g\left(g\left(\frac{x}{\alpha}\right)\right) + \epsilon\alpha\left[\theta\left(g\left(\frac{x}{\alpha}\right)\right) + g'\left(g\left(\frac{x}{\alpha}\right)\right)\theta\left(\frac{x}{\alpha}\right)\right] \end{aligned} \quad (2.55)$$

where, in the third and fourth lines we've Taylor expanded and dropped terms of order ϵ^2 .

When you Taylor expand a function, the first correction is the derivative. Here we've Taylor expanded a *functional* $\text{Ren}[\phi(x)]$, which is a function of a function. The term multiplying the ϵ parameter should be thought of as the functional derivative of $\text{Ren}[\phi(x)]$. In fancy maths words, it's called the Fréchet derivative. We define

$$\text{DRen}_g[\theta(x)] = \alpha \left[\theta \left(g \left(\frac{x}{\alpha} \right) \right) + g' \left(g \left(\frac{x}{\alpha} \right) \right) \theta \left(\frac{x}{\alpha} \right) \right] . \quad (2.56)$$

We can then write

$$\text{Ren}[\phi(x)] = g(x) + \epsilon \text{DRen}_g[\theta(x)] . \quad (2.57)$$

where we've used the fact that $g(x)$ is a fixed point to get the first term.

How should we think of this? If this were a dynamical system with a finite number of degrees of freedom then, when we linearise around a fixed point, we get a matrix. And to understand the behaviour of the fixed point we have to look at the eigenvectors and eigenvalues of that matrix. The same is true here, except that we should look for eigenfunctions, $\theta_i(x)$ of the weird operator DRen . These obey

$$\text{DRen}_g[\theta_i(x)] = \lambda_i \theta_i(x) . \quad (2.58)$$

Here λ_i is the corresponding eigenvalue and i is an (infinite) index that labels the different eigenthings.

Our next step would be to solve for the different eigenvalues using (2.58). And we will, eventually, do this. The trouble is that this is a hard equation to solve and it's useful to have some motivation for doing so! That's where we're going next. We will argue that, out of the infinite number of eigenvalues, there is just single one λ_0 that has $|\lambda_0| > 1$. And, rather wonderfully, this eigenvalue coincides with the Feigenbaum constant δ .

Expanding About the Edge of Chaos

To extract the Feigenbaum constant, we need to look more carefully at the original map $f(x; r)$ and, in particular, this map evaluated at the superstable points $r = R_n$.

To start, we look at the map when it sits at the edge of chaos at $r = R_\infty$. We define

$$F(x) = f(x; R_\infty) . \quad (2.59)$$

At this point, we need to make an assumption: we assume that if you act with successive iterations of the map Ren on the function $F(x)$, then you will quickly converge towards the universal function $g(x)$ defined in (2.45),

$$\text{Ren}^n[F](x) \approx g(x) \quad \text{for } n \text{ suitable large.} \quad (2.60)$$

Said in more sophisticated language, we assume that $F(x)$ lies on the stable manifold of the fixed point $g(x)$. This assumption is, it turns out, true, but we will not prove it here. It seems plausible because the renormalisation map Ren differs from our original renormalisation by not changing the value of r , but we've already tuned the value of r to its final resting place R_∞ when considering $F(x)$.

Next we ask: what if we act with the renormalisation map on $f(x; r)$ with r close to R_∞ ? We write

$$f(x; r) = F(x) + \epsilon \left. \frac{\partial f(x; r)}{\partial r} \right|_{r=R_\infty} \quad \text{with } \epsilon = r - R_\infty . \quad (2.61)$$

Here we've dropped terms of order ϵ^2 and higher. Acting with the renormalisation map on $f(x; r)$ is just a matter of repeating the calculation (2.55), with $g(x)$ replaced by $F(x)$ and $\theta(x)$ replaced by $\partial f/\partial r$. We can express the result in terms of our operator DRen defined in (2.56),

$$\text{Ren}[f(x; r)] = \alpha F \left(F \left(\frac{x}{\alpha} \right) \right) + \epsilon \text{DRen}_F \left[\frac{\partial f(x; r)}{\partial r} \right] . \quad (2.62)$$

Now we act with successive renormalisation maps on this function. We know that $F(x)$ tends towards the universal function $g(x)$, as in (2.60). Acting successively on the function $f(x; r)$, with r close to R_∞ then gives

$$\text{Ren}^n[f(x; r)] \approx g(x) + (r - R_\infty) \text{DRen}_g^n \left[\frac{\partial f(x; r)}{\partial r} \right] \quad \text{for suitable large } n. \quad (2.63)$$

Admittedly, that \approx sign is doing some heavy lifting here. We've taken n iterations of the map DRen_g , each of them evaluated around the function $g(x)$ rather than the function $F(x)$ or some intermediate. That should really be justified. But we're not going to.

To proceed, we expand the function $\partial f/\partial r$ in terms of eigenfunctions of the operator DRen . We write

$$\frac{\partial f(x; r)}{\partial r} = \sum_i a_i \theta_i(x) . \quad (2.64)$$

Clearly there's yet another assumption here that the $\theta_i(x)$ form a complete basis of functions. You may have guessed by now that it's not an assumption we're going to justify. If we substitute this ansatz into the result (2.63), we have

$$\text{Ren}^n[f(x; r)] \approx g(x) + (r - R_\infty) \sum_i \lambda_i^n a_i \theta_i(x) . \quad (2.65)$$

Now we're in good shape to make the final argument.

The Feigenbaum Constant is an Eigenvalue

We make the following claim:

Claim: If there is just a single eigenvalue λ_0 with $|\lambda_0| > 1$, then

$$\lambda_0 = \delta \quad (2.66)$$

with δ the Feigenbaum constant.

Proof: The key idea is to get different expressions for the renormalisation $\text{Ren}^n[f(x; r)]$. First, if the assumption is correct, and all eigenvalues other than λ_0 have modulus $|\lambda_i| < 1$, then the iterations in (2.65) quickly kill all but the $\theta_0(x)$ eigenfunction,

$$\text{Ren}^n[f(x; r)] \approx g(x) + (r - R_\infty) \lambda_0^n a_0 \theta_0(x) . \quad (2.67)$$

Now we think about this result applied to the case with $r = R_n$. The nice thing about the function $f(x; R_n)$ is that it has a superstable n -cycle and, moreover, we know that $x = 0$ is one of the points on this n cycle. In other words, if we act with $f(x; R_n)$ a total of 2^n times, starting at $x = 0$, then we get back to $x = 0$,

$$f^{2^n}(x = 0; R_n) = 0 . \quad (2.68)$$

But this means that the left-hand side of (2.67) vanishes, with the extra scaling by α in the renormalisation group map unimportant because $0/\alpha = 0\alpha = 0$. So we have

$$R_n - R_\infty = -\frac{g(0)}{a_0 \theta_0(0)} \frac{1}{\lambda_0^n} = \frac{\text{constant}}{\lambda_0^n} . \quad (2.69)$$

But this is precisely the geometric progression (2.52) seen in the bifurcations, with $\lambda_0 = \delta$. □

It remains to find the eigenvalue $\lambda_0 = \delta$. For this, we need to solve the eigenfunction equation (2.58)

$$\text{DRen}_g[\theta(x)] = \alpha \left[\theta \left(g \left(\frac{x}{\alpha} \right) \right) + g' \left(g \left(\frac{x}{\alpha} \right) \right) \theta \left(\frac{x}{\alpha} \right) \right] = \delta \theta(x) . \quad (2.70)$$

This too can be found using a power series ansatz for $\theta(x)$, together with our previous expansion for the universal function $g(x)$. Expanding to order $N = 6$ is sufficient to give $\delta \approx 4.66914$, accurate to one part in 10^5 .