Chapter 1

What is a mathematical model?

Learning outcomes

By the end of this chapter you should be able to:

- describe the modelling cycle and apply it to a simple real-world situation;
- identify the dimensions of physical quantities and check dimensional homogeneity;
- make and justify simplifying assumptions when building a model;
- derive and solve first-order ordinary differential equations (ODEs) arising from rate laws and conservation principles;
- formulate and solve simple difference equations (recurrence relations) and compare discrete-time models with their continuous counterparts;
- interpret parameters in context, assess a model's limitations, and suggest refinements.

1.1 The modelling cycle

Mathematical modelling is the process of translating a real-world phenomenon into the language of mathematics, analysing the result, and then translating back to make predictions or gain understanding.

Definition: Mathematical model

A **mathematical model** is a set of equations, together with assumptions and parameter values, that describes or approximates the behaviour of a real-world system.

No model is a perfect description of reality, and that is fine. A model is a tool, judged by whether it is *useful* for the question at hand. As the statistician George Box put it: *"All models are wrong, but some are useful."*

1.1.1 The steps of the modelling cycle

Building and using a model typically follows a cycle:

1. **Observe**: identify the phenomenon and the question you want to answer.
2. **Simplify**: make assumptions to reduce the problem to something tractable.

3. **Formulate**: express the assumptions as equations.
4. **Solve**: find solutions analytically or numerically.
5. **Interpret**: translate the mathematical output back into the real-world context.
6. **Validate**: compare predictions with data or known behaviour.
7. **Refine**: relax assumptions, add complexity, and repeat.

We will follow this cycle throughout the course, starting with very simple models and building up to richer ones.

1.1.2 The cycle in action: how fast does a rumour spread?

To see what the modelling cycle looks like in practice, consider this question: *a rumour starts with one person in a town of 10 000, how long before half the town has heard it?*

1. Observe. The rumour spreads by word of mouth. People who know the rumour tell people who don't. Eventually, most people have heard it and there are fewer people left to tell, so the spread slows down.

2. Simplify. Let's make some assumptions:

- The population is fixed at $P = 10\,000$ (no one arrives or leaves).
- Everyone mixes equally (no social clusters or geography).
- Each person who knows the rumour tells it to others at a constant rate.
- Once you hear the rumour, you remember it and start spreading it.

3. Formulate. Let $N(t)$ be the number of people who have heard the rumour at time t (in days). The rate of new hearers depends on two things: how many spreaders there are (N) and how many listeners remain ($P - N$). Each encounter between a spreader and a non-spreader has some probability of transmitting the rumour. The simplest model is

$$\frac{\mathrm{d}N}{\mathrm{d}t} = \beta\, N\,(P - N),$$

where $\beta > 0$ is a rate constant (units: per person per day). This is the **logistic equation**; we will study it in detail later in the course.

4. Solve. We will learn how to solve this equation (and others with a similar structure) later in the course. For now we note that the solution is an S-shaped curve: a slow initial spread, followed by a rapid middle phase, and then *saturation* as $N \to P$.

5. Interpret. If $\beta = 0.0002$ per person per day and $N(0) = 1$, the model predicts that half the town ($N = 5000$) has heard the rumour after about 4.6 days.

6. Validate. We could compare with data on how information actually spreads (e.g. from public health campaigns). The S-shaped pattern is widely observed.

7. Refine. Our assumptions were quite crude. Possible refinements include:

- People forget the rumour, or stop spreading it after a while.
- The population has social structure (friend groups, workplaces).
- Some people are more connected than others.

Each refinement adds complexity but may improve accuracy for a particular question.

Key idea

The modelling cycle is *iterative*: you rarely get it right the first time. Start with the simplest model that captures the essential behaviour, check whether it is good enough, and refine if necessary. The first model is typically not the final answer!

1.2 The art of simplification

The hardest (and most creative) step in modelling is deciding what to include and what to leave out. Every assumption you make simplifies the mathematics but moves the model further from reality. The goal is to find the right balance: simple enough to solve, rich enough to capture the behaviour you care about.

1.2.1 What makes a good assumption?

A good simplifying assumption is one that:

- makes the problem tractable,
- can be stated precisely ("friction is small" is vague; "the friction force is negligible compared to gravity" is testable),
- can be checked after the fact.

1.2.2 Common simplifications

Many models in this course will use one or more of the following:

- **Constant parameters.** Birth rates, drag coefficients, thermal conductivities, etc. are assumed independent of time or state. This is often valid over limited ranges.
- **Homogeneity.** The system is well-mixed (no spatial variation). Valid when mixing is fast compared to the process of interest.
- **Neglecting small effects.** Air resistance on a falling ball, heat loss from a well-insulated container, evaporation from a deep tank, if an effect contributes less than a few percent we would typically ignore it initially.
- **Linearity.** Assume relationships are proportional (e.g. Hooke's law, Newton's cooling). Valid for small deviations from equilibrium but breaks down for large perturbations.
- **Continuity.** Treat discrete quantities (individuals in a population, molecules in a gas) as continuous variables. Valid when the numbers are large.

1.2.3 How to check your assumptions

After solving the model, go back and check:

1. **Self-consistency.** If you assumed "the population is large," does the solution ever predict $N < 10$? If so, the assumption breaks down in that regime.
2. **Sensitivity.** Change a parameter by 10%. Does the answer change by 1% (robust) or by 200% (sensitive)? If the answer is sensitive to a parameter you estimated roughly, the prediction is unreliable.

3. **Limiting cases.** Does the model reduce to something known in extreme cases? For example, does a model of population growth with limited resources reduce to exponential growth when resources are abundant? These sanity checks are powerful.

Key idea

Modelling is not about being "right", it is about being *useful*. A simple model that captures the dominant effect and whose limitations you understand is more valuable than a complicated model you cannot solve or interpret.

1.3 Units and dimensional homogeneity

Every physical quantity has **dimensions**, the type of measurement it represents, and is expressed in particular **units**. The most common dimensions are:

Dimension	Symbol	SI unit
Length	L	metre (m)
Mass	M	kilogram (kg)
Time	T	second (s)
Temperature	Θ	kelvin (K)

We write $[Q]$ for the dimensions of a quantity Q. For example, [speed] $= L\,T^{-1}$, i.e. metres per second. Derived quantities have dimensions that are products of powers of the base dimensions: [force] $= M\,L\,T^{-2}$, [energy] $= M\,L^2\,T^{-2}$, etc.

Key idea

Dimensional homogeneity: every term in a valid equation must have the same dimensions. You can use this as a quick sanity check on any formula. If the dimensions do not match, the equation is certainly wrong. If they do match, the equation *might* be right (but could still have a wrong numerical factor).

Example Checking a formula

Newton's second law states $F = ma$. Check dimensional homogeneity.
$[F] = M\,L\,T^{-2}$ (force, in newtons), $[m] = M$, $[a] = L\,T^{-2}$. So $[ma] = M \cdot L\,T^{-2} = M\,L\,T^{-2}$. ✓

Example Catching an error

A student claims the distance fallen under gravity is $s = \frac{1}{2}g^2t$. Check the dimensions.
$[s] = L$, but $[g^2t] = (L\,T^{-2})^2 \cdot T = L^2\,T^{-3}$. These do not match, so the formula is wrong. The correct formula is $s = \frac{1}{2}gt^2$: check that $[gt^2] = L\,T^{-2} \cdot T^2 = L$. ✓

A note on differential equations

Most of the models in this chapter describe how a quantity changes over time, so they take the form of equations involving derivatives.

Definition: Ordinary differential equation

An **ordinary differential equation (ODE)** is an equation relating an unknown function $y(x)$ of a single independent variable x to one or more of its derivatives. The **order** of the ODE is the order of the highest derivative that appears. A **first-order** ODE involves only $\mathrm{d}y/\mathrm{d}x$, and can usually be written as

$$\frac{\mathrm{d}y}{\mathrm{d}x} = F(x, y).$$

A **solution** is a function $y(x)$ that satisfies the equation.

In our models the independent variable will usually be time t, as in $\mathrm{d}N/\mathrm{d}t = rN$. The word *ordinary* distinguishes these from *partial* differential equations, in which the unknown depends on several variables. We study ODEs systematically in Chapter 4. For now we need one technique for solving simple first-order ODEs, called **separation of variables**.

Definition: Separation of variables

If a first-order ODE can be written in the form

$$\frac{\mathrm{d}y}{\mathrm{d}x} = g(x)\,h(y),$$

collect all y-terms on the left and all x-terms on the right, then integrate both sides:

$$\int \frac{\mathrm{d}y}{h(y)} = \int g(x)\,\mathrm{d}x.$$

Adding an arbitrary constant c on one side and solving for y gives the **general solution**; a known value of y at a specific x determines c and gives a **particular solution** (see Section 4.1).

1.4 A first model: exponential growth

1.4.1 Deriving the model from a rate law

Suppose a population of bacteria doubles every hour under ideal conditions. Let $N(t)$ denote the population size at time t (measured in hours).

We follow the modelling cycle:

Observe. The population grows over time. We want to predict how many bacteria there will be at any future time.

Simplify. Assume: unlimited nutrients, no death, no spatial variation, and the population is large enough to treat as a continuous variable.

Formulate. The simplest assumption consistent with the observations is that the *rate of change* of the population is proportional to the current population:

$$\frac{\mathrm{d}N}{\mathrm{d}t} = r\,N,$$

where $r > 0$ is the **growth rate** (units: per hour, i.e. $[r] = T^{-1}$). This is a first-order ODE: it involves only the first derivative of N.

The reasoning is: each bacterium divides independently at the same average rate, so twice as many bacteria produce twice as many new bacteria per unit time.

1.4.2 Solving by separation of variables

Separate and integrate:

$$\frac{\mathrm{d}N}{N} = r\,\mathrm{d}t \implies \log N = r\,t + C \implies N(t) = N_0\,e^{rt},$$

where $N_0 = N(0)$ is the initial population.

Interpret. The population grows exponentially. The growth is slow when N is small (few bacteria to divide) and fast when N is large but always at the same *proportional* rate r.

Validate. Bacterial populations do grow exponentially in the early stages when nutrients are plentiful. The model fails when resources run out; this motivates the logistic model later in the course.

Note. In this course log will be used to denote the natural logarithm (i.e. base e).

Example Bacterial growth

Suppose $N_0 = 100$ and the population doubles every hour, so $N(1) = 200$. Find r. From $200 = 100\,e^{r\cdot 1}$ we get $r = \log 2 \approx 0.693$ per hour. The population at time t is $N(t) = 100\,e^{0.693\,t}$.

1.4.3 Interpreting the parameters

- N_0: the initial value, sets the scale of the population.
- $r > 0$: exponential growth; $r < 0$: exponential decay; $r = 0$: constant.
- **Doubling time**: the time t_d such that $N(t_d) = 2N_0$ is $t_d = \log 2/r$.
- **Half-life**: for decay ($r < 0$), the time to halve is $t_{1/2} = \log 2/|r|$.

The same equation $\dfrac{\mathrm{d}N}{\mathrm{d}t} = rN$ models radioactive decay (with $r < 0$), continuously-compounded interest, and Newton's law of cooling (with a shift). This universality is one reason ODEs are so powerful.

1.5 A second model: the draining tank

To practise the modelling cycle on a different kind of problem, consider a cylindrical tank of water with a hole in the bottom. *How long does it take to drain?*

1.5.1 Setting up the model

Observe. Water flows out through the hole under gravity. The flow is fast when the tank is full and slows as the water level drops.

Simplify. We assume:

- The tank is a vertical cylinder with constant cross-sectional area A.
- The hole at the bottom has area $a \ll A$.

- The water surface remains horizontal (no swirling or sloshing).
- The flow speed through the hole depends only on the current water depth.

Formulate. Let $h(t)$ be the depth of water at time t. We need two ingredients:

Conservation of volume. The volume of water in the tank is $V = A\,h$. The rate of change of volume equals the (negative) outflow rate Q:

$$A\,\frac{\mathrm{d}h}{\mathrm{d}t} = -Q. \tag{1.1}$$

Constitutive law (Torricelli's law). The speed of water exiting through a small hole at the bottom of a tank of depth h is approximately $v = \sqrt{2gh}$, where g is gravitational acceleration. This follows from energy conservation (Bernoulli's equation) and is named after Evangelista Torricelli (1608–1647). The outflow rate is then

$$Q = a\,v = a\sqrt{2gh}.$$

Substituting into (1.1):

$$\boxed{\frac{\mathrm{d}h}{\mathrm{d}t} = -\frac{a}{A}\sqrt{2g\,h}.} \tag{1.2}$$

Check dimensions. $[a/A] = 1$ (dimensionless ratio of areas), $[g\,h] = L\,T^{-2}\cdot L = L^2\,T^{-2}$, so $[\sqrt{2gh}] = L\,T^{-1}$. The right-hand side has dimensions $L\,T^{-1}$, which matches $[\mathrm{d}h/\mathrm{d}t]$. ✓

1.5.2 Solving the model

Equation (1.2) is separable. Let $c = (a/A)\sqrt{2g}$ for convenience:

$$\frac{\mathrm{d}h}{\mathrm{d}t} = -c\sqrt{h}.$$

Separate:

$$\frac{\mathrm{d}h}{\sqrt{h}} = -c\,\mathrm{d}t \qquad \Longrightarrow \qquad 2\sqrt{h} = -c\,t + C.$$

With initial condition $h(0) = h_0$, we find $C = 2\sqrt{h_0}$, giving

$$\boxed{h(t) = \left(\sqrt{h_0} - \frac{c\,t}{2}\right)^2 = h_0\left(1 - \frac{t}{t_{\mathrm{drain}}}\right)^2,} \tag{1.3}$$

valid for $0 \le t \le t_{\mathrm{drain}}$, where the **drain time** is

$$t_{\mathrm{drain}} = \frac{2\sqrt{h_0}}{c} = \frac{2A}{a}\sqrt{\frac{h_0}{2g}} = \frac{A}{a}\sqrt{\frac{2h_0}{g}}. \tag{1.4}$$

1.5.3 Interpreting the result

- The water depth decreases as a *quadratic* in t (not an exponential) the shape of the curve is different from the exponential models we saw earlier.
- The draining is fastest at the start (when h is large and the pressure drives fast flow) and slowest near the end (the last trickle takes a long time).
- The drain time scales as $\sqrt{h_0}$: doubling the initial depth does *not* double the drain time, it increases it by a factor of $\sqrt{2} \approx 1.41$.
- The ratio A/a appears naturally: a larger tank or a smaller hole means a longer drain time.

1.5.4 Validating and refining

Limiting cases. As $a \to 0$ (tiny hole), $t_{\text{drain}} \to \infty$, the tank takes forever to drain. As $a \to A$ (the entire bottom is open), $t_{\text{drain}} \to \sqrt{2h_0/g}$, which is the free-fall time for the water surface, both make physical sense.

What did we neglect? Viscosity (friction in the hole), surface tension, the shape of the hole, turbulence, and the fact that real tanks are rarely perfect cylinders. Torricelli's law also assumes the hole is small compared to the tank ($a \ll A$); if a is comparable to A, the analysis is more involved.

Comparison with experiment. In practice, Torricelli's law slightly overestimates the outflow speed because it ignores viscous losses. Engineers correct for this with a **discharge coefficient** $C_d < 1$, replacing $Q = a\sqrt{2gh}$ with $Q = C_d\, a\sqrt{2gh}$. The model structure is unchanged; only the parameter c is modified.

Example Numerical prediction

A cylindrical tank has radius 0.5 m and a circular drain hole of radius 1 cm. The initial water depth is 2 m. How long does it take to drain?
The areas are $A = \pi(0.5)^2 = 0.785$ m^2 and $a = \pi(0.01)^2 = 3.14 \times 10^{-4}$ m^2. With $g = 9.81$ m/s^2 and $h_0 = 2$ m:

$$t_{\text{drain}} = \frac{0.785}{3.14 \times 10^{-4}} \sqrt{\frac{2 \times 2}{9.81}} \approx 2500 \times 0.639 \approx 1597 \text{ s} \approx 27 \text{ minutes}.$$

Key idea

The draining tank illustrates several important modelling ideas: combining a conservation law with a constitutive relation, checking dimensions, identifying the key parameter groups (A/a and $\sqrt{h_0/g}$), and using limiting cases to validate. These are techniques you will use throughout the course.

1.6 Formulating models from scratch

The examples above followed a familiar pattern: rate of change equals something. Here is a general strategy for building ODE models:

1. **Choose your variable.** What quantity do you want to track? Population, temperature, depth, concentration, bank balance, ...
2. **Write a word equation.** Before any symbols, write the balance in words:

   $$\text{rate of change of [variable]} = \text{[gains]} - \text{[losses]}.$$

3. **Model each term.** For each gain or loss, decide: is it constant? Proportional to the variable? Proportional to some other quantity? Write each as a mathematical expression.
4. **Assemble and check dimensions.** Combine into an ODE and verify that every term has the same dimensions.
5. **Identify parameters and initial conditions.** List what you know and what you would need to measure.

Example Drug concentration in the bloodstream

A patient receives a constant intravenous drip of a drug at rate D (mg per hour). The body eliminates the drug at a rate proportional to the current concentration $C(t)$ (mg per litre), with elimination rate k (per hour).
Word equation:

$$\text{rate of change of drug in body} = \text{infusion rate} - \text{elimination rate}.$$

Model: Let V be the volume of distribution (litres). The total amount of drug is $V\,C$, so

$$V\frac{\mathrm{d}C}{\mathrm{d}t} = D - k\,V\,C \qquad \Longrightarrow \qquad \frac{\mathrm{d}C}{\mathrm{d}t} = \frac{D}{V} - k\,C.$$

Check dimensions: $[\mathrm{d}C/\mathrm{d}t] = (\mathrm{mg/L})/\mathrm{hr}$, $[D/V] = (\mathrm{mg/hr})/\mathrm{L} = (\mathrm{mg/L})/\mathrm{hr}$, $[kC] = \mathrm{hr}^{-1}\cdot\mathrm{mg/L} = (\mathrm{mg/L})/\mathrm{hr}$. ✓
Solve: This is a first-order linear ODE. With $C(0) = 0$ (no drug initially):

$$C(t) = \frac{D}{kV}(1 - e^{-kt}).$$

As $t \to \infty$, $C \to D/(kV)$: the **steady-state concentration**. The time to reach 95% of steady state is $t \approx 3/k$ (see if you can work out why).
Refine: In reality, the drug distributes unevenly between blood and tissue, leading to multi-compartment models with systems of ODEs.

1.7 Discrete-time models: difference equations

All of the models so far have been *continuous-time* models: the variable changes smoothly and we describe its evolution with a differential equation. But not every situation is naturally continuous. Many real-world processes evolve in distinct steps:

- An insect population breeds once per year, we observe it at yearly intervals.
- A bank account earns interest at the end of each month.
- A drug is taken as a daily pill rather than a continuous drip.
- A salmon population is counted at each spawning season.

In these settings, a **difference equation** is the natural modelling framework.

Definition: Difference equation

A **difference equation** is a rule that expresses the value of a sequence at step $n+1$ in terms of its value at step n (and possibly earlier steps):

$$x_{n+1} = f(x_n), \qquad n = 0,\, 1,\, 2,\, \ldots$$

together with an initial value x_0. This is the discrete analogue of the initial value problem $x'(t) = g(x)$, $x(0) = x_0$.

The index n plays the role of time, but it takes only integer values: $n = 0,\, 1,\, 2,\, \ldots$. Given the rule f and the starting value x_0, we can compute the entire sequence by iteration:

$$x_1 = f(x_0), \quad x_2 = f(x_1), \quad x_3 = f(x_2), \quad \ldots$$

1.7.1 A first discrete model: geometric growth

Consider a population of insects that breeds once per year, with each generation producing R times as many offspring as the previous one. Let N_n denote the population in year n. Then

$$N_{n+1} = R\,N_n, \qquad N_0 \text{ given.} \tag{1.5}$$

This is the discrete analogue of the exponential growth ODE $\mathrm{d}N/\mathrm{d}t = rN$.

Solving by iteration. Applying the rule repeatedly:

$$N_1 = R\,N_0, \quad N_2 = R\,N_1 = R^2\,N_0, \quad N_3 = R^3\,N_0, \quad \ldots$$

The pattern is clear:

$$\boxed{N_n = R^n\,N_0.} \tag{1.6}$$

This is a geometric sequence, the discrete counterpart of the exponential function $N(t) = N_0\,e^{rt}$.

The long-term behaviour depends on R:

- $R > 1$: the population grows without bound ($N_n \to \infty$).
- $R = 1$: the population is constant ($N_n = N_0$ for all n).
- $0 < R < 1$: the population decays to zero ($N_n \to 0$).
- $R = 0$: the population is wiped out immediately ($N_1 = 0$).

Example Annual interest

A savings account pays 5% interest per year. The balance B_n after n years satisfies

$$B_{n+1} = 1.05\,B_n, \qquad B_0 = 1000.$$

The solution is $B_n = 1000 \times 1.05^n$. After 10 years: $B_{10} = 1000 \times 1.05^{10} \approx 1628.89$.

1.7.2 The connection between discrete and continuous models

The geometric growth model $N_{n+1} = R\,N_n$ and the exponential ODE $\mathrm{d}N/\mathrm{d}t = rN$ are closely related. To see the connection, suppose the discrete time step corresponds to an interval of length h. Then the ODE model predicts

$$N(t+h) = N(t)\,e^{rh},$$

so the multiplication factor per step is $R = e^{rh}$, or equivalently $r = (\log R)/h$. When the step size h is small and R is close to 1, we can write $R = 1 + rh$ (the first-order Taylor approximation of e^{rh}), and the difference equation $N_{n+1} = (1 + rh)\,N_n$ becomes the Euler method approximation to the ODE. This link between difference equations and numerical methods will reappear when we study Euler's method later in the course.

1.7.3 First-order linear difference equations

A more general class of difference equations takes the form

$$x_{n+1} = a\,x_n + b, \qquad x_0 \text{ given,} \tag{1.7}$$

where a and b are constants. This is the discrete analogue of the first-order linear ODE $x' = -kx + c$. The case $b = 0$ is the geometric growth model we already solved.

Solution: We use a technique that mirrors the method for ODEs: find the equilibrium (steady state) and then track departures from it. An equilibrium is a value x^* such that $x_{n+1} = x_n = x^*$. Substituting into (1.7):

$$x^* = a\,x^* + b \qquad \Longrightarrow \qquad x^* = \frac{b}{1-a}, \quad \text{provided } a \neq 1.$$

Now define the deviation from equilibrium: $e_n = x_n - x^*$. Substituting $x_n = x^* + e_n$ into (1.7):

$$x^* + e_{n+1} = a(x^* + e_n) + b = a\,x^* + b + a\,e_n = x^* + a\,e_n,$$

so $e_{n+1} = a\,e_n$, which is simply geometric growth in the deviation. Therefore $e_n = a^n\,e_0 = a^n(x_0 - x^*)$, and

$$\boxed{x_n = x^* + a^n\,(x_0 - x^*) = \frac{b}{1-a} + a^n\left(x_0 - \frac{b}{1-a}\right).} \tag{1.8}$$

Stability: The equilibrium x^* is stable (attracting) if $|a| < 1$, because $a^n \to 0$ and therefore $x_n \to x^*$ regardless of the starting value. If $|a| > 1$, then $|a^n| \to \infty$ and the solution diverges: the equilibrium is unstable. If $|a| = 1$ the equilibrium is neutrally stable (the deviation neither grows nor decays). Note an important difference from the continuous case: when $-1 < a < 0$ the terms a^n alternate in sign, so the solution *oscillates* around the equilibrium while converging to it. This oscillatory behaviour has no analogue in first-order ODEs, where solutions approach equilibria monotonically.

Example Drug dosing — discrete model

A patient takes 100 mg of a drug every morning. Between doses, the body eliminates 30% of the drug present. Let D_n be the amount of drug in the body just *after* the n-th dose.

Just before the n-th dose, the amount remaining from the previous dose is $0.7\,D_{n-1}$ (70% survives). Adding the new dose:

$$D_n = 0.7\,D_{n-1} + 100, \qquad D_1 = 100.$$

This has the form (1.7) with $a = 0.7$ and $b = 100$. The equilibrium is

$$D^* = \frac{100}{1 - 0.7} = \frac{100}{0.3} \approx 333 \text{ mg}.$$

Since $|a| = 0.7 < 1$, the drug level converges to about 333 mg regardless of the starting dose. The solution is

$$D_n = 333 + (0.7)^{n-1}(100 - 333) = 333 - 233 \times 0.7^{n-1}.$$

After 5 doses: $D_5 = 333 - 233 \times 0.7^4 \approx 333 - 56 = 277$ mg.

Compare this with the continuous drug model from the previous section: the steady state $D/(kV)$ plays the same role as $b/(1-a)$, and the exponential approach e^{-kt} is replaced by the geometric decay a^n.

1.7.4 Worked example: the discrete logistic model

The geometric growth model $N_{n+1} = R\,N_n$ predicts unlimited growth when $R > 1$. In reality, no population grows forever: food runs out, space becomes scarce, disease

spreads more easily in crowded conditions. There is a maximum population size that the environment can sustain, which we call the carrying capacity K. When the population is small relative to K, resources are plentiful and growth is essentially unrestricted. As the population approaches K, competition for resources intensifies and the effective growth rate must decrease.

The simplest way to build this into the geometric model is to multiply the growth factor by a term that equals 1 when N_n is small and 0 when $N_n = K$:

$$N_{n+1} = R\,N_n\left(1 - \frac{N_n}{K}\right), \tag{1.9}$$

where $R > 0$ is the intrinsic growth rate (the growth rate in the absence of competition for resources) and $K > 0$ is the carrying capacity. This is the **discrete logistic model**. When $N_n \ll K$, the factor $(1 - N_n/K)$ is close to 1 and we recover geometric growth $N_{n+1} \approx R\,N_n$. When N_n is close to K, the factor suppresses growth, and at $N_n = K$ it shuts growth off entirely. We will reinforce this idea when we meet a continuous-time version of this idea, the logistic ODE, later in the course.

Equilibria: An equilibrium is a population level N^* that, once reached, persists: $N^* = R\,N^*(1 - N^*/K)$. There are two solutions:

- $N^* = 0$ (trivial equilibrium — extinction).
- $N^* = K(1 - 1/R)$ (non-trivial, exists only for $R > 1$).

For moderate growth factors ($1 < R < 3$), the population converges to the equilibrium $K(1 - 1/R)$, settling down to a steady level just as we might expect. For larger values of R, however, the dynamics become more complicated: the population may oscillate between two values, then four, then eight, and eventually behave in a way that appears chaotic, completely unpredictable despite the rule being perfectly deterministic. This phenomenon, in which a simple rule generates complex behaviour, is one of the most striking discoveries in applied mathematics and you will explore this further in the exercises. A more thorough treatment of chaos will be provided in an optional 3$^{\text{rd}}$ year course at Lancaster.

Key idea

Discrete-time models can exhibit behaviours (oscillations and chaos) that have no counterpart in first-order continuous models. This is not a flaw of the model; it reflects genuine features of systems that evolve in discrete steps. The choice between a continuous and a discrete model should be guided by the nature of the process being modelled.

Exercise 1.1

The period of a simple pendulum of length l in gravitational field g is claimed to be

$$\tau = 2\pi\sqrt{\frac{\ell}{g}}.$$

Check dimensional homogeneity. What are the dimensions of g?

Exercise 1.2

A radioactive isotope has a half-life of 5730 years (carbon-14).

(a) Find the decay rate r.

(b) What fraction of the original material remains after 1000 years?

(c) Write a sentence explaining what the sign of r means physically.

Exercise 1.3

A cup of coffee is at 90°C in a room at 20°C. Assume Newton's law of cooling:

$$\frac{\mathrm{d}T}{\mathrm{d}t} = -k(T - T_{\text{room}}), \quad k > 0.$$

Let $\theta = T - T_{\text{room}}$. Show that $\theta(t) = \theta_0\, e^{-kt}$ and hence $T(t) = T_{\text{room}} + (T_0 - T_{\text{room}})\, e^{-kt}$.

Exercise 1.4

A cylindrical tank of cross-sectional area $A = 1\ \text{m}^2$ has a drain hole of area $a = 5 \times 10^{-4}\ \text{m}^2$ at the bottom. The initial water depth is $h_0 = 3\ \text{m}$.

(a) Write down the ODE governing $h(t)$ and verify dimensional homogeneity.

(b) Solve the ODE with initial condition $h(0) = h_0$.

(c) Calculate the drain time.

(d) Sketch $h(t)$ and explain why the graph is concave up (curves upward).

(e) A discharge coefficient $C_d = 0.6$ is introduced to account for viscous losses. How does this change the drain time?

Exercise 1.5

A lake has volume $V = 10^6\ \text{m}^3$. A factory discharges pollutant into the lake at a constant rate P kg/day. Clean water flows in at rate $Q\ \text{m}^3$/day, and the lake drains at the same rate (so the volume stays constant). Assume the lake is well-mixed.

(a) Let $C(t)$ be the pollutant concentration (kg/m^3). Write a word equation for the rate of change of pollutant mass in the lake.

(b) Formulate this as an ODE for $C(t)$. State clearly what assumptions you are making.

(c) Solve the ODE with $C(0) = 0$ and find the steady-state concentration.

(d) The government requires $C < 0.01$ kg/m^3. If $Q = 5000\ \text{m}^3$/day, what is the maximum permitted discharge rate P?

(e) What assumptions might break down in practice? Suggest one refinement to the model.

Exercise 1.6

A savings account earns 3% annual interest, and the holder deposits an additional £500 at the end of each year. Let B_n be the balance after n years.

(a) Write down a difference equation for B_n with $B_0 = 2000$.

(b) Find the general solution.

(c) What is the balance after 20 years?

(d) Compare with the continuous model $\mathrm{d}B/\mathrm{d}t = 0.03\,B + 500$, $B(0) = 2000$. Solve the ODE and evaluate $B(20)$. How do the two answers differ, and why?

Exercise 1.7

A patient takes 200 mg of a drug every 8 hours. The body eliminates 40% of the drug between doses. Let D_n be the amount of drug in the body just after the n-th dose.

(a) Write down the difference equation for D_n.

(b) Find the steady-state drug level D^*.

(c) If the therapeutic window requires the drug level to be between 250 mg and 600 mg, is this dosing regimen safe? Justify your answer.

(d) How many doses does it take for the drug level to reach 95% of the steady state?

Exercise 1.8

The discrete logistic model is $N_{n+1} = R\,N_n(1 - N_n/K)$.

(a) Nondimensionalise by setting $u_n = N_n/K$. Show that the equation becomes $u_{n+1} = R\,u_n(1 - u_n)$.

(b) Find the equilibria of the nondimensional model and show that the non-trivial equilibrium $u^* = 1 - 1/R$ exists only for $R > 1$.

(c) Using a calculator or computer, iterate the map $u_{n+1} = R\,u_n(1 - u_n)$ starting from $u_0 = 0.2$ for $R = 2.0$, $R = 3.2$, and $R = 3.8$. In each case, compute the first 30 terms and describe the long-term behaviour.

(d) Based on your observations, explain in a sentence or two why the discrete logistic model can behave very differently from the continuous logistic ODE $\mathrm{d}N/\mathrm{d}t = rN(1 - N/K)$.