Applied Computational Mathematics Lab

EPIDEMIC MODELING · MODULE 3

Compartmental models

SIR and its family, R_0 and the next-generation matrix



6 EPISODES

1. SIR from scratch and the meaning of R_0
2. More states: births, deaths, reinfection, latency
3. More classes: stages, outcomes, structure, vectors
4. When transmission changes
5. Chance matters: stochastic models
6. The next-generation matrix

Based on Chapter 3 of T. Lazebnik, Epidemic Modeling: Principles, Algorithms, and Applications (Academic Press, 2026)

acml.teddylazebnik.com

BEFORE WE START

Warm-up: true or false?

CHECK YOURSELF

(a) In the SIR model, the epidemic stops because everyone has been infected.

(b) Adding a latent period (SEIR) changes how many people are infected in the end.

(c) With $R_0 = 2$, a single case in a large town always starts a big outbreak.

HOW TO USE THIS

Guess now. (a) is settled in Part 1, (b) in Part 2, (c) in Part 3; all answers are at the end of the deck.

WHY THIS MODULE

This module is the mathematical core of the course: one small model, then the family of extensions that every later module builds on.

BEFORE WE START

This module's map

PART 1

SIR from scratch

Assumptions, equations, R_0, the threshold, final size, observation.

PART 2

More states and classes

Births and deaths, reinfection, latency, stages, outcomes, groups, vectors.

PART 3

Changing rates and chance

Seasons, interventions, behavior; stochastic models and early extinction.

PART 4

The next-generation matrix

R_0 for any number of infected types, and what it can and cannot tell you.

GOALS

By the end you should be able to write and solve an SIR-family model, compute R_0 and the final size, explain early extinction, and build a next-generation matrix.

ONE WORD

Compartment: a group of people in the same epidemiological state. The model tracks how many are in each, and the rates at which they move.

Applied Computational Mathematics Lab

PART 1

SIR from scratch

Three boxes, two rates, a surprising amount of insight.



Three boxes and five assumptions



IN PLAIN WORDS

Removed can mean recovered with immunity, isolated, or dead, depending on the question.

ASSUMPTION	WHAT IT IGNORES	RELAXED IN
everyone mixes equally	households, ages, places	Parts 2 and 4; Module 4
mass action	saturation, crowd effects	this part
closed population	births, deaths, travel	Part 2
permanent immunity	waning, reinfection	Part 2
exponential infectious period	regular durations	Part 2 (stages)

The SIR equations

$$\frac{dS}{dt} = -\beta SI, \qquad \frac{dI}{dt} = \beta SI - \gamma I, \qquad \frac{dR}{dt} = \gamma I$$

- β: transmission rate = how often people meet × chance that a meeting transmits.
- γ: recovery rate; the average time infectious is $1/\gamma$.
- S, I, R as shares of the population add up to 1 at all times.

WHICH INCIDENCE?

Density-dependent: βSI with S, I as counts: transmission grows with crowd size.

Frequency-dependent: $\beta SI/N$: each person makes a fixed number of contacts, whatever the population. With shares instead of counts the two look the same, but β has different units and scaling.

TRY IT

A town doubles in size and people keep the same daily contacts. Which form keeps the epidemic's speed unchanged?

PART 1 · SIR

R_0 and the epidemic threshold

$$R_0 = \frac{\beta}{\gamma}, \qquad \frac{dI}{dt} > 0 \iff S > \frac{1}{R_0}, \qquad R(t) = R_0\, S(t)$$

- At the start S ≈ 1, so one case infects about β/γ others over its infectious period.
- Example: 10 contacts a day × 3% chance per contact × 10 days = 3.
- Growth starts only if the susceptible share is above $1/R_0$; R(t) crosses 1 exactly at the peak.

B PER DAY	DAYS INFECTIOUS	R_0	THRESHOLD S	DOUBLING
0.3	10	3	33%	3.5 d
0.6	5	3	33%	1.7 d
0.2	10	2	50%	6.9 d

SPEED VS R_0

R_0 is not speed. Early growth rate is β − γ: with β = 0.3, γ = 0.1, that is 0.2 per day, doubling every ln 2 / 0.2 ≈ 3.5 days. Same R_0 with β = 0.6, γ = 0.2 doubles in 1.7 days.

TRY IT

With β = 0.25 per day and 8 days infectious on average, find R_0 and the susceptible share below which the outbreak shrinks.

Our run: the peak at S = 1/R_0

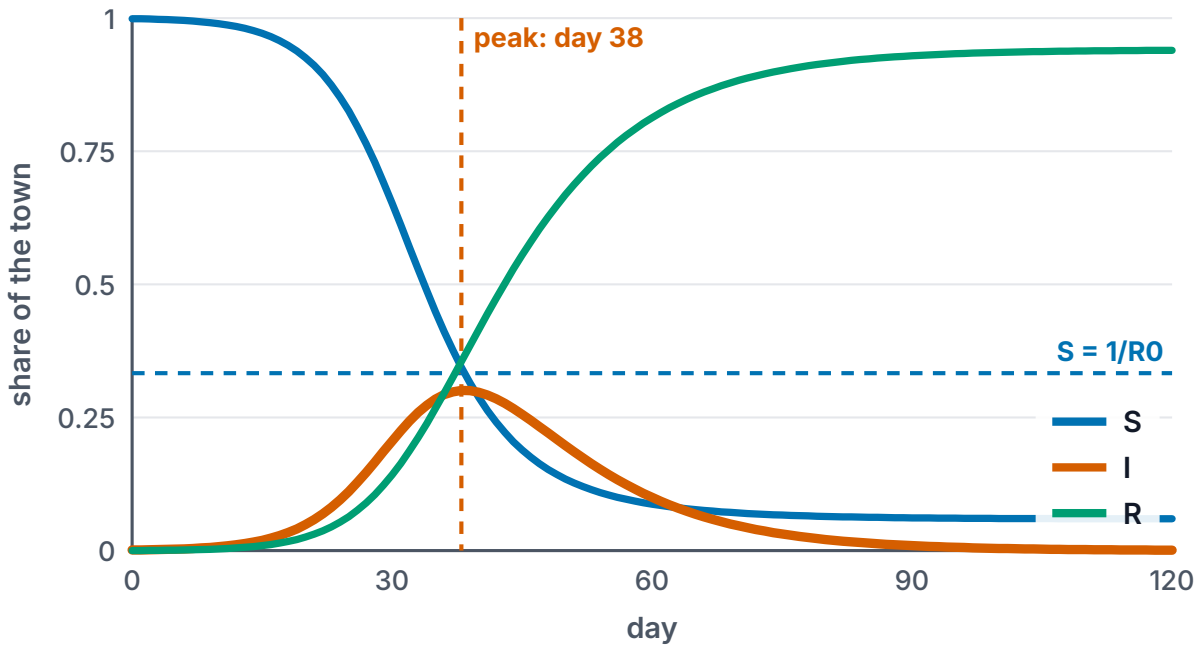

Our run: β = 0.3, γ = 0.1 per day (R_0 = 3), 1 in 1,000 infected on day 0; RK4 with steps of 0.02 day.

- Infections peak on day 38, with 30% of the town infected at the same time.
- At the peak, about a third of the town is still susceptible: S = 1/R_0, so R(t) = 1.
- After the peak each case infects fewer than one other, and infections fall, though people keep getting infected.

IN PLAIN WORDS

The peak is not a tired virus. It is the moment the susceptible pool gets too small for each case to replace itself.

Final size and overshoot

$$\ln s_\infty = -R_0\,(1 - s_\infty)$$

- Divide the S equation by the R equation and integrate: time drops out, and the share never infected s_∞ depends only on R_0 (starting with almost everyone susceptible).
- The epidemic ends with susceptibles left: it runs out of chains, not of people.

R_0	1.5	2	3	5
final size	58%	80%	94%	99%
herd immunity	33%	50%	67%	80%

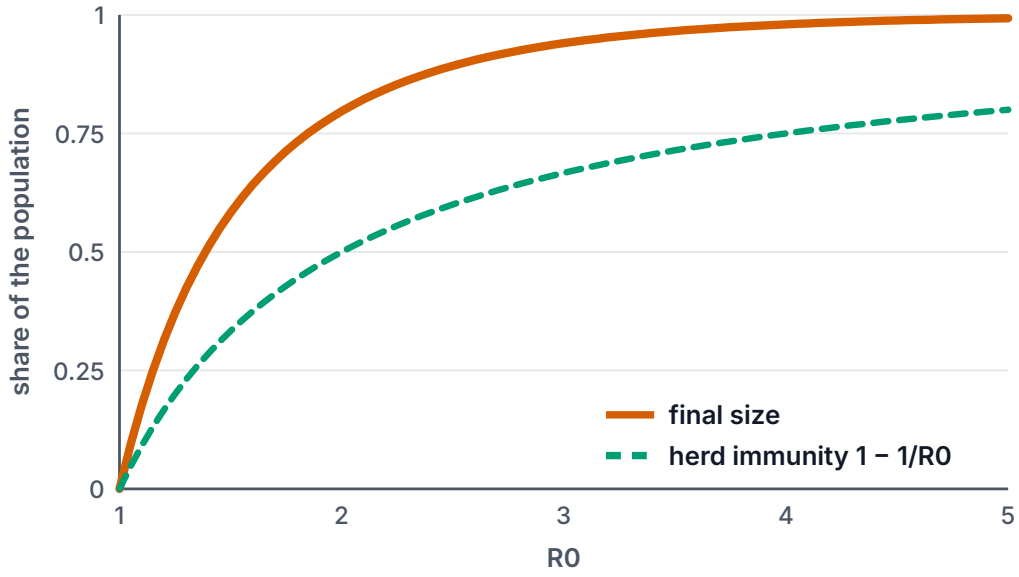

Our numbers: the final-size relation solved by fixed-point iteration.

IN PLAIN WORDS

Overshoot: with R_0 = 3, herd immunity is reached at 67% immune, but the epidemic runs on to 94%: about a quarter of the town is infected after the threshold has passed.

Counts stay non-negative

- In the exact equations, S only decreases toward 0, R only increases, and I cannot cross below 0: the region where S, I and R are all at least 0 is forward-invariant.
- So a negative count signals a broken numerical method or a wrong model, never a feature of SIR.
- Our run with the simplest method (forward Euler) and 10-day steps: S goes from 99% to -6.7% by day 50.

IN PLAIN WORDS

Small steps, or a better method (Runge–Kutta), keep S positive. Module 12 covers numerical stability.

DAY	S (EULER, 10-DAY STEPS)	I
0	99.0%	1.0%
10	96.0%	3.0%
20	87.5%	8.6%
30	65.0%	22.5%
40	21.2%	43.8%
50	-6.7%	27.9%
60	-1.1%	-5.6%

TRY IT

Why does one Euler step of size Δt send S negative when $\beta I \Delta t > 1$?

We never see S, I and R

- Data are reported cases, hospital admissions, deaths: indirect, delayed and biased views of the infection process.
- Applied models add an **observation layer**: a reporting share applied to true incidence, a delay distribution (convolution), and count noise (often negative binomial).
- The flows of the SIR model do not change; the layer links model states to what is measured.

A SIMPLE OBSERVATION MODEL

$\text{reported}(t) \approx \rho \times \Sigma_d f(d) \times \text{new infections}(t - d)$, plus noise. ρ: reporting share; f: delay distribution.

LOOKING AHEAD

Modules 7 and 8 fit models through such layers; Module 9 looks at what each data stream really measures.

Check yourself: Part 1

CHECK YOURSELF

1. $\beta = 0.4$ per day and an infectious period of 5 days: what is R_0, and below what susceptible share does the outbreak shrink?

2. For $R_0 = 2$, compare the final size with the herd immunity threshold. Why do they differ?

3. Give one reason a simulation of SIR can show negative numbers.

TIP

Answers are at the end of the deck. Try them before you look.

Applied Computational Mathematics Lab

PART 2

More states and classes

Relaxing the assumptions, one box at a time.



Births and deaths: endemic disease

$$\frac{dS}{dt} = \mu - \beta SI - \mu S, \quad \frac{dI}{dt} = \beta SI - (\gamma + \mu)I, \quad \frac{dR}{dt} = \gamma I - \mu R$$

- R_0 = β / (γ + μ). Our town: β = 0.5, γ = 1/10 per day, life expectancy 50 years: $R_0 \approx 5$.
- Endemic balance: $S^* = 1/R_0 = 20\%$, $I^* = \mu (R_0 - 1) / \beta$.
- Average age at infection ≈ life expectancy / (R_0 − 1) = 12.5 years: the higher R_0, the earlier people catch it.

IN PLAIN WORDS

No seasons and no behavior change: births alone refill the susceptible pool, and the disease comes back in multi-year waves that slowly settle toward balance.

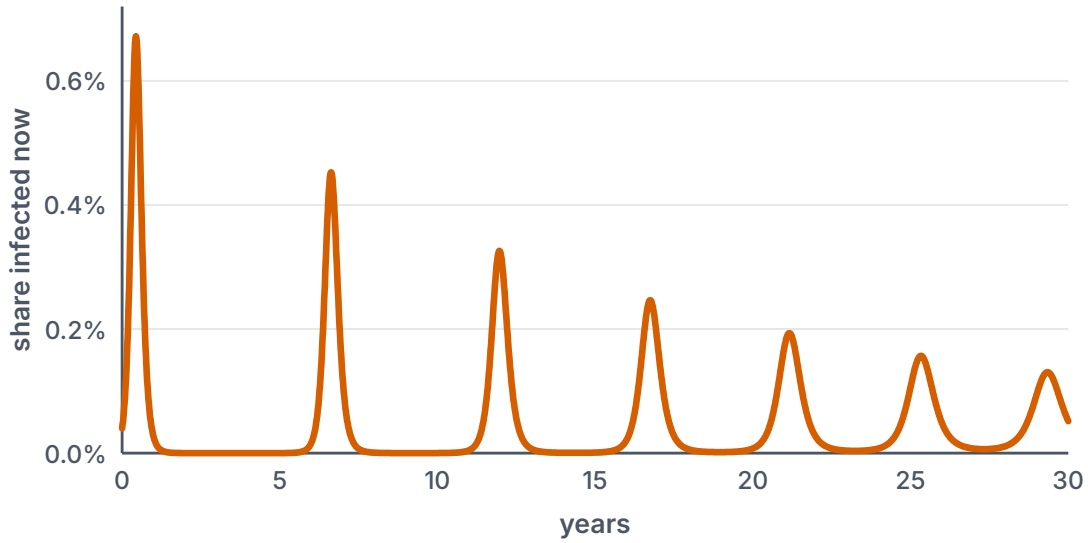

Our run, started near balance (25% susceptible): waves at years 7, 12, 17, 21… settling down; theory near balance: a period of about 3.7 years.

How much vaccination?

$$R_{\text{eff}} = R_0\,(1-p) < 1 \iff p > p_c = 1 - \frac{1}{R_0}$$

- Vaccinating a share p of newborns (with a fully protective vaccine) moves them straight to R.
- Our town (R_0 = 5): p_c = 80%.
- An imperfect vaccine with efficacy e needs p > p_c / e, which can exceed 100%: then vaccination alone cannot eliminate.
- Vaccinating adults or children later in life also counts, but people infected before vaccination are already removed: timing matters.

HERD IMMUNITY

Herd immunity: enough of the population is immune that each case infects fewer than one other on average, so the vaccinated also protect the unvaccinated.

R_0	1.5	2	3	5	15
p_c	33%	50%	67%	80%	93%

TRY IT

A vaccine is 90% effective. With R_0 = 5, what share must be vaccinated? And with R_0 = 15?

Reinfection: SIS and SIRS

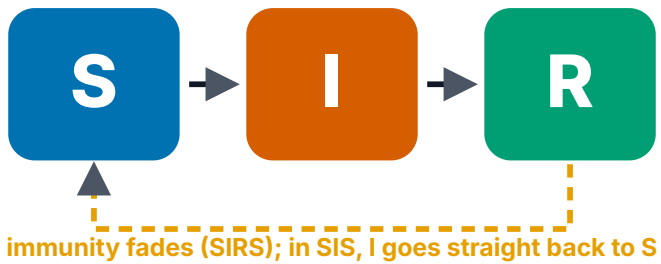

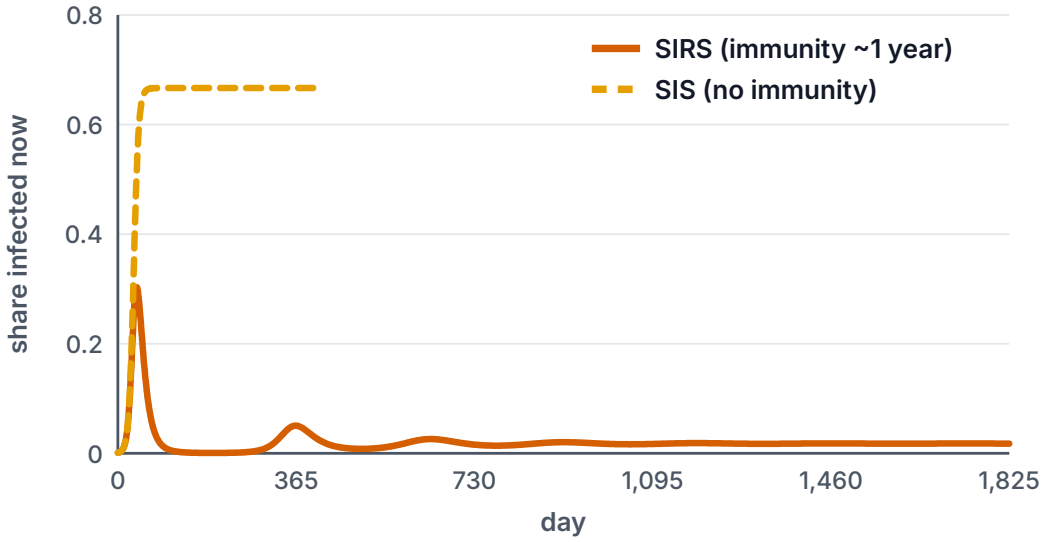

- **SIS** (no immunity): endemic level $I^* = 1 - 1/R_0$; with $R_0 = 3$ our run settles at 67%, with no waves. Some bacterial infections behave roughly like this.
- **SIRS** (immunity fades at rate ω): $I^* = (1 - 1/R_0)\ \omega / (\omega + \gamma)$. With immunity lasting a year: 1.8%.
- SIRS brings damped waves: in our run peaks on days 38, 364, 641, 913...
- Seasonal coronaviruses and common colds are closer to SIRS: the same virus can infect you again a year or two later.

Latency: the SEIR model

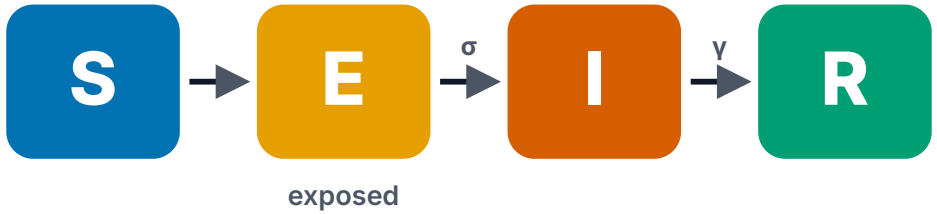

$$\frac{dE}{dt} = \beta SI - \sigma E, \qquad \frac{dI}{dt} = \sigma E - \gamma I, \qquad (r+\sigma)(r+\gamma) = \beta\sigma$$

- Exposed people are infected but not yet infectious; the latent period averages 1/σ.
- R_0 stays β/γ (without deaths), but early growth r is slower: it solves the quadratic above.
- Fitting an SIR model to SEIR-like data gets the speed right only with a wrong β or γ: the latent period matters for timing and for control (contact tracing has a head start).

R_0 = 3	SIR	SEIR (5 DAYS LATENT)
growth per day	0.2	0.1
doubling time	3.5 d	6.9 d
peak day	38	74
peak infected	30%	20%
final size	94%	94%

IN PLAIN WORDS

Latency changes the timing, not the total.

Stages: more regular durations

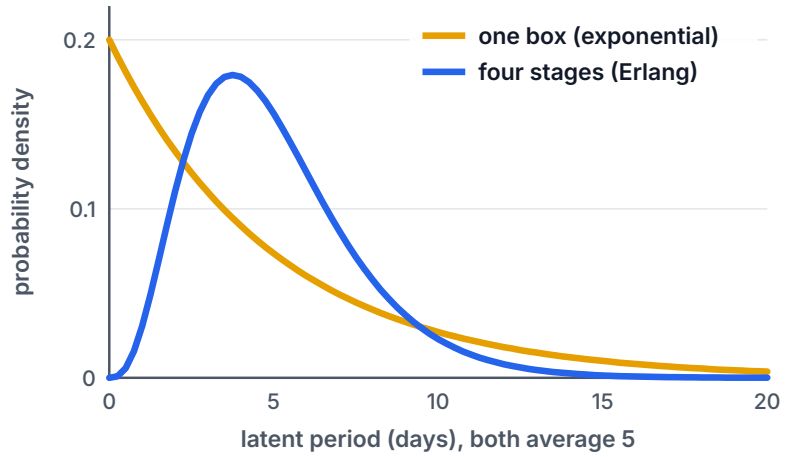

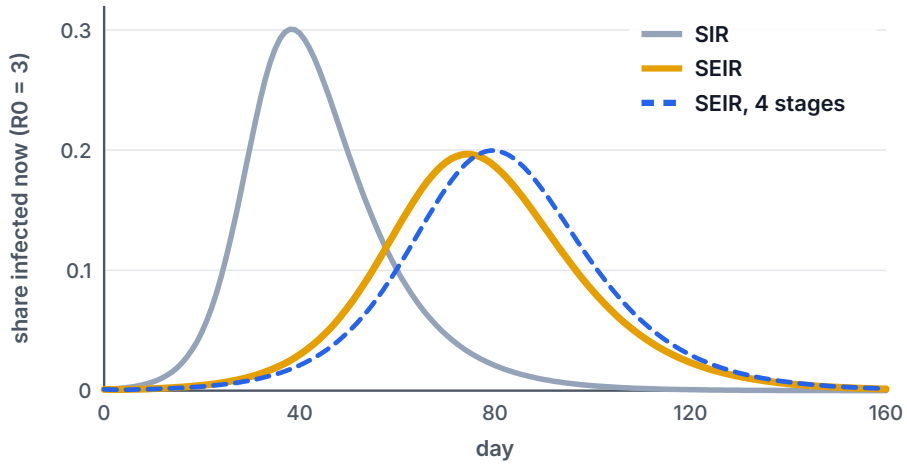

- One exposed box means a constant exit rate: the latent period is exponential, and some people leave after a day, some after twenty.
- Splitting E into k boxes in a row, each left at rate kσ, keeps the mean 1/σ but gives an Erlang distribution with much less spread.
- In our runs (4 stages), the peak moves from day 74 to day 80; the final size is unchanged.
- The same trick works for infectious periods and hospital stays.

CHOOSING

Which box to add? Births and deaths for long time scales; waning for diseases you catch again; latency or stages whenever delays matter for timing or control.

PART 2 · CLASSES

Several infectious classes

STAGE (TOY DISEASE)	WHO	DAYS	INFECTIOUSNESS	SHARE OF R_0
before symptoms	everyone	2	full	32%
with symptoms	7 in 10	5	full	56%
no symptoms	3 in 10	5	half	12%

$$R_0 = \beta\,(2 + 0.7 \cdot 5 + 0.3 \cdot 0.5 \cdot 5) = 0.3 \times 6.25 \approx 1.9$$

- If people with symptoms isolate after one day, R_0 falls to 1.03: much lower, still above 1.
- The hidden stages (before symptoms, never ill) keep transmission going.

TRY IT

What share of symptomatic people would have to isolate after one day to bring R_0 exactly to 1?

IN PRACTICE

Measuring each stage's duration and infectiousness (from contact tracing and viral load studies) is what turns this toy into a usable model.

Clinical outcomes and delays

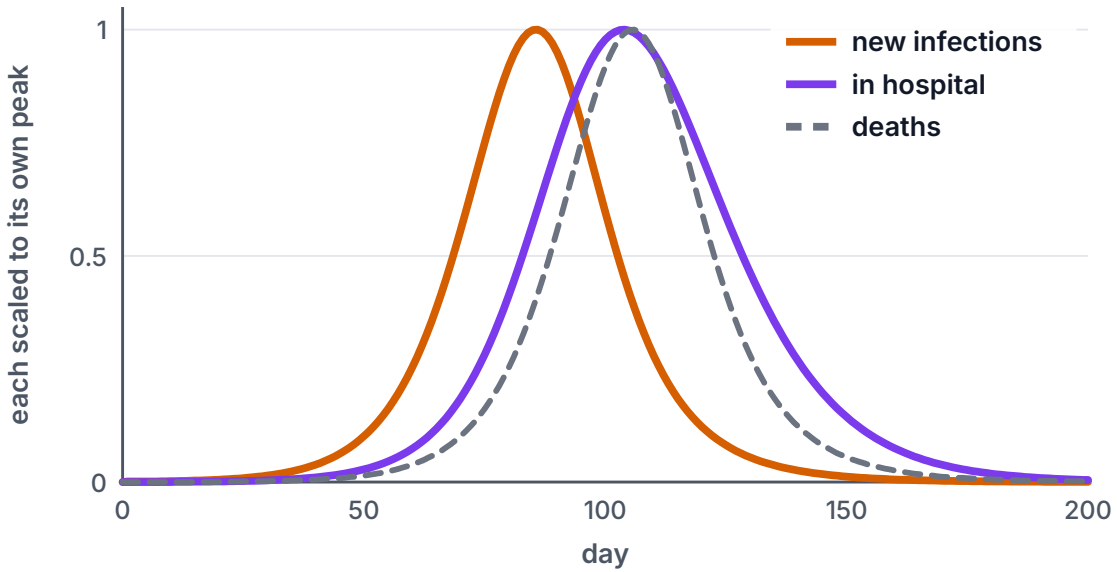

Our SEIR run (R_0 3): 3% of infections admitted about a week after becoming infectious, for 10 days; 0.5% die about 20 days after infection.

- New infections peak on day 86, hospital occupancy on day 104 (about 551 beds per 100,000), deaths on day 106.
- Total deaths: about 470 per 100,000.
- Hospitals are fullest weeks after new infections have started to fall: decisions based on hospital data are late by design.

TWO RATIOS

IFR: deaths ÷ all infections. **CFR**: deaths ÷ reported cases. Early in an outbreak, when few mild cases are tested, CFR can be many times the IFR.

Structured populations

- Split the population into groups (ages, risk, places); each has its own S, I, R, linked by a contact matrix.
- Our two-group town: children are a quarter of the people; contacts per day: children 10 with children and 6 with adults; adults 2 with children and 6 with adults (reciprocity: 0.25 × 6 = 0.75 × 2).
- With 3% transmission per contact and 7 days infectious: 94% of children and 73% of adults are infected in the end.

IN PLAIN WORDS

One R_0 for the whole town hides who drives transmission. Part 4 computes R_0 from the matrix.

CONTACTS PER DAY	WITH CHILDREN	WITH ADULTS
a child	10	6
an adult	2	6

SCALING UP

Ten age bands give ten copies of the equations and a 10 × 10 matrix; contact surveys supply the numbers.

Vector-borne diseases

- Malaria and dengue pass from people to mosquitoes and back. Couple SIR for people with SI for mosquitoes (infected for life).
- Mosquitoes add their own clock: the parasite needs days to develop inside the mosquito, and many mosquitoes die first.
- Toy numbers: one infected person infects 10 mosquitoes; each infected mosquito infects 0.3 people: 3 people per person over two steps.

CHANGE	PEOPLE PER PERSON (TWO STEPS)
none	3.00
half the mosquitoes	1.50
half the bites (bed nets)	0.75

IN PLAIN WORDS

Bites count twice: once to infect the mosquito, once to infect the next person. In the Ross–Macdonald formula the biting rate appears squared.

Check yourself: Part 2

CHECK YOURSELF

4. $R_0 = 4$ and a fully protective vaccine: what share of newborns must be vaccinated?

5. An SIS disease has $R_0 = 2.5$. What share is infected at the endemic balance?

6. SEIR vs SIR with the same R_0: which changes, the timing or the final size?

7. One person infects 8 mosquitoes and each mosquito infects 0.25 people. What is the two-step number? What if bed nets halve the bites?

TIP

Answers are at the end of the deck.

Applied Computational Mathematics Lab

PART 3

Changing rates and chance

β that moves, and outbreaks that may fizzle.



Seasons

$$\beta(t) = \beta_0 \left(1 + a \cos \frac{2\pi t}{365}\right), \qquad \frac{\beta_{\max}}{\beta_{\min}} = \frac{1+a}{1-a}$$

- Crowding indoors and cold, dry air raise winter transmission for many respiratory infections; school terms set another rhythm (classic measles studies).
- Our run: a = 0.2 (ratio 1.5), β_0 = 0.25, γ = 0.1, immunity lasting about a year.
- Without seasons: a steady 1.6% infected. With seasons: yearly waves peaking at about 8%, locked to the yearly high of β.

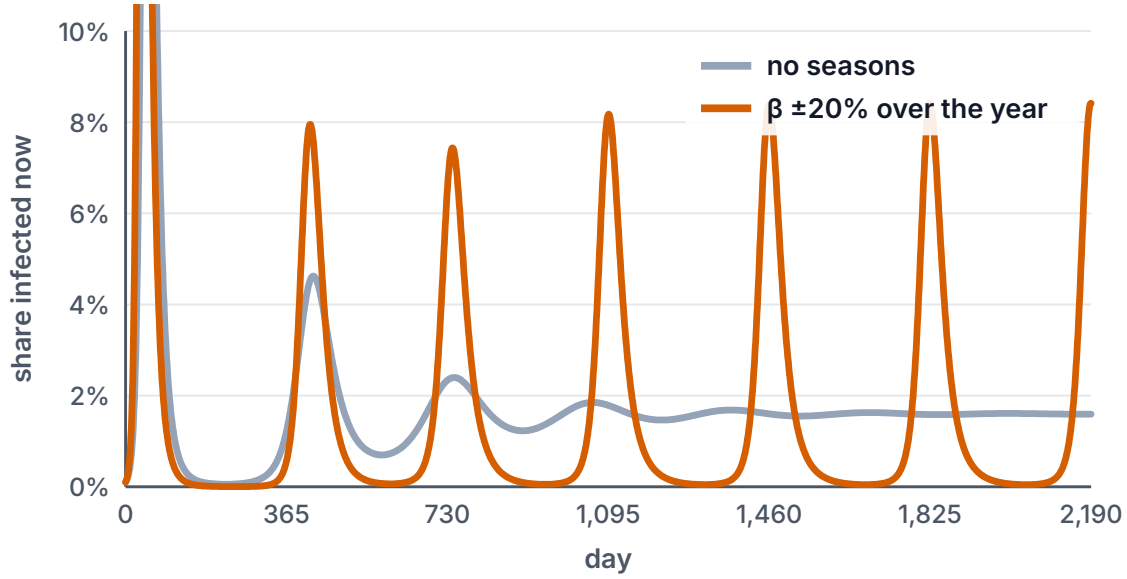

TRY IT

With a = 0.25, what is the peak-to-trough ratio of β? Why does a modest swing in β give a large swing in cases?

Interventions as steps

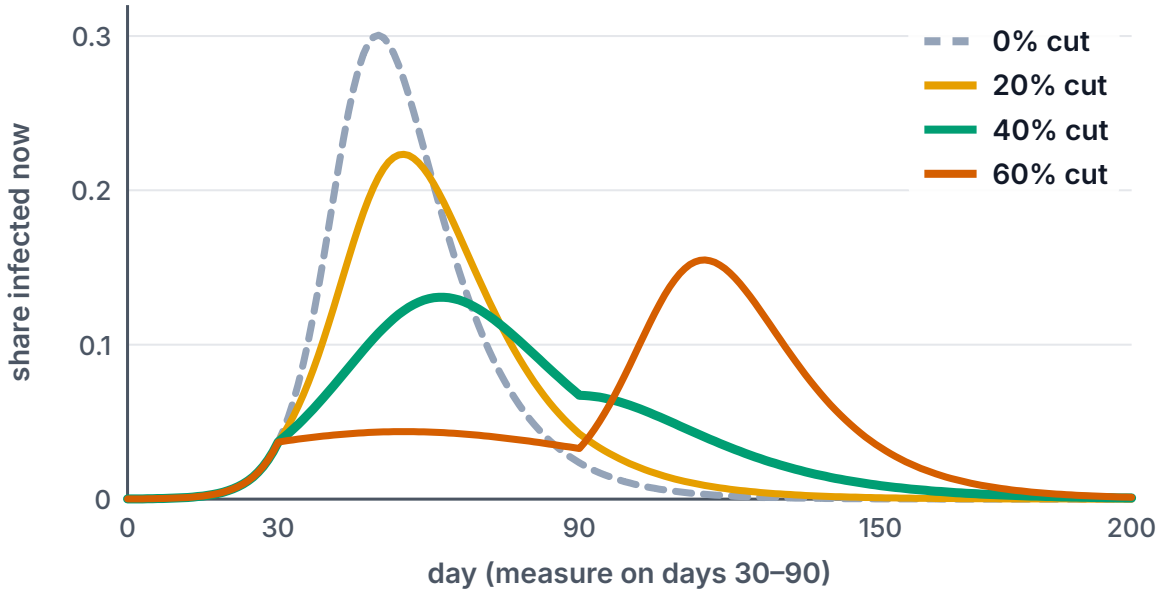

Our runs: β = 0.3, γ = 0.1; β cut on days 30–90; 1 in 10,000 infected on day 0.

CUT	PEAK	PEAK DAY	FINAL SIZE
0%	30%	50	94%
20%	22%	55	89%
40%	13%	62	84%
60%	15%	115	89%

- A 60% cut does worse than 40%: strong but short, it leaves most people susceptible when it ends, and the wave comes later and higher.
- Real measures ramp up and compliance is partial: a smooth ramp or a compliance share captures that.

Feedback, and R_0 vs R(t)

$$\beta(t) = \beta_0 \times \text{season}(t) \times \text{policy}(t) \times \text{behavior}(S, I, t)$$

- When β depends on the state (behavior reacting to cases), the model becomes a feedback system: plateaus and new waves (Module 2's lab).
- R_0 is defined at the start: everyone susceptible, nothing changed.
- R(t) is what we measure later: immunity, interventions, behavior and season all at once.

READING THE NEWS

"R is now 0.9" in the news means R(t). Estimating it from data needs the reporting delays and the generation interval (Module 7).

TRY IT

A disease has R_0 = 2.5. Half the population is immune and contacts are 20% lower than normal. What is R(t)?

TWO KINDS

A model is **time-varying** if β depends on t, and **state-dependent** if β depends on S, I or R. The second kind can generate its own waves.

Random events: the Gillespie algorithm

EVENT	CHANGE	RATE
infection	$S - 1, I + 1$	$\beta S I / N$
recovery	$I - 1, R + 1$	γI

ONE STEP OF THE GILLESPIE ALGORITHM

1. Add up the rates: total = $\beta SI/N + \gamma I$.
2. Draw the waiting time from an exponential distribution with that total rate.
3. Pick the event with probability rate / total.
4. Update the counts; repeat until $I = 0$.

- People are whole numbers; outcomes are random. Each run is one possible epidemic.
- Faster cousins: tau-leaping (draw how many events happen in a short step) and chain-binomial models such as Reed–Frost (whole generations).

IN PLAIN WORDS

Run it many times and report distributions, not one curve.

Early extinction

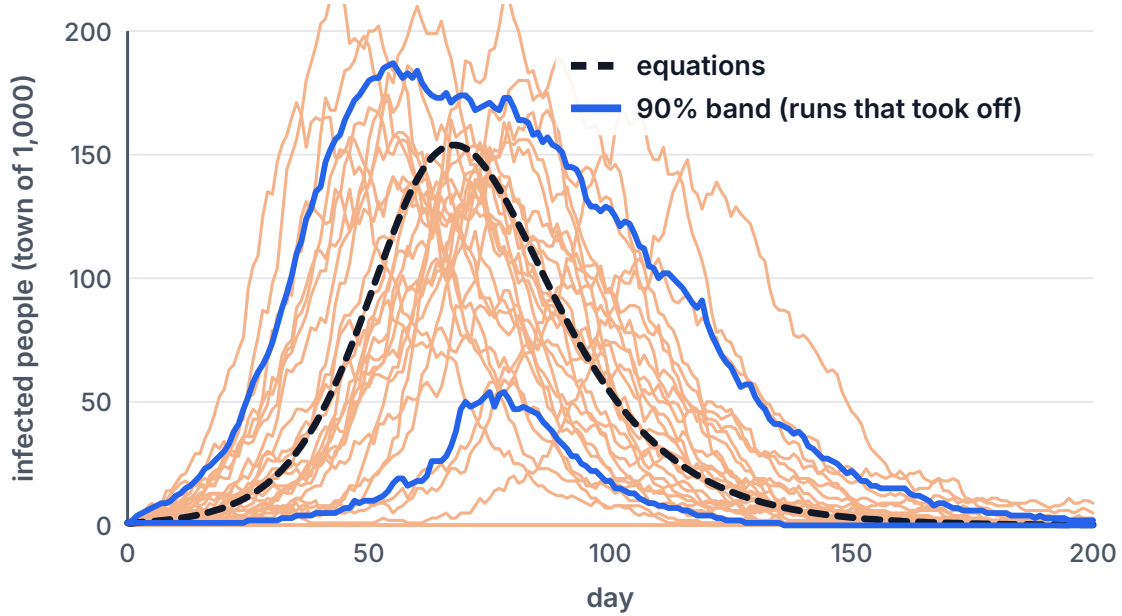

Our Gillespie runs: town of 1,000, R_0 = 2, one first case. 40 runs drawn; band from the runs that took off.

R_0	RUNS THAT FIZZLED	THEORY
1.5	66.3%	66.7%
2	49.2%	50.0%
3	31.6%	33.3%
2, three first cases	12.6%	12.5%

WHY $1/R_0$

Early on, the outbreak is a branching process. With exponential infectious periods, each case's offspring count is geometric and the extinction chance is $1/R_0$ per first case: $(1/R_0)^i$ for i first cases.

Two kinds of outcome

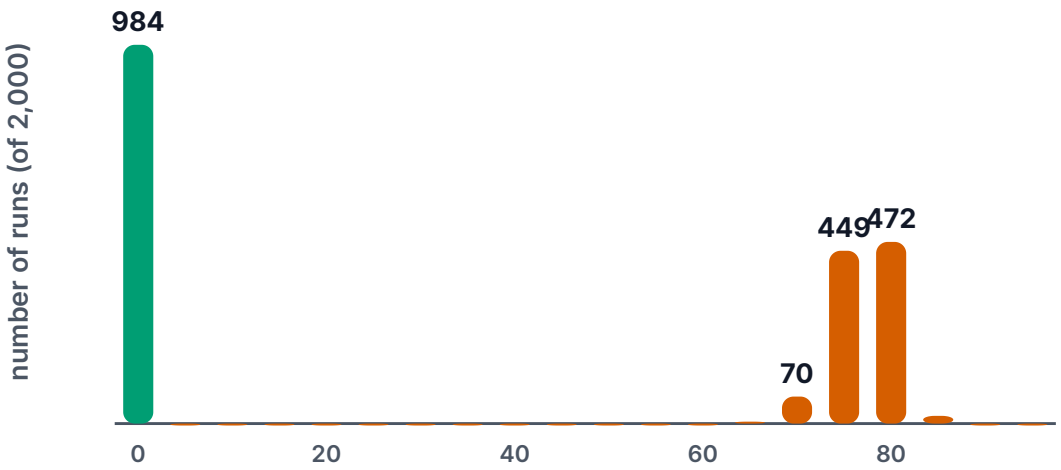

Final share infected in 2,000 runs (R_0 = 2), in 5% bins. Green: runs that fizzled.

- Outcomes split: either a tiny outbreak (49%) or a big one, almost nothing in between.
- Big outbreaks infect 80% on average; the equations say 80%.
- Their timing varies: peak on day 66 ± 15.
- Chance matters when numbers are small: first cases, small communities, the tail of an epidemic, elimination. After a wave, chains can break by chance (fade-out); measles persisted only in cities of a few hundred thousand or more.

IN PLAIN WORDS

The equations describe what happens if the outbreak takes off; chance decides whether it does.

Check yourself: Part 3

CHECK YOURSELF

8. With $\beta(t) = \beta_0(1 + 0.25 \cos(2\pi t/365))$, what is the ratio between the highest and lowest β?

9. In a stochastic SIR with $R_0 = 4$ and two first cases, roughly how likely is early extinction?

10. Why can a stronger but shorter intervention lead to a later and higher peak?

TIP

Answers are at the end of the deck.

Applied Computational Mathematics Lab

PART 4

The next-generation matrix

R_0 for any number of infected types.



The idea

- With several infected types (stages, groups, hosts and vectors), β/γ no longer works.
- Ask, for each type: one typical infected individual of this type, over its whole infection, causes how many new infections of each type?
- Those numbers form the next-generation matrix K. Generations evolve as $x_{n+1} = K\, x_n$ (with K written "to row, from column").
- R_0 is the dominant eigenvalue (spectral radius) of K: the long-run growth factor per generation. Its eigenvector is the stable mix of infected types.

THE RECIPE (NEAR THE DISEASE-FREE STATE)

1. List the infected compartments.
2. F: rates at which new infections appear in each.
3. V: all other transitions between and out of them.
4. Linearize; take Jacobians F and V at the disease-free state.
5. $K = F\, V^{-1}$; $R_0 = \rho(K)$.

Worked example: SEIR

$$F = \begin{pmatrix} 0 & \beta \\ 0 & 0 \end{pmatrix}, \quad V = \begin{pmatrix} \sigma & 0 \\ -\sigma & \gamma \end{pmatrix}, \quad V^{-1} = \begin{pmatrix} 1/\sigma & 0 \\ 1/\gamma & 1/\gamma \end{pmatrix}$$

$$FV^{-1} = \begin{pmatrix} \beta/\gamma & \beta/\gamma \\ 0 & 0 \end{pmatrix} \quad \Rightarrow \quad R_0 = \frac{\beta}{\gamma}$$

- Order of infected states: (E, I). Only I creates new infections, and they arrive in E.
- The latent rate σ cancels: latency delays infections but does not change how many each case causes.
- The same recipe adds up the contributions of Part 2's three infectious stages to 1.9.

TRY IT

Redo the example for an SEIR model with deaths at rate μ in E and I. Where do the factors σ/(σ + μ) and 1/(γ + μ) appear?

IN PLAIN WORDS

Read K column by column: column j lists the new infections, by type, caused by one individual entering infected state j.

PART 4 · NEXT-GENERATION MATRIX

The two-group town

ONE INFECTED…	INFECTS CHILDREN	INFECTS ADULTS
child	2.10	1.26
adult	0.42	1.26

Entries: 3% per contact × contacts per day × 7 days infectious (contacts from Part 2).

- Each group alone: 2.1 (children) and 1.26 (adults).
- Dominant eigenvalue: $R_0 \approx 2.52$.
- Eigenvector: the infected are 50% children and 50% adults, generation after generation, though children are 25% of the town.

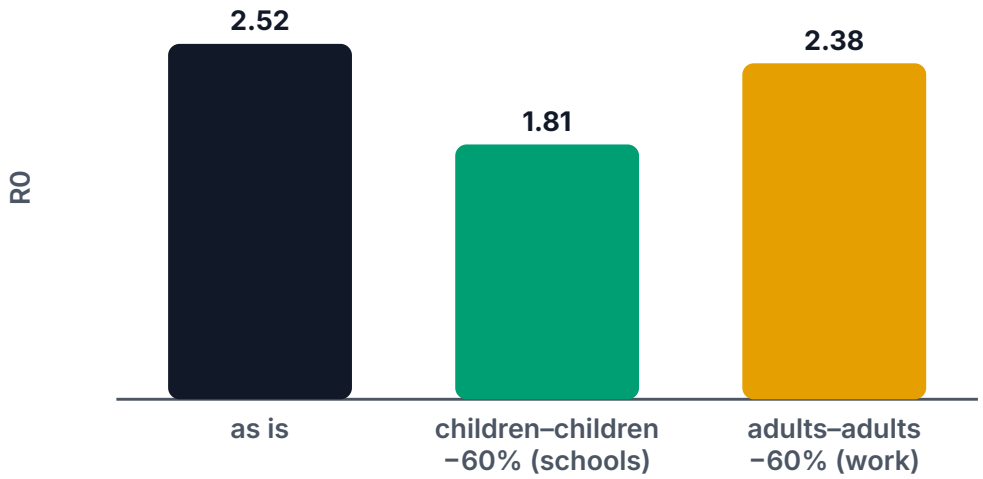

Our arithmetic: R_0 after cutting one kind of contact by 60%.

IN PLAIN WORDS

Same cut, very different effect: the matrix shows where transmission lives. Sensitivities of R_0 to each entry point to the routes worth targeting first.

Conventions and the fine print

- Vectors: in K one generation is one step (person to mosquito, or mosquito to person). Our toy gives the square root of 3, about 1.73, not 3 over two steps. Both have the threshold at 1; check which one a paper reports.
- K describes the start of an outbreak near the disease-free state: not the peak, not the final size.
- The split into F and V is a modeling choice; different splits can give different numbers with the same threshold.

A TINY EXAMPLE

K = [[1.5, 0.5], [0.5, 1.5]]: the eigenvalues are 2 and 1, so $R_0 = 2$, and the stable mix is half and half.

TRY IT

For K = [[2, 1], [0.5, 1]], compute R_0. Which entry would you reduce first?

Check yourself: Part 4

CHECK YOURSELF

11. Compute R_0 for K = [[1.2, 0.3], [0.3, 1.2]].

12. Why does the latent period cancel in the SEIR next-generation matrix?

13. A vector-borne model reports $R_0 = 1.5$ from the next-generation matrix. What is the two-step (person-to-person) number?

TIP

Answers are on the next slides.

GOING FURTHER

For each question, write which assumption of the basic SIR model the matrix approach relaxes.

ANSWERS

Answers: warm-up and Parts 1–2

- **(a)** False: it stops while some people are still susceptible (6% for $R_0 = 3$). **(b)** False: same final size, different timing. **(c)** False: with $R_0 = 2$, about half of single introductions fizzle.
- **1.** $\gamma = 0.2$, $R_0 = 2$; it shrinks when $S < 1/2$.
- **2.** Final size 80%, herd threshold 50%: infections continue after the threshold (overshoot).
- **3.** Too-large time steps with a simple method (Euler), or an invalid model.
- **4.** $1 - 1/4 = 75\%$. **5.** $1 - 1/2.5 = 60\%$. **6.** The timing; the final size is the same.
- **7.** $8 \times 0.25 = 2$ over two steps; halving bites halves both steps: $4 \times 0.125 = 0.5$.

TRY-IT ANSWERS

Try-it answers: frequency-dependent keeps speed unchanged. $\beta = 0.25$, 8 days: $R_0 = 2$, threshold $S = 50\%$. Euler: $S \cdot (1 - \beta I \Delta t) < 0$ when $\beta I \Delta t > 1$. Vaccine 90%: $80\%/0.9 \approx 89\%$; $R_0 = 15$ needs 104%, impossible. Isolation share: $1.875 - 0.3 \times 0.7 \times 4 \times q = 1$ gives $q \approx 1.04$: even full isolation after one day leaves R_0 just above 1.

Answers: Parts 3–4

- **8.** 1.25 / 0.75 ≈ 1.67.
- **9.** $(1/4)^2 \approx 6\%$.
- **10.** It suppresses transmission while almost everyone is still susceptible; when it ends, the epidemic resumes from a large susceptible pool.
- **11.** Eigenvalues 1.5 and 0.9: $R_0 = 1.5$.
- **12.** Everyone who enters E reaches I (without deaths in E), so σ only changes when, not how many.
- **13.** $1.5^2 = 2.25$.

TRY-IT ANSWERS

Try-it answers: R(t) = 2.5 × 0.5 × 0.8 = 1.0. With deaths: $R_0 = \beta\,\sigma / ((\sigma + \mu)(\gamma + \mu))$. K = [[2, 1], [0.5, 1]]: trace 3, determinant 1.5, $R_0 = (3 + \sqrt{3})/2 \approx 2.37$; the child–child-like entry 2 matters most.

CREDITS

Sources and credits

- This deck accompanies Chapter 3 of T. Lazebnik, Epidemic Modeling: Principles, Algorithms, and Applications (Academic Press, 2026), the course textbook. It follows the chapter's topics in our own words; no text, figures, tables, exercises or code are copied from the book.
- The text, figures, examples and simulations were written, drawn and run for ACML; nothing is copied from the book. Numbers labelled "our run" come from our own code; counts from agencies and papers are credited where they appear.
- All numbers, curves and toy calculations are our own runs (production/explains/sims/epi_m03_sims.py), chosen to show mechanisms, not to estimate any real disease.
- Historical notes (measles and school terms, critical community size, the Ross–Macdonald formula, Reed–Frost) are standard results of the field, described in our own words.
- Companion videos: **Epidemic Modeling**, module 3, episodes 7–12. More from the lab: acml.teddylazebnik.com

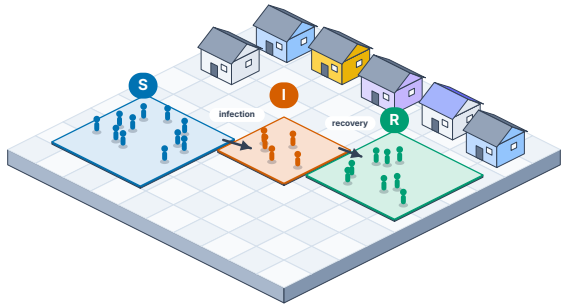

Applied Computational Mathematics Lab

Module 3 complete

END OF MODULE 3

Compartmental models

YOU COVERED

- SIR from scratch and the meaning of R_0
- More states: births, deaths, reinfection, latency
- More classes: stages, outcomes, structure, vectors
- When transmission changes
- Chance matters: stochastic models
- The next-generation matrix



UP NEXT · MODULE 4

Spatio-temporal models

Cities, networks and continuous space

Based on Chapter 3 of T. Lazebnik, Epidemic Modeling: Principles, Algorithms, and Applications (Academic Press, 2026)