Applied Computational Mathematics Lab

EPIDEMIC MODELING · MODULE 5

Multi-strain and evolving pandemics

Strains side by side, immune history, and mutation



3 EPISODES

1. Strains side by side
2. Reinfection and immune history
3. Mutation and the emergence of new strains

Based on Chapter 5 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) Two strains with the same R_0 make an epidemic twice as big as one strain.

(b) A fitter strain (higher R_0) always takes over.

(c) If you had strain A, you are fully protected against strain B.

(d) In a deterministic model, a new strain exists as soon as mutation is allowed.

HOW TO USE THIS

Guess now; each part settles one or two statements. Answers at the end.

WHY THIS MODULE

Until now the pathogen was one fixed type. Here several strains circulate, immunity depends on which ones you had, and new strains appear during spread.

BEFORE WE START

This module's map

PART 1

Strains side by side

A shared pool of susceptibles, replacement, timing, and reading a variant's advantage.

PART 2

Immune history

Who had which strain, cross-immunity, the invasion threshold, waves of strains, shortcuts.

PART 3

Mutation

A mutation matrix, emergence in the equations and by chance, the supply of variants, diversity.

GOALS

By the end you should be able to write a multi-strain SIR model, explain when a new strain can invade, and estimate how often successful variants appear.

ONE WORD

Strain: a variant of the pathogen that we track separately, because it differs in transmission, severity or the immunity it meets.

Applied Computational Mathematics Lab

PART 1

Strains side by side

Competition for the same people.



PART 1 · PARALLEL STRAINS

One pool, several strains

$$\frac{dS}{dt} = -S\sum_k \beta_k I_k, \qquad \frac{dI_k}{dt} = \beta_k S I_k - \gamma_k I_k, \qquad \frac{dR}{dt} = \sum_k \gamma_k I_k$$

- Each strain k has its own infected box and rates; all strains infect from one susceptible pool S.
- Strong assumption: recovering from any strain protects against all (one R box).
- Strains interact only through S. Their effective reproduction numbers are $R_{e,k} = (\beta_k/\gamma_k)\, S(t)$: when one strain uses up susceptibles, every strain slows.

IDENTICAL STRAINS (R_0 2 EACH)	FINAL SIZE	PEAK	PEAK DAY
1	80%	15.4%	68
2	80%	15.4%	68
3	80%	15.4%	68
4	80%	15.4%	68

Our runs: the same total seed split over 1–4 identical strains.

IN PLAIN WORDS

A strain label does not create new susceptibles. Identical strains share one epidemic.

TRY IT

With an exposed stage per strain (parallel SEIR), how many boxes does a model with K strains have?

Timing can beat fitness

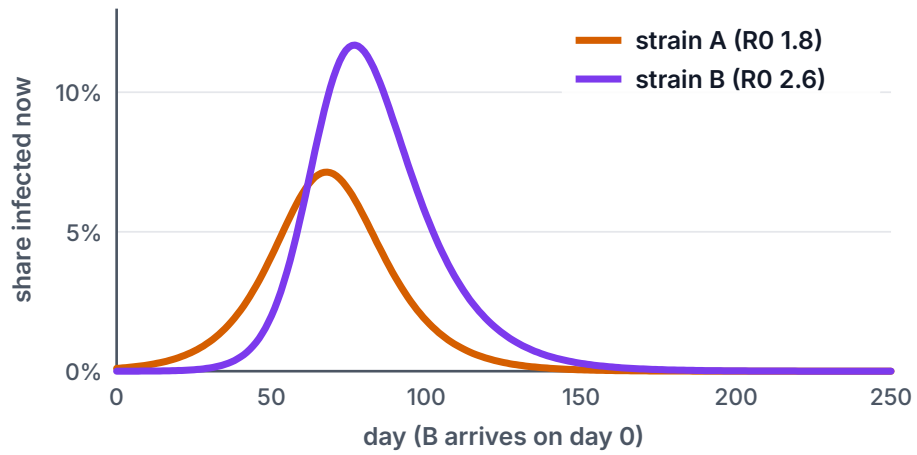

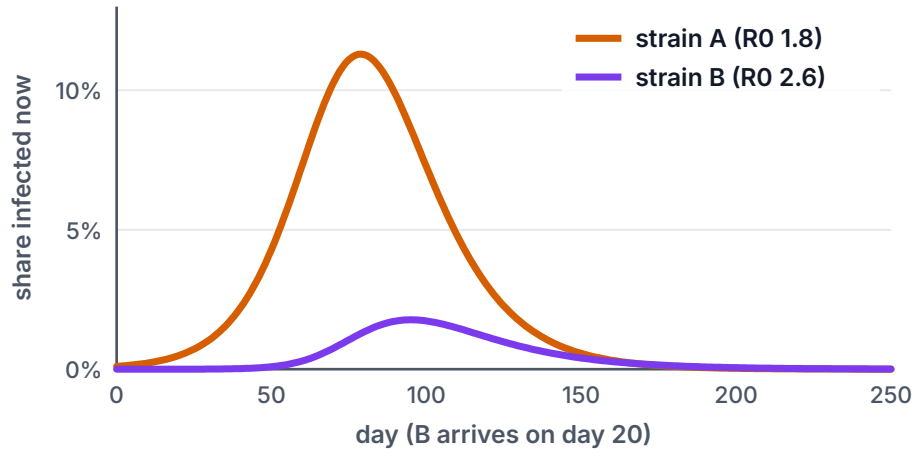

B ARRIVES ON	B'S EFFECTIVE NUMBER THEN	B'S SHARE OF INFECTIONS	TOWN INFECTED
day 0	2.60	60.5%	86%
day 20	2.58	14.5%	77%
day 40	2.47	1.0%	74%
day 60	2.10	0.1%	73%
day 80	1.46	0.0%	73%
day 120	0.80	0.0%	73%

Our runs: A (R_0 1.8) seeded at 1 in 1,000 on day 0; B (R_0 2.6) arrives as 1 in 100,000.

- B is fitter, but arriving late with a small seed it misses its window: A's wave ends before B can grow.
- Even on day 60, B's effective number is above 2, yet B causes 0.1% of infections.

Reading a variant's advantage

$$\frac{d}{dt}\ln\frac{p_B}{1-p_B} = (\beta_B - \beta_A)\,S(t) - (\gamma_B - \gamma_A)$$

- p_B: B's share among the infected. While S stays near 1, the log odds rise in a straight line.
- Our run: slope 0.08 per day, exactly $\beta_B - \beta_A$; B passes half of all infections on day 58.
- This is how a variant's advantage is read from data: the share of sequenced cases, on a log-odds scale.

DATA NEEDED

Case counts cannot show it: you need the strain of a sample of cases (sequencing, Module 9).

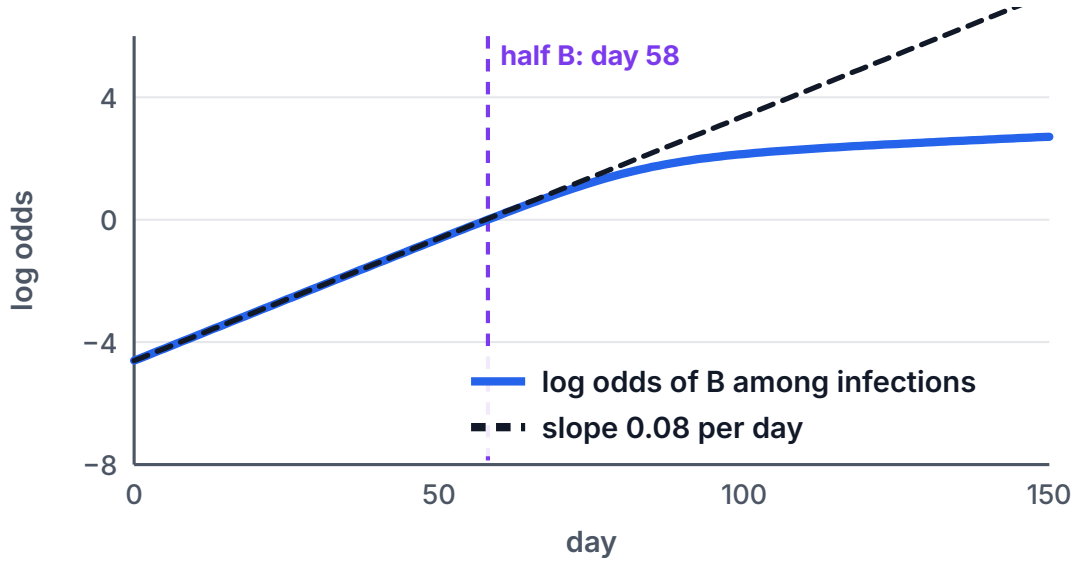

Our run: A at 1 in 10,000, B at 1 in 1,000,000 on day 0, γ = 0.1 for both.

TRY IT

Derive the equation above from dI_A/dt and dI_B/dt.

Hint: differentiate $\ln I_B - \ln I_A$.

Replacement and coexistence

- With one shared immunity and no new susceptibles, there is one epidemic: strains split it, and the faster one takes the larger share.
- Add births or waning immunity and the epidemic keeps going. Then, with full cross-protection, the strain with the largest R_0 wins in the long run: it can grow at a susceptible level where the others shrink (competitive exclusion).
- Partial cross-immunity changes the picture: strains can coexist, alternate or oscillate (for example Andreasen, Lin and Levin, 1997).

IN PLAIN WORDS

Who wins depends on the immunity the strains create for each other, not only on their own R_0.

WHY THE LARGEST R_0 WINS

In an endemic setting a strain persists only if it can replace itself when susceptibles sit at the level set by the others: $1/R_0$ of the dominant strain. A strain with a smaller R_0 cannot.

IN PRACTICE

Seasonal influenza: several subtypes co-circulate for years, which full cross-protection would not allow. Partial, strain-specific immunity is one reason.

TRY IT

Strain A has R_0 4 and strain B R_0 3, with full cross-protection and births. At the endemic balance of A, what is B's effective number?

Own clocks and own severity

- Strains can differ in more than transmission: latent period (an exposed box per strain), infectious period, severity, and how well tests detect them.
- A parallel SEIR model with K strains has S, R, and an exposed and an infected box per strain: 2K + 2 boxes.
- A strain that is more transmissible per day is not always more transmissible per generation: a shorter generation time also raises the daily growth rate.

READING THE SLOPE

Our slope of 0.08 per day multiplies B's odds by $e^{0.08 \times 7}$, about 1.75, every week. Reports often quote this weekly factor.

ODDS

Odds of B = share of B / share of A. A straight line of log odds against time means the odds grow by a constant factor per day.

CAREFUL

Two strains with the same R_0 but different generation times grow at different speeds: R_0 alone does not decide which strain rises faster.

TRY IT

A variant's odds double every 5 days. What is the slope of the log odds per day?

Check yourself: Part 1

CHECK YOURSELF

1. Strain A has R_0 2 and strain B R_0 3, both with $\gamma = 0.1$ per day. What is the early slope of the log odds of B's share?

2. Why does adding a third identical strain not change the final size?

3. Strain B (R_0 3) arrives when 70% of the town is already immune (to all strains). What is its effective number, and can it spread?

TIP

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

Applied Computational Mathematics Lab

PART 2

Immune history

Who had which strain, and how protected that makes them.



Histories as sets

- Track, for each person, the set J of strains already had. R_J: not infected now, history J. $R_J I_i$: infected now with strain i (i not in J).
- $R_\varnothing$ plays the role of S: no strain yet. Recovering from i moves $R_J I_i$ to $R_{J \cup \{i\}}$.
- History can change three things: the chance of infection, how sick a person gets, and how infectious they are.

STRAINS M	HISTORIES 2^M	INFECTED BOXES $M \cdot 2^{M-1}$	PARALLEL MODEL M + 2
1	2	1	3
2	4	4	4
3	8	12	5
4	16	32	6
6	64	192	8
10	1,024	5,120	12

TWO STRAINS

For m = 2: histories none, A, B, both; infected boxes I_A from none or B, I_B from none or A. Four of each.

THE PRICE

Faithful histories grow exponentially with the number of strains. Part 2 ends with the shortcuts used in practice.

TRY IT

Count the histories and infected boxes for 5 strains. How many would a parallel SEIR model need?

The history model in equations

$$\frac{dR_J I_i}{dt} = \beta_{J,i}\, R_J \sum_{K \not\ni i} R_K I_i \;-\; \gamma_{J,i}\, R_J I_i$$

$$\frac{dR_J}{dt} = \sum_{i \in J} \gamma_{J\setminus\{i\},i}\, \phi_{J\setminus\{i\},i}\, R_{J\setminus\{i\}} I_i \;-\; \sum_{i \notin J} \beta_{J,i}\, R_J \sum_{K \not\ni i} R_K I_i$$

- Infection pressure of strain i comes from everyone now infected with i, whatever their own history.
- The recipient's history J sets its susceptibility ($\beta_{J,i}$); recovery adds i to the history; ϕ is the share who survive.

OUR SPECIAL CASE

All quantities are shares of a closed population. With $\beta_{J,i} = \beta_i(1 - \sigma)$ for any non-empty J that lacks i, this is the model behind our runs.

IN CODE

Code stores histories as bitmasks: J = 0b0101 means strains 1 and 3 already had. Recovery from strain i turns J into J OR (1 << i).

PART 2 · IMMUNE HISTORY

Cross-immunity

$$R_{e,B} = R_{0,B}\left[R_{\varnothing} + (1-\sigma)\,R_{\{A\}}\right]$$

- σ: how much an earlier infection protects against a different strain (0: none, 1: full).
- The bracket is B's effective susceptible share: everyone who never had B, weighted by protection.
- Our town: A (R_0 2.5) leaves 10.7% never infected and 89.2% who had A by day 150. B (R_0 2.5) arrives then.
- B can invade only if $\sigma < \sigma^*$ = 0.672.

σ	$R_{E,B}$ AT ARRIVAL	B INFECTS	B'S PEAK
0.0	2.50	89.3%	23.4%
0.3	1.83	73.7%	12.2%
0.5	1.38	47.3%	4.0%
0.6	1.16	22.9%	0.9%
0.7	0.94	0.1%	0.0%
0.8	0.71	0.0%	0.0%
1.0	0.27	0.0%	0.0%

Our runs, 800 days.

TRY IT

Derive σ* from $R_{e,B}$ = 1. Check it against the table.

PART 2 · IMMUNE HISTORY

A second strain arrives

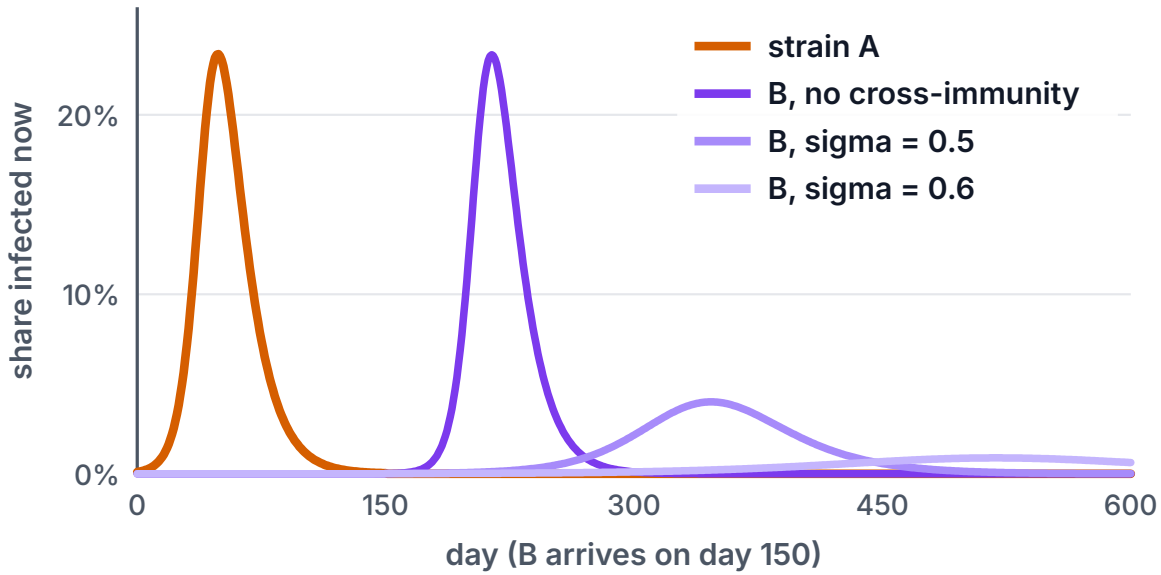

Our runs: A (R_0 2.5) from day 0; B (R_0 2.5) arrives on day 150 at 1 in 10,000.

- No cross-immunity: B infects 89% again, in a wave as high as the first.
- σ = 0.5: 47% in a lower, slower wave (peak on day 347).
- σ = 0.6: a creep of more than a year, 23% in the end.
- σ = 0.7 and above: B never takes off.

IN PLAIN WORDS

Close to the threshold, waves are slow and long: small changes in protection make a big difference in timing.

Severity and infectiousness by history

- A second infection is often milder: in the model, a lower death share (higher ϕ) or a shorter infection (higher γ) for histories that are not empty.
- Prior infection can also cut infectiousness: then β depends on the infector's history too, not only on the recipient's.
- Waves driven by reinfection can be large in infections but small in hospital and death counts: what you measure decides what you see.

TOY ARITHMETIC

Two waves of equal size in infections: if second infections are half as severe, the second wave brings half the hospital admissions.

HONEST MODELING

A model where history changes susceptibility, severity and infectiousness has three sets of parameters per history and strain. Data rarely pin them all down: report which ones you assumed.

TRY IT

A first wave infected 60% and killed 0.5% of those infected. A second strain reinfects 40%, with a death share half as large among people infected before. How many deaths per 100,000 does the second wave cause, if 10% of it hits never-infected people?

PART 2 · IMMUNE HISTORY

Waves of strains

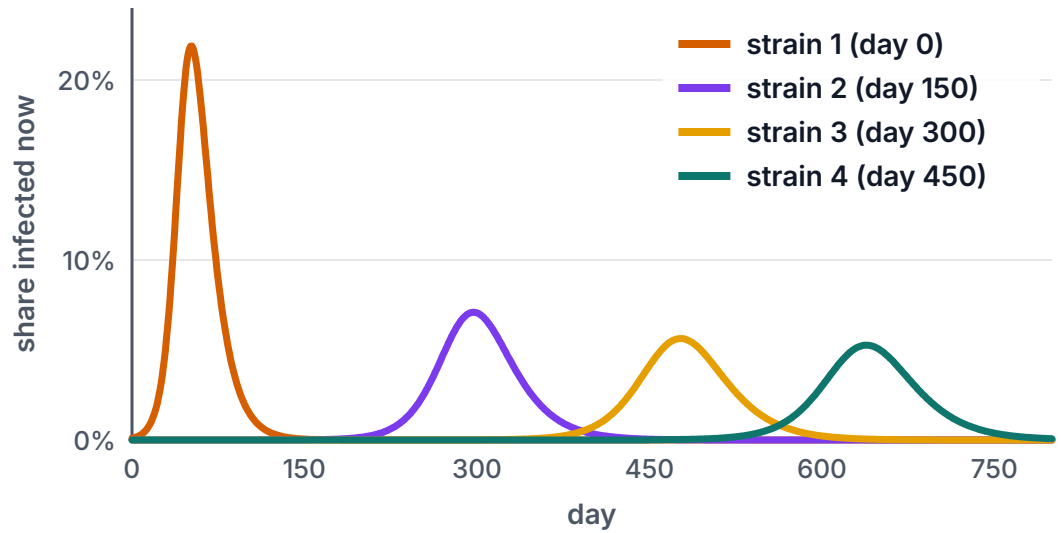

Our run: four strains (R_0 2.4 each), arriving on days 0, 150, 300 and 450; σ = 0.4.

- Wave peaks: 21.9%, 7.1%, 5.7%, 5.3% infected at the same time.
- Each new strain finds people its predecessors left only partly protected.

BY DAY 800	NEVER	ONCE	TWICE	3 TIMES	4 TIMES
share of the town	0.2%	11.1%	34.6%	39.0%	15.1%

IN PLAIN WORDS

Immune history turns one epidemic into a sequence of waves: one ingredient behind repeated influenza seasons and variant waves (seasons and behaviour are others).

Shortcuts when histories explode

- **Status-based models** track how protected a person is, not every history (Gog and Grenfell, 2002, among others).
- **Antigenic space** places strains on a map; protection falls with distance (antigenic cartography: Smith et al., 2004).
- **Effective susceptible share per strain**: $S_i^{\text{eff}} = \sum_J s_{J,i} R_J$, with $s_{J,i}$ the susceptibility of history J to strain i. Then $R_{e,i} \approx R_{0,i}\, S_i^{\text{eff}}$.

IN PLAIN WORDS

A single S can be nearly used up while S^{eff} for a new strain is still large: that is how a variant can surge in a well-infected population.

LINK TO THE THRESHOLD

With $s_{J,i}$ = 0 if i is in J, 1 if J is empty and 1 − σ otherwise, the threshold above is the special case of two strains.

TRY IT

20% never infected, 80% had A; B has R_0 2 and σ = 0.5. Compute S^{eff}_B and $R_{e,B}$.

Check yourself: Part 2

CHECK YOURSELF

4. How many histories and how many infected boxes does a full immune-history model with 5 strains have?

5. 20% of a town were never infected and 80% had A. B has $R_0 = 2$ and cross-immunity $\sigma = 0.5$. Can B invade?

6. In question 5, what is the largest σ that still lets B invade?

TIP

Answers are at the end of the deck.

Applied Computational Mathematics Lab

PART 3

Mutation and emergence

New strains made during spread.



The mutation matrix

$$\Lambda_{J\to j} = \sum_i \beta_i \, I_i^{\mathrm{tot}} \, q_{i\to j}, \qquad \sum_j q_{i\to j} = 1$$

- $q_{i\to j}$: chance that a transmission by strain i founds strain j. A ladder: $q_{i\to i} = 1-\mu,\ q_{i\to i+1} = \mu$.
- Rows sum to one, so without immunity Q only relabels infections: the total pressure is unchanged.
- With immunity (no infection by a strain already in J), relabelling matters: a new strain can reach people the old one cannot.

OTHER KERNELS

A ladder is the simplest kernel. Others: a mutant can jump several steps, mutation can go back, or Q can follow distances in antigenic space.

WHERE MUTATION HAPPENS

Mutation here happens at transmission. Within-host evolution during long infections is another route the simple Q leaves out.

TRY IT

With rows summing to one, show that the total infection pressure on a person with no immunity does not depend on Q.

PART 3 · MUTATION

A ladder of strains

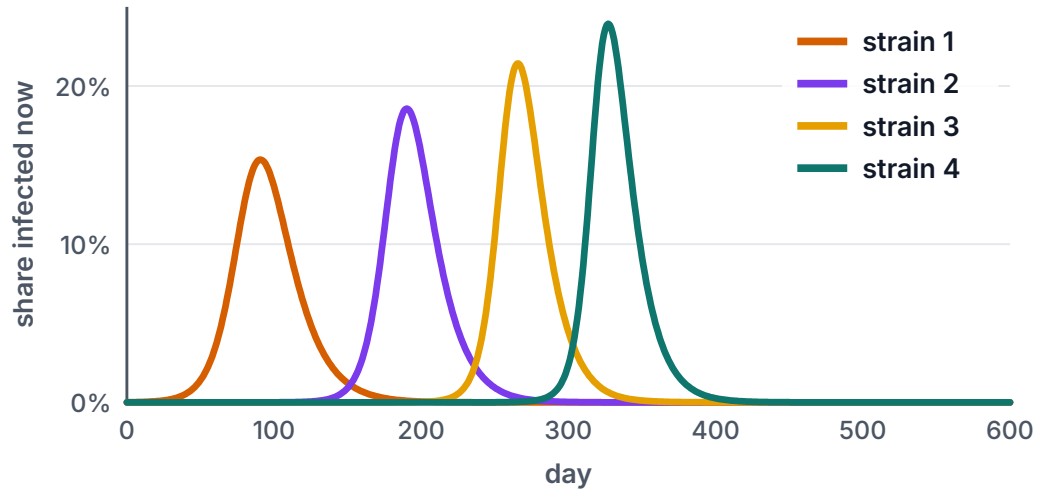

Our run: one mutation in a million transmissions; four separate waves.

MUTATION RATE M	STRAIN 2 REACHES 1 IN 10,000	STRAINS ABOVE 0.1%	SEPARATE WAVES
none	never	1	1
1 in a million	day 109	4	4
1 in 100,000	day 84	4	3
1 in 10,000	day 62	4	2
1 in 1,000	day 44	4	1

Our runs: four strains (R_0 2, 2.2, 2.4, 2.6), each escaping immunity to the earlier ones; only strain 1 at the start.

- More mutation brings strains sooner, and their waves merge: at 1 in 1,000 it is one long wave.
- Without cross-immunity people are infected again and again: about 3.4 infections per person over 900 days in these runs.

PART 3 · MUTATION

Merged waves at high mutation

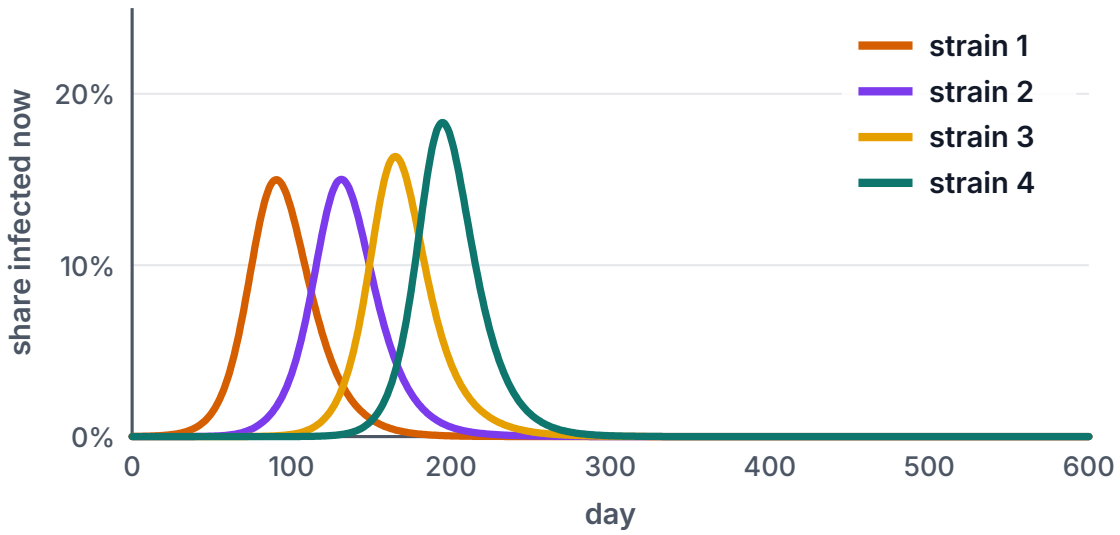

Our run: one mutation in 1,000 transmissions.

- When new strains arrive faster than a wave lasts, their epidemics overlap and the total looks like one long wave.
- Counting waves is therefore a crude summary: the same number of strains can show as one wave or four.
- Reporting strain shares over time (from sequencing) reveals the turnover that total cases hide.

IN PLAIN WORDS

Total cases can look smooth while the strain mix underneath changes completely.

Emergence: smooth in equations, random in reality

- In the equations, each mutant strain exists from day 0 as a tiny fraction and grows deterministically once its effective number exceeds 1.
- So the wait for the next strain grows only with log(1/μ): about 24 days per factor of ten in our runs.
- Real emergence starts from one infected person. Early on, a new strain is a branching process and survives with chance about $1 - 1/R_e$ (Module 3).

IN PLAIN WORDS

The equations describe what a variant does once it is established; chance decides whether and when it establishes.

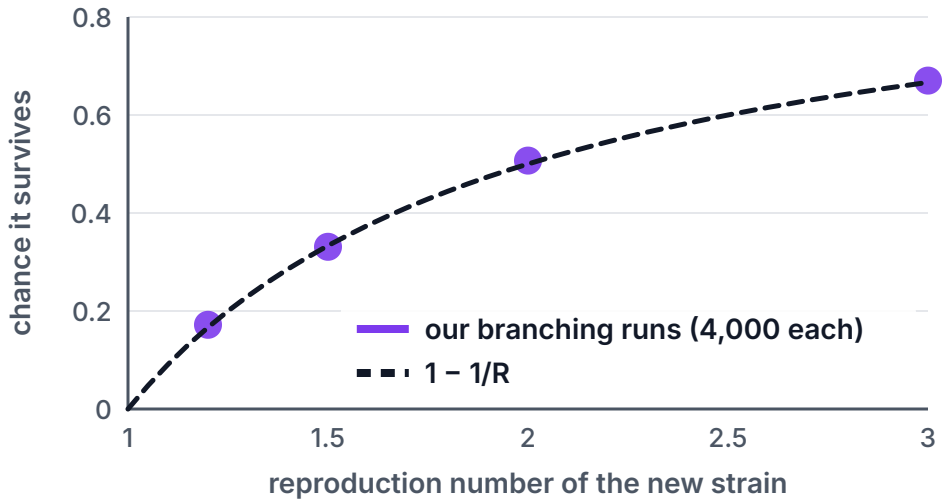

Our branching runs: geometric offspring with mean R, 4,000 chains per point; a chain counts as surviving above 200 cases.

R	1.2	1.5	2.0	3.0
survived	17.2%	33.1%	50.7%	67.0%
1 − 1/R	16.7%	33.3%	50.0%	66.7%

The supply of variants

$$\lambda_{\text{succ}} \approx \mu \cdot C(t) \cdot p_{\text{est}}, \qquad \Pr(\text{at least one}) = 1 - e^{-\mu Z p_{\text{est}}}$$

- C(t): transmissions per day; Z: total transmissions over the epidemic; p_{est}: survival chance of a mutant.
- Expected waiting time for a successful variant (roughly constant C): 1/(μ C p_{est}).

EPIDEMIC IN A TOWN OF A MILLION	INFECTIONS	CHANCE OF A VARIANT
R_0 1.3 (controlled)	422,970	19.1%
R_0 2.0 (partly controlled)	796,812	32.9%
R_0 3.0 (uncontrolled)	940,480	37.5%

Our arithmetic: μ = 1 in a million, p_{est} = 0.5 (a variant with R_e 2).

IN PLAIN WORDS

Fewer infections mean fewer chances to mutate: control also slows evolution.

TRY IT

1,000 transmissions a day, μ = 1 in 100,000 and p_{est} = 0.5. What is the average wait for a successful variant?

CAVEAT

Long infections in a few people (for example with weak immune systems) can matter more than many short ones; the simple formula ignores that.

Measuring diversity

$$H(t) = -\sum_i p_i \ln p_i, \qquad N_{\text{eff}} = e^H, \qquad 0 \le H \le \ln m$$

- p_i: share of strain i among the infected. One strain dominating: H = 0 and N_{eff} = 1; m equally common strains: N_{eff} = m.
- N_{eff} reads as “the number of equally common strains with the same diversity”.
- Waves can be counted from peaks of total incidence with a height floor and a minimum gap: crude, but comparable across runs.

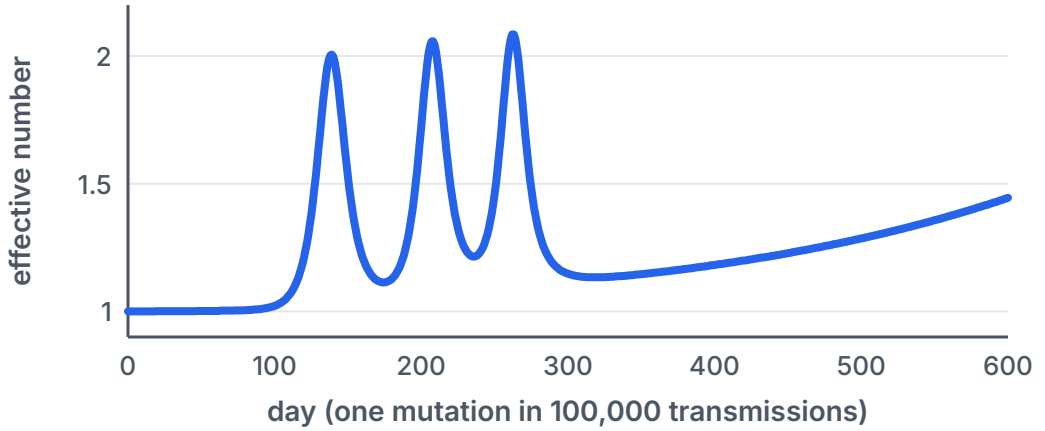

Our run (one mutation in 100,000 transmissions): N_{eff} rises to about 2 at each handover between strains.

BEYOND THE LADDER

Real evolution is messier: transmission gains and immune escape come together, strains recombine, and selection depends on the immunity left by earlier waves.

Forecasting an evolving pathogen

- Short term: once a variant is spreading, its share often follows a logistic trend: the log-odds slope from Part 1 can be extrapolated for a few weeks.
- Long term: which variant comes next, and when, is not predictable from these models alone; use scenarios (for example, a variant with immune escape arriving in month 6).
- Every scenario needs the immunity landscape: who is protected against what, from infections and vaccines.

IN PLAIN WORDS

A model can say how much a new strain would matter, given its properties; it cannot promise when such a strain appears.

TWO KINDS OF OUTPUT

Forecast: a statement about what will happen, with uncertainty. Scenario: a conditional statement, “if a variant like this arrives, then…”. Module 11 returns to this.

WHAT CONTROL BUYS

Our supply arithmetic shows the lever a model can give: fewer infections lower the chance of a new variant, whatever it turns out to be.

Check yourself: Part 3

CHECK YOURSELF

7. $\mu = 1$ in 100,000 per transmission, 200,000 transmissions, $p_{est} = 0.4$.

Expected number of successful mutants? Chance of at least one?

8. Strain shares are 0.5, 0.25 and 0.25. Compute H and N_{eff}.

9. Why does a deterministic model make emergence look more predictable than it is?

TIP

Answers are on the next slides.

Answers: warm-up and Parts 1–2

- **(a)** False: identical strains share one epidemic. **(b)** False: a late arrival can miss its window. **(c)** False: protection depends on cross-immunity. **(d)** True, as a tiny fraction: that is why chance must be added for realistic emergence.
- **1.** $\beta_A = 0.2$, $\beta_B = 0.3$: slope 0.1 per day. **2.** The strains draw from the same susceptibles; the total behaves like one SIR model.
- **3.** $3 \times 0.3 = 0.9$: below 1, it cannot spread.
- **4.** $2^5 = 32$ histories and $5 \times 2^4 = 80$ infected boxes.
- **5.** $R_{e,B} = 2 \times (0.2 + 0.5 \times 0.8) = 1.2$: yes, slowly. **6.** $2\,(0.2 + 0.8\,(1 - \sigma)) = 1$ gives $\sigma = 0.625$.

TRY-IT ANSWERS

Try-it answers: parallel SEIR has 2K + 2 boxes. Log odds: d/dt (ln I_B − ln I_A) = (β_B S − γ_B) − (β_A S − γ_A). Five strains: 32 histories, 80 infected boxes, versus 12 boxes for parallel SEIR. $S^{eff}_B = 0.6$, $R_{e,B} = 1.2$. $\sigma^* = 1 - (1/R_{0,B} - R_\varnothing)/R_{\{A\}}$. Endemic A: S = 1/4, so B's effective number is 3/4: it cannot persist. Doubling every 5 days: slope ln 2 / 5, about 0.14 per day. Second wave: 0.4 × [0.1 × 0.5% + 0.9 × 0.25%] = 0.11%, about 110 deaths per 100,000.

Answers: Part 3

- **7.** $10^{-5} \times 200{,}000 \times 0.4 = 0.8$ expected; chance of at least one: $1 - e^{-0.8}$, about 55%.
- **8.** $H = -(0.5 \ln 0.5 + 2 \times 0.25 \ln 0.25) = 1.04$; $N_{eff} = e^{1.04}$, about 2.83.
- **9.** The equations hold every strain at a tiny fraction from the start, so the timing depends only on growth; in reality the first mutants usually die out by chance, and which strain emerges, and when, varies from run to run.

TRY-IT ANSWERS

Try-it answers: rate $10^{-5} \times 1{,}000 \times 0.5 = 0.005$ per day, so about 200 days on average. With rows summing to one, the pressure $\sum_j \sum_i \beta_i I_i q_{i \to j} = \sum_i \beta_i I_i$.

CREDITS

Sources and credits

- This deck accompanies Chapter 5 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_m05_sims.py), chosen to show mechanisms, not to estimate any real pathogen.
- Standard ideas described in our own words: competition for susceptibles and strain replacement, cross-immunity and invasion thresholds, status-based reductions (Gog and Grenfell, 2002), antigenic cartography (Smith et al., 2004), branching-process establishment.
- Companion videos: **Epidemic Modeling**, module 5, episodes 17–19. More from the lab: acml.teddylazebnik.com

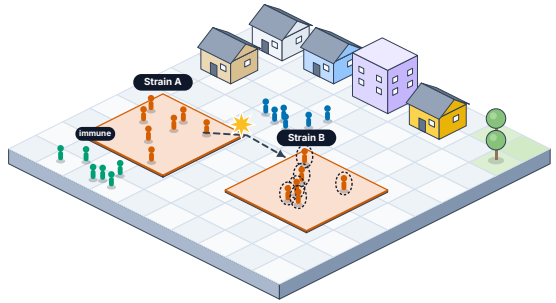

Applied Computational Mathematics Lab

Module 5 complete

END OF MODULE 5

Multi-strain and evolving pandemics

YOU COVERED

- Strains side by side
- Reinfection and immune history
- Mutation and the emergence of new strains



UP NEXT · MODULE 6

Multi-agent simulation

Building, running and trusting agent-based models

Based on Chapter 5 of T. Lazebnik, Epidemic Modeling: Principles, Algorithms, and Applications (Academic Press, 2026)