Applied Computational Mathematics Lab

EPIDEMIC MODELING · MODULE 6

Multi-agent simulation

Building, running and trusting agent-based models

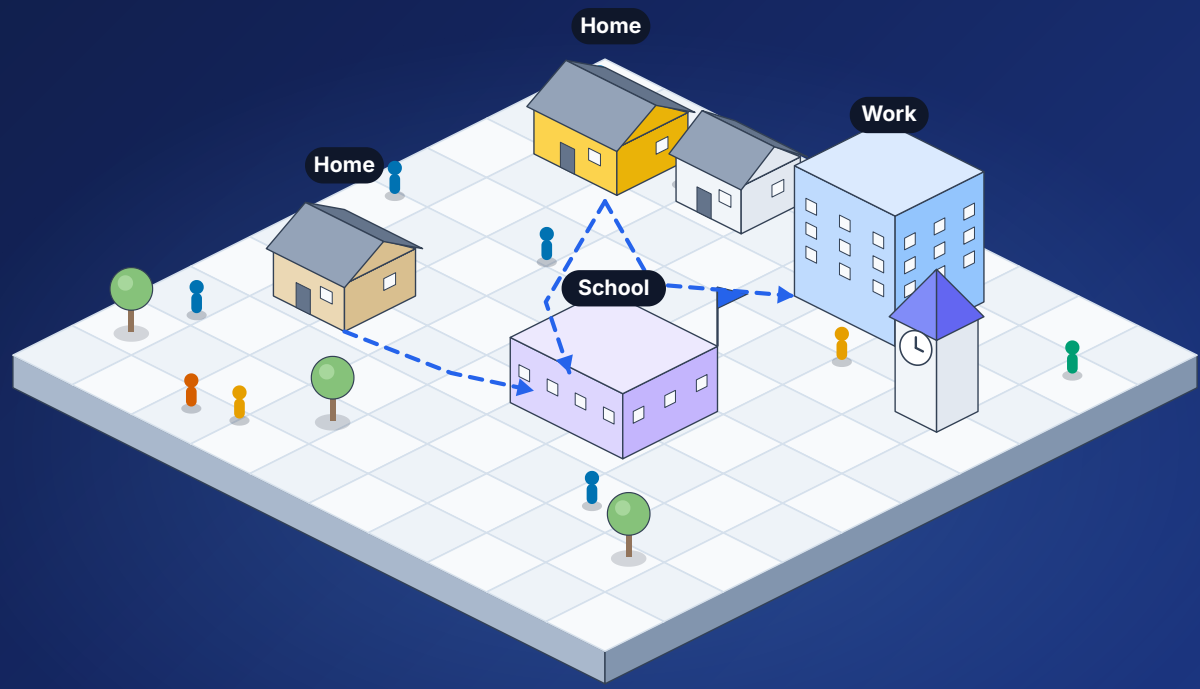

6 EPISODES

1. Why agents? Their power and their cost
2. Building an agent-based model
3. Randomness and reproducibility
4. Two case studies: flu-like illness and COVID-19
5. Designing your own simulation
6. Sensitivity, ethics and policy risk

Based on Chapter 6 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) An agent model where everyone is average and mixes freely behaves differently from the SIR equations.

(b) The time step of an agent model is a technical detail with no effect on results.

(c) Updating agents one at a time, always in the same order, is harmless.

(d) Two runs with the same seed give exactly the same epidemic.

HOW TO USE THIS

Guess now; the module settles each one. Answers at the end.

WHY THIS MODULE

Until now people were shares in boxes. Here every person is simulated: a state, rules, neighbours. That buys heterogeneity and explicit contacts, at a price in design effort, computing time and noise.

BEFORE WE START

This module's map

PART 1

Why agents

What an agent is, agents versus equations, heterogeneity, what agents can compute, and their cost.

PART 2

Building blocks

Time steps and events, update order, geometry, random numbers, common random numbers, richer agents.

PART 3

In practice

Two case studies, seven design steps, sensitivity analysis, ethics and policy risk.

GOALS

By the end you should be able to design a small agent-based epidemic model, choose its time, order and geometry on purpose, make it reproducible, compare policies efficiently, and report its uncertainty honestly.

ONE PHRASE

Agent-based simulation: a model made of many individual units, each with its own state and rules, whose interactions produce the population-level outcome.

Applied Computational Mathematics Lab

PART 1

Why agents

Their power and their cost.



What an agent is

- **State**: health state (S, E, I, R) plus anything the question needs: age, household, a timer, a belief.
- **Rules**: what the agent does each step: meet people, infect or be infected, recover, change behaviour.
- **Neighbours**: who it can meet: everyone (well mixed), a list of contacts (a network), or the cells around it (a grid).
- **Scheduler**: when agents act: all at once each step, one by one, or event by event.

IN PLAIN WORDS

Nobody writes the epidemic curve down. It comes out of the rules, which is why the same model can show surprises the equations do not.

ONE DAY OF OUR SIMPLEST AGENT MODEL

1. for every infectious agent: pick 10 people at random
2. each picked person who is susceptible becomes infected with chance 2.5%
3. each infectious agent recovers with chance 10%
4. apply all changes together, then record the counts

LINK TO SIR

Expected new infections per infectious agent per day: 10 × 2.5% = 0.25 = β; recovery chance 0.1 = γ; so R_0 = 2.5.

TRY IT

Each agent meets 12 people a day and recovers with chance 1/8 per day. What chance per contact gives R_0 = 3?

PART 1 · AGENTS

Everyone average: agents give back SIR

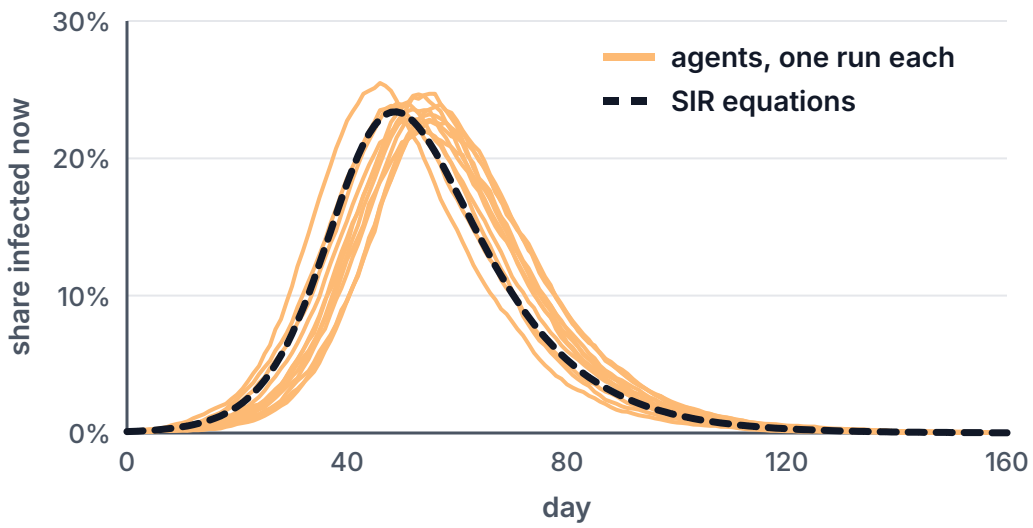

Our runs: 10,000 agents, 10 first cases, 12 of 40 runs shown.

	AGENTS	EQUATIONS
share infected in the end	89.3% (average of 40)	89.3%
peak day	46–60	49
runs that fizzled	0	–

- With homogeneous agents and random mixing, the average run follows the equations; single runs scatter around them.
- The scatter shrinks as the town grows: noise is the price of counting people one by one.
- Start from one case instead of ten and many runs would fizzle (Module 3).

PART 1 · AGENTS

Nobody is average

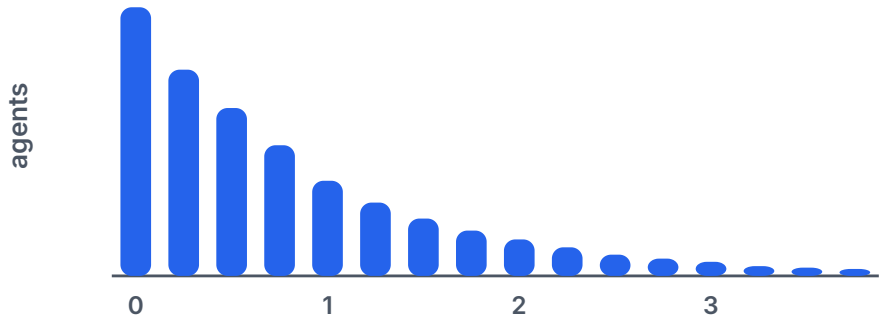

Susceptibility of 10,000 agents: spread out (gamma distribution, coefficient of variation 1), average 1. Bins of 0.25.

SAME AVERAGE SUSCEPTIBILITY	SHARE INFECTED
everyone at 1 (the SIR town)	89%
spread out	60%
least susceptible fifth	14%
most susceptible fifth	96%

Our runs: 20 runs, same contacts and recovery as before.

- The most susceptible are infected first; the wave then meets a less susceptible population and stops sooner.
- Agents keep who is who. A compartment model would need extra boxes to do the same.

TRY IT

Why would spreading out infectiousness alone (who infects, not who is infected) change the chance of take-off more than the final size?

Agents can compute anything

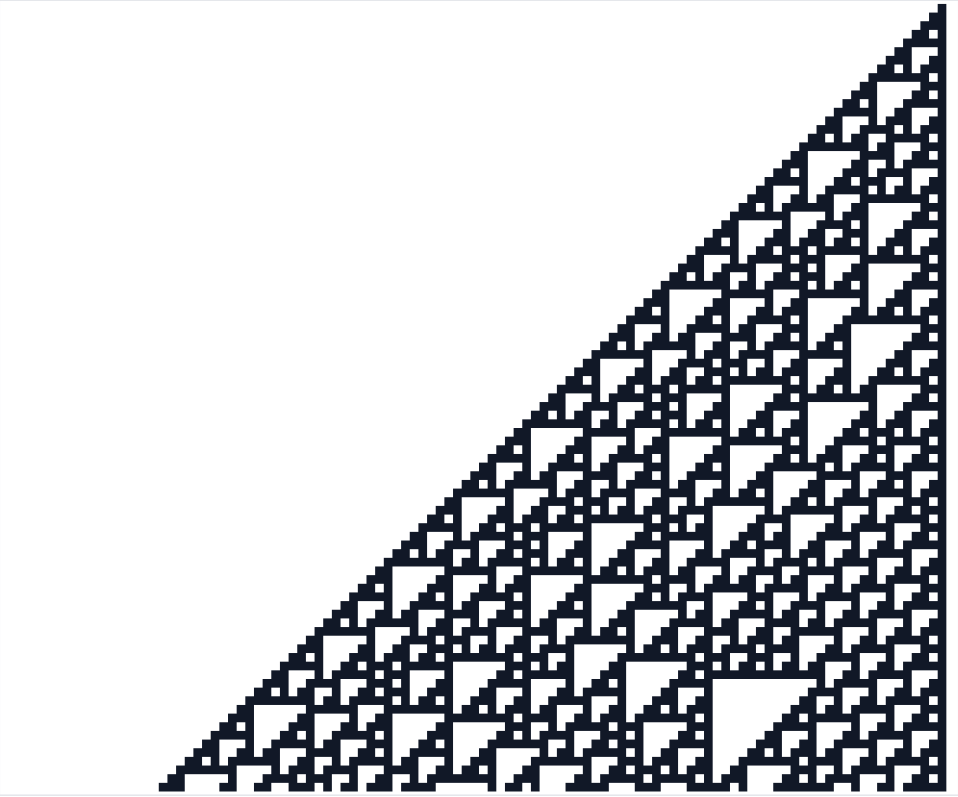

Our run of Rule 110 from one live cell; time runs downwards, 90 steps.

- A one-dimensional cellular automaton: each cell is 0 or 1 and updates from itself and its two neighbours with a fixed table. Rule 110 is one such table.
- Cook (2004) proved that Rule 110 can carry out any computation, given the right start and enough room. Treat each cell as an agent: agent models can be Turing complete.
- Universality needs unbounded room (agents or space). A fixed number of agents with finite states must eventually repeat itself.

IN PLAIN WORDS

The price of that power: no general method can decide whether an arbitrary agent model will stop (the halting problem). In practice we reason about the restricted models we actually build.

PART 1 · AGENTS

The cost of contacts

every pair: $\Theta(TN^2)$ k contacts: $\Theta\left(k\sum_t I_t\right)$ network: $\Theta\left(\sum_t \sum_{a\in I_t} \deg a\right)$

OUR TOWN OF 10,000, ONE EPIDEMIC	CHECKS
every pair, every day	7.9 billion
10 random contacts per infectious agent	891,326
ratio	8,830

Our counts: 157 days, 89,132 infectious agent-days.

- Memory: agent states need Θ(N); a stored network adds Θ(|E|). Event-driven schedulers add a queue.

AGENTS	PAIRS PER DAY	CONTACT CHECKS PER RUN
1,000	499,500	89,132
10,000	50.0 million	891,326
100,000	5.0 billion	8.9 million
1,000,000	500.0 billion	89.1 million

Larger towns: our counts scaled with N, assuming the same epidemic shape.

IN PLAIN WORDS

Every pair grows with N^2, contacts with N. Design the contact loop first: it usually dominates the running time.

TRY IT

A network of 1 million people with mean degree 20, 30% infected over 100 days, 8 infectious days each. How many edge checks?

What to report

- Population size N and the peak number infectious.
- Contact structure: well mixed with k contacts, a network with |E| edges, or a grid with M cells and a neighbourhood.
- Scheduler: synchronous steps, random sequential, or event-driven; and the time step.
- The dominant loops and their scaling, and peak memory in terms of N, |E|, M and the event queue.
- Number of runs and how seeds were set.

IN PLAIN WORDS

These few lines tell a reader whether your model can be rerun, scaled, and compared with others.

A STANDARD

A widely used template for describing agent models is the ODD protocol (overview, design concepts, details; Grimm et al., 2006 and later updates). It asks for the same things in a fixed order.

ONE-LINE REPORT

“10,000 agents; 10 random contacts per infectious agent per day; synchronous daily steps; about 0.9 million contact checks per run; 40 runs, seeds 1000–1039.”

Check yourself: Part 1

CHECK YOURSELF

1. 5,000 agents each meet 8 random people a day; chance per contact 3%; recovery chance 1/6 per day. What is R_0?

2. Why can an agent model with a fixed number of agents and finite states not be universal?

3. Doubling N doubles the work of the k-contact model but quadruples the naive one. Why?

TIP

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

Applied
Computational
Mathematics
Lab

PART 2

Building blocks

Time, order, geometry and randomness.



PART 2 · TIME

Steps, events, or leaps

	FIXED STEPS	EVENT BY EVENT (GILLESPIE)	TAU LEAPING
time	jumps by Δt	jumps to the next event	jumps by τ
exact?	no: error grows with Δt	yes, for the stochastic model	approximately
cost	steps × agents	one update per event	leaps × event types
good for	simple, synchronous rules	small or detailed models	large populations

Our town of 10,000: about 17,856 events (infections and recoveries) in one epidemic, against 200 daily steps or 100 two-day leaps.

$$\Pr(\text{infected during } \Delta t) = 1 - e^{-\lambda \Delta t} \quad \text{not} \quad \lambda \Delta t$$

IN PLAIN WORDS

With a fixed step, everything that happens inside one step is decided from the state at its start. If a lot can happen in one step, the model drifts from the process it describes.

WHY THE EXPONENTIAL

A rate λ = 0.5 per day with weekly steps: λΔt = 3.5, which is not a probability; $1 - e^{-3.5} = 0.97$ is.

TRY IT

With λ = 0.2 per day, compare λΔt and $1 - e^{-\lambda \Delta t}$ for Δt = 1 and Δt = 7.

PART 2 · TIME

The step size changes the epidemic

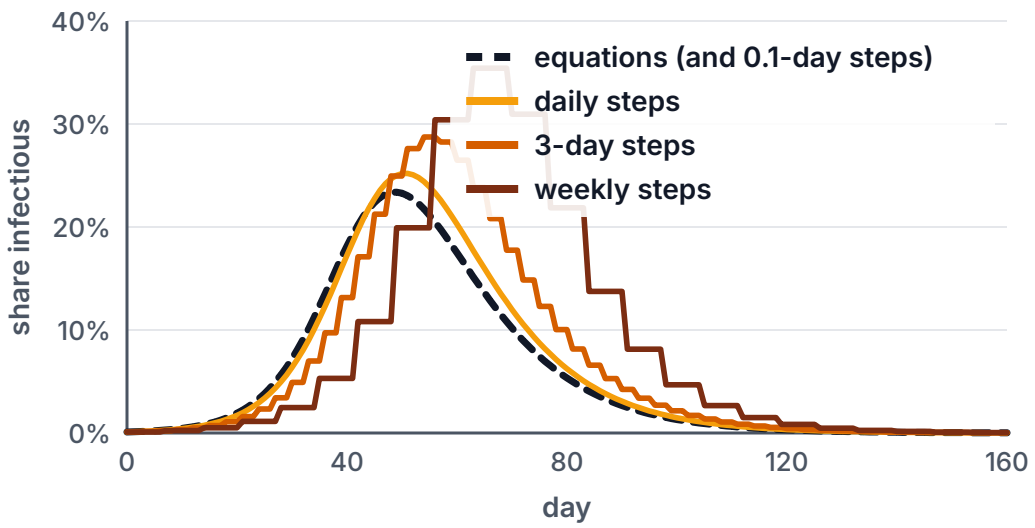

Our runs: the same expected-value model, steps of 0.1, 1, 3 and 7 days; β 0.25, γ 0.1.

STEP	PEAK INFECTIOUS	PEAK DAY	PEAK (ΛΔT VERSION)
0.1 day	23.6%	49	23.5%
1 day	25.2%	51	24.1%
3 days	28.7%	54	25.7%
7 days	35.4%	63	28.4%

The equations: peak 23.4% on day 49.

- Weekly steps make the peak higher and two weeks late: newly infected people cannot recover within the step in which they were infected.
- A step much shorter than the infectious period (here a day or less) keeps the error small.

PART 2 · ORDER

All at once, or one by one?

UPDATE SCHEME	FRONT SPEED (CELLS PER STEP)
synchronous: everyone from the old state	1
one by one, left to right	all 200 in step 1
one by one, right to left	1
one by one, fresh random order	1.70 (average of 200 runs)

Our runs: a line of 200 cells, left end infected, infection certain between neighbours.

- Sequential updates let a newly infected agent act within the same step: chains form inside one step.
- A fixed order adds a direction to the world that the model never meant to have. Shuffle the order every step.

WHERE E − 1 COMES FROM

In a random order, the front passes cell after cell while the next cell happens to come later in the order: an increasing run. Its expected length is 1/1! + 1/2! + 1/3! + … = e − 1, about 1.72.

IN PLAIN WORDS

Synchronous updates are easy to parallelise (high-performance computing needs them), but they forbid chains within a step. Random sequential updates allow chains but depend on the order. Choose on purpose and report it.

TRY IT

Why does the right-to-left sweep move only one cell per step?

PART 2 · ORDER

The same effect in a town

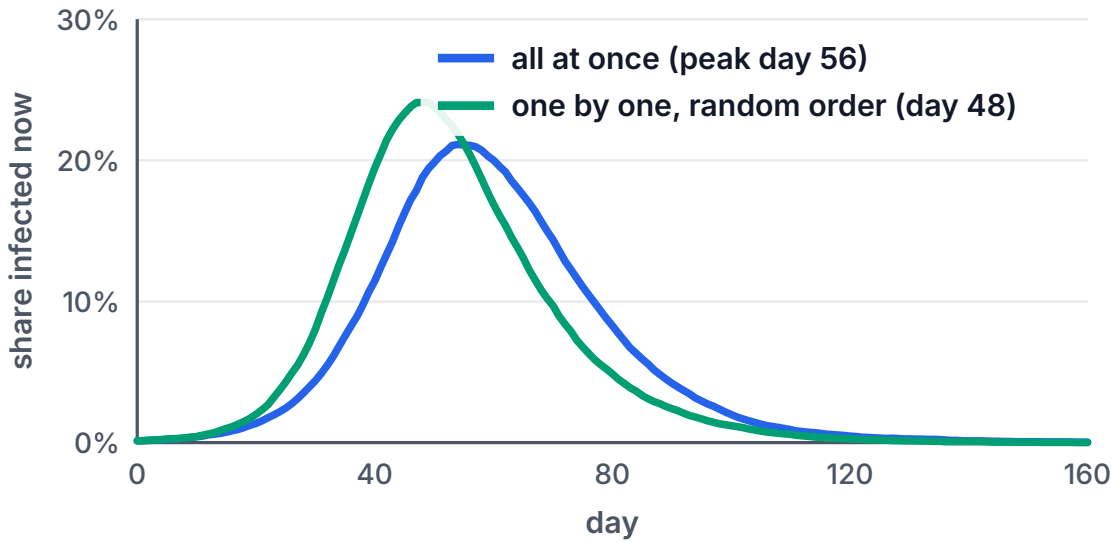

Our runs: 4,000 agents, 8 contacts a day, chance 3% per contact, γ 0.1; averages of 12 and 11 runs that took off.

- Random sequential updates: peak 25.0% on day 48. Synchronous: 22.0% on day 56.
- Same rules, same parameters: only the order changed. Someone infected in the morning can pass it on that afternoon.
- With a shorter step the two schemes get closer: fewer events fall inside one step.

IN PLAIN WORDS

Neither is wrong. But a result that depends on the update scheme must say which scheme it used.

PART 2 · GEOMETRY

Who is a neighbour?

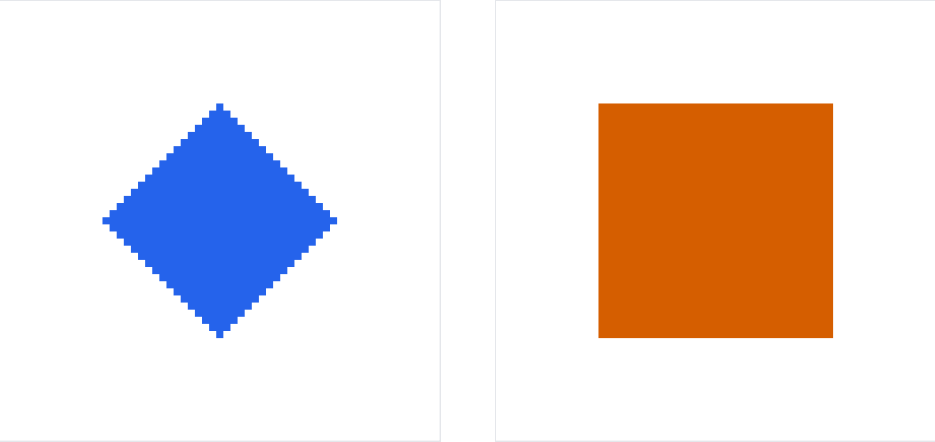

Our runs after 16 steps, certain infection: four neighbours (left) make a diamond, eight neighbours (right) a square.

AFTER 16 STEPS	4 NEIGHBOURS	8 NEIGHBOURS
certain infection	545	1,089
30% chance per neighbour	118	507

- With certain infection the counts follow formulas: $2t^2 + 2t + 1$ cells for four neighbours, $(2t + 1)^2$ for eight. Check $t = 16$.
- Grids suit fields, farms and buildings. For people, a contact network often matters more than distance (Module 4).
- Continuous space: agents move and meet within a radius. Finding neighbours fast needs a spatial index (a grid of buckets), or every agent checks every other.
- Networks can also change over time (temporal networks): who meets whom today is not who met yesterday.

TRY IT

How many cells does each neighbourhood cover after 10 steps of certain infection?

Seeds and reproducibility

- Computers use pseudo-random number generators: formulas that produce numbers that look random, from a starting value, the seed.
- Same seed, same run. Our 2,000-agent town with seed 42, twice: identical, agent for agent. Seed 43 ends at 58.9% instead of 58.7%.
- In parallel runs, give each process its own stream (generators with independent streams or keys). Otherwise the result can depend on which process ran faster.

IN PLAIN WORDS

Save the seeds, the code version and the parameters with every result. That is what makes a stochastic model reproducible.

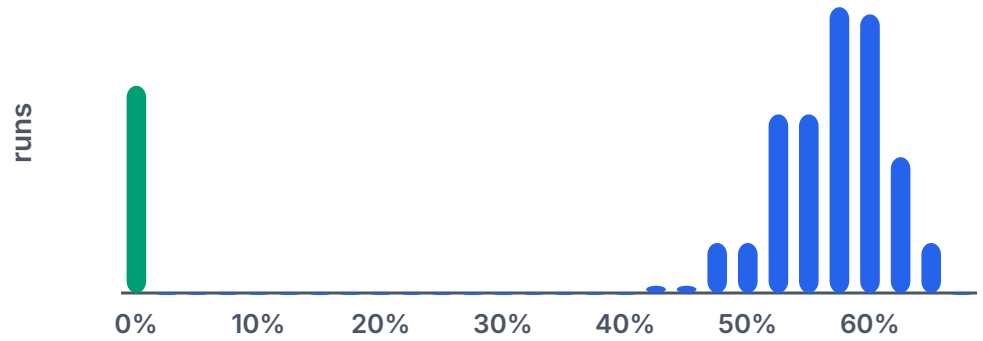

Our runs: 200 runs, 2,000 agents, R_0 1.5, 5 first cases. Orange: 29 runs that fizzled.

RUNS	10	100	1,000
uncertainty of the average (points)	±1.3	±0.4	±0.1

Single runs that took off: 45% to 65%, spread 4.2 points.

Common random numbers

$$\mathrm{Var}(A - B) = \mathrm{Var}(A) + \mathrm{Var}(B) - 2\,\mathrm{Cov}(A, B)$$

- Give both policies the same random numbers, agent by agent and day by day. Their runs then share most of their luck, the covariance is large, and the difference is sharp.
- Our trick: random numbers drawn in fixed blocks (per agent, per day) whatever the state, so a policy changes outcomes, not the stream.

POLICY (REMOVED CONTACTS)	SPREAD OF THE DIFFERENCE	PAIRS BELOW ZERO	RUNS SAVED
10%	2.2 to 1.5 points	0 to 0	2.4×
3%	2.1 to 0.9 points	37 to 6	5.6×

Our runs: 4,000 agents, R_0 2, 10 first cases, 200 pairs; independent runs, then common random numbers.

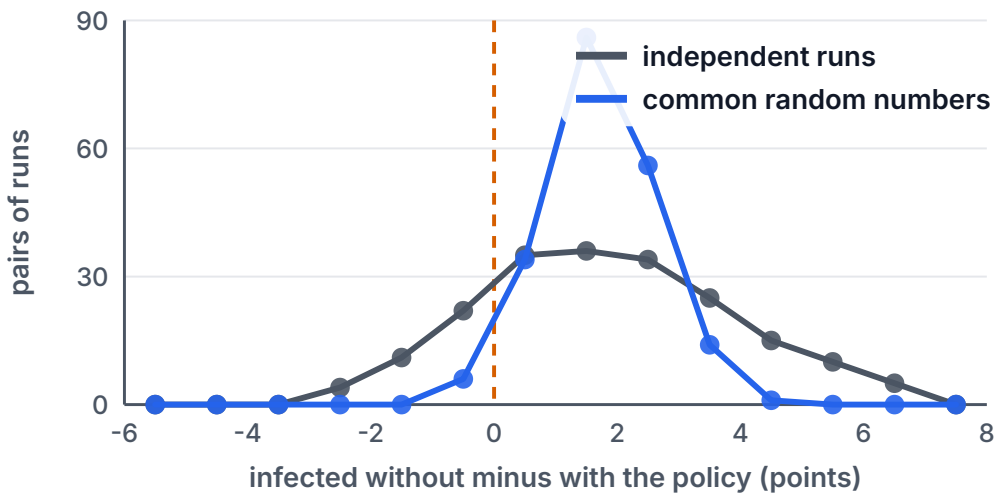

Removing 3% of contacts saves about 1.7% of the town on average. Independent pairs (grey) often point the wrong way.

TRY IT

Two policies each have run-to-run spread 4 points. With correlation 0.9 between paired runs, what is the spread of the difference? With independent runs?

PART 2 · RULES

Richer agents

- **Beliefs, desires, intentions**: agents hold beliefs about the world, have goals, and commit to plans (the BDI architecture).
- **Learning agents**: agents adapt their behaviour from experience, for example with reinforcement learning; or a controller learns a policy for the whole town.
- **Language-model agents**: a language model chooses what each agent does from a written description of its situation. Promising and expensive, and hard to validate.

IN PLAIN WORDS

Each richer rule is a new assumption with new parameters. It needs data behind it, or at least a clear statement that it is a scenario.

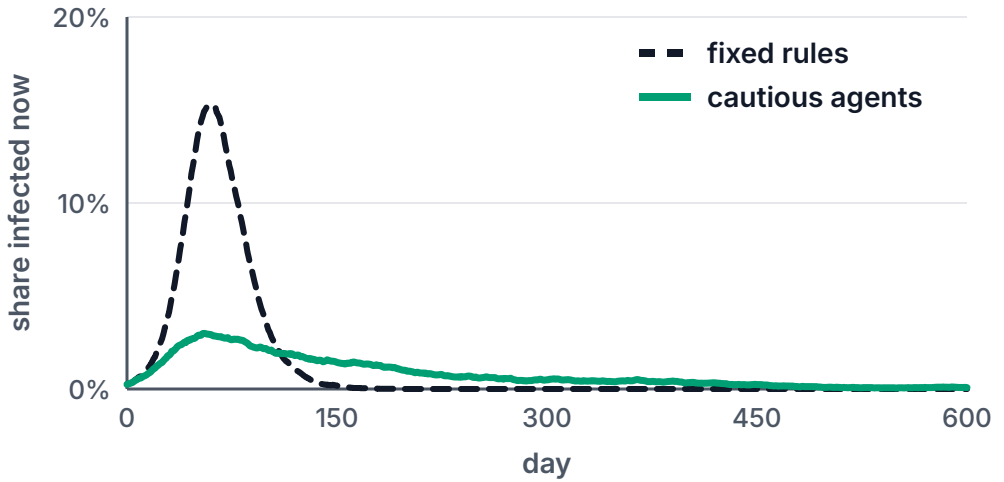

Our runs: agents who met an infectious person in the last week halve their chance of infection. Average of 10 runs; same town as before.

	FIXED RULES	CAUTIOUS AGENTS
share infected in the end	80%	49%

The cautious epidemic is slower and much longer: still going on day 700.

Check yourself: Part 2

CHECK YOURSELF

4. A model uses weekly steps for a disease with a 5-day infectious period. What goes wrong, and how would you fix it?

5. Why does a fixed left-to-right update order create an artefact, and what is the usual remedy?

6. Two policies differ by 1 point on average. Single runs spread by 3 points, and paired runs with common random numbers correlate at 0.95. How many pairs give a standard error of 0.25 points on the difference?

TIP

Answers are at the end of the deck.

Applied Computational Mathematics Lab

PART 3

In practice

Case studies, design, sensitivity and responsibility.



PART 3 · CASE STUDIES

A toy town with settings

OUR SYNTHETIC TOWN (MADE UP)	
people	4,578 (children 1,322, adults 2,676, seniors 580)
households	1,900 (1 to 6 people)
school classes	53 (up to 25 children of one age)
workplaces	168 (about 12 adults each)
around town	8 random meetings a day

CHANCE OF INFECTION PER DAY, PER INFECTIOUS PERSON	
household member	10%
classmate	1.2%
co-worker	1.2%
town meeting	1.2%

Susceptibility: children 1.0, adults 0.8, seniors 0.7. All values are made up to show mechanisms.

THE REAL TOOLS

FRED (Grefenstette et al., 2013) builds synthetic populations of real cities with households, schools and workplaces. Covasim (Kerr et al., 2021) adds detailed COVID-19 stages, testing and behaviour. We borrow their main ideas only; none of our numbers come from them.

HOW SETTINGS COMBINE

Force of infection on a susceptible person, per setting: $1 - (1 - p)^{(\text{infectious people there})}$, then weighted by the person's susceptibility.

A flu-like illness

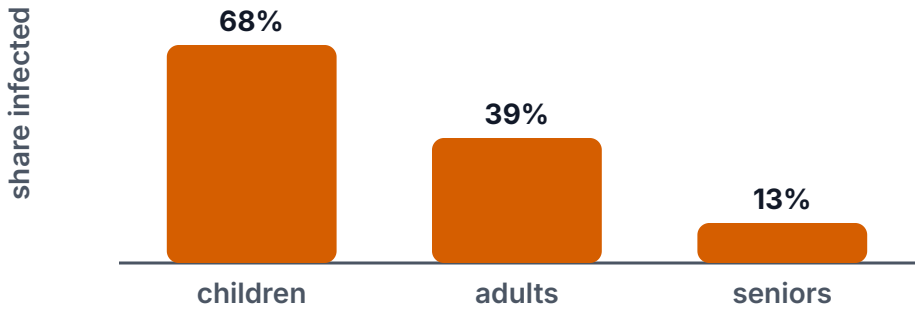

Our runs: 2 latent days, 4 infectious days (four stages each); 27 of 30 runs took off, 5 first cases.

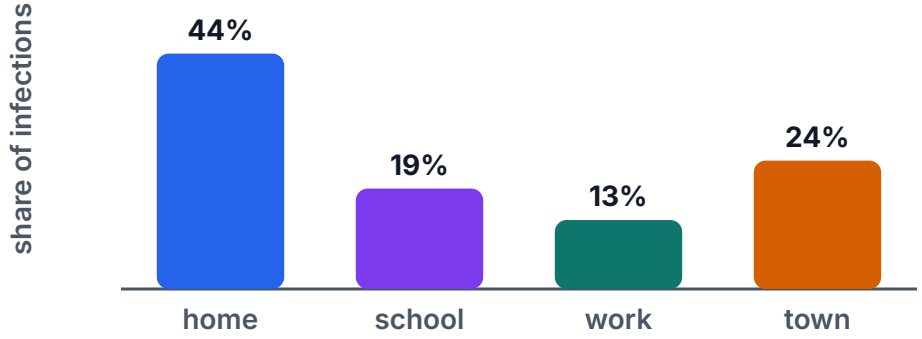

Where infections happened. Overall 44% infected; peak 3.8% infectious on day 51.

- Children are infected most: they meet more people at school and are more susceptible.
- A one-box model gives the total; only the structure can say where transmission happens.

Closing schools

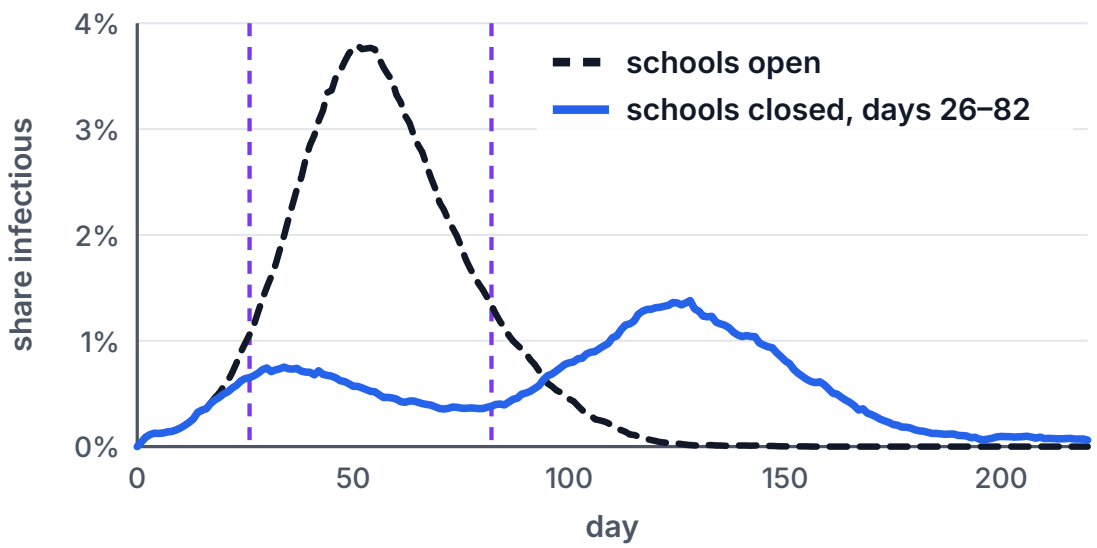

Our runs: schools close on the day 2% of the town has been infected (day 26 on average) for 8 weeks.

	SCHOOLS OPEN	WITH CLOSURE
share infected	44%	30%
children	68%	46%
adults	39%	27%
peak	3.8%, day 51	1.4%, day 128

- While schools are shut the epidemic stays low; when they reopen, it returns.
- The closure lowered the total and moved the wave by more than two months. When to reopen matters as much as when to close.

PART 3 · CASE STUDIES

A COVID-like illness and isolation

- Natural history (made up): 3 latent days, 2 infectious days before symptoms, then 5 more. 35% never get symptoms and are half as infectious.
- Without isolation: 66% infected; 27% of infections caused before symptoms.
- Isolation (80% comply) removes school, work and town contacts and halves household contact. The faster it starts, the larger the share of what remains that comes from people without symptoms: 39% with same-day isolation.

IN PLAIN WORDS

Isolation on symptoms cannot reach transmission that happens before symptoms. Testing contacts, or early testing, targets that part.

ISOLATION AFTER FIRST SYMPTOMS	SHARE INFECTED	FROM PEOPLE WITHOUT SYMPTOMS YET
none	66%	27%
4 days late	60%	29%
2 days late	51%	33%
1 day late	44%	35%
same day	34%	39%

Our runs: 20 runs per row, the same seeds in every row.

TRY IT

Which number in the table would change most if people without symptoms were as infectious as people with symptoms? Why?

Behaviour feeds back



Our runs: work and town contacts are divided by 1 + c × (cases per 1,000 reported a week earlier), c = 0.02.

	NO REACTION	PEOPLE REACT
peak infectious	12.3%	8.6%
share infected	66%	49%

- Behaviour lowers the peak and the total; the reporting delay makes the reaction late.
- This is the fear feedback of Module 2, now inside each agent. Fatigue can be added the same way.
- Covasim-style models also make infectiousness depend on time since infection; our stages are a coarse version of that.

PART 3 · DESIGN

Seven steps, tied to the question

DESIGNING AN AGENT-BASED SIMULATION

1. a question that names its output: quantity, population, horizon, comparison, and what the model must reproduce
2. the agents and their state: the coarsest agent that still expresses the mechanism
3. who meets whom: contacts organised by setting if an intervention acts on a setting
4. the illness: a state machine and stage durations with the right shape
5. behaviour and interventions, written as rule changes
6. the starting point: synthetic population and first cases
7. outputs and runs: what is logged, how often, and how many runs

IN PLAIN WORDS

The steps are presented in order, but design is iterative: a choice at step 4 can send you back to step 2.

VAGUE

How do schools affect transmission?

ANSWERABLE

How much does closing schools for 8 weeks change infections among adults over 200 days in this town?

DO YOU NEED AGENTS?

If everyone can be treated as average and the question is about totals (final size, peak, R_0), a compartmental model is clearer and faster. Agents earn their cost when heterogeneity, settings or individual rules matter.

PART 3 · DESIGN

Stage durations: the shape matters

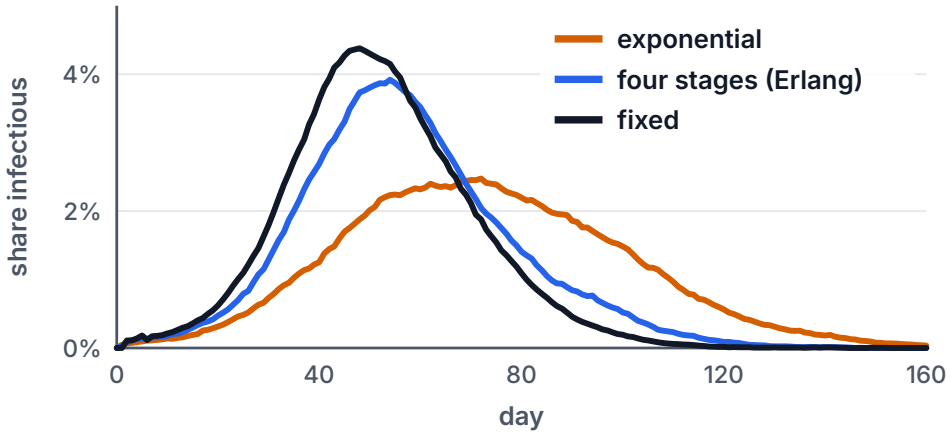

Our runs: same average durations (2 latent, 4 infectious days), three shapes; same seeds.

DURATIONS	PEAK DAY	PEAK INFECTIOUS	SHARE INFECTED	10 TO 200 CASES
exponential	72	2.5%	41.2%	34 days
four stages	54	3.9%	43.6%	29 days
fixed	48	4.4%	44.9%	27 days

- Exponential (memoryless) durations: many very short infections and a few very long ones. In a household, long infections waste days on people already infected; rounding to whole days also lengthens the short stages.
- Erlang durations (a chain of k stages) sit between exponential (k = 1) and fixed (k very large): the linear chain trick from Module 3.

TRY IT

A four-stage Erlang infectious period has mean 4 days. What is its standard deviation? (Each stage is exponential with mean 1 day.)

PART 3 · DESIGN

Rules, the start, and how many runs

INTERVENTION	AS A RULE CHANGE
vaccination	lower susceptibility (and maybe infectiousness)
isolation	remove contacts of isolated agents
masks	lower the chance per contact
school closure	switch off one contact layer

TWO FIRST CASES IN	OUTBREAKS THAT TOOK OFF	DAY 2% INFECTED
one school class	67 of 100	29
two random adults	57 of 100	30
two seniors	27 of 100	33

Our runs: 100 outbreaks per row, flu-like town.

RUNS	5	10	30	100
uncertainty of the average (points)	±1.6	±1.1	±0.7	±0.4

Single runs spread by 1.8 points: ±1 point needs 14 runs, ±0.5 points 53.

- Output contract: decide in advance what is logged (incidence, prevalence, by group and setting) and how often. Logging everything slows the model and invites fishing.
- Seeding is a modeling decision: where the first cases are changes whether and how fast an outbreak starts.

TRY IT

How many runs give ±0.25 points with a single-run spread of 1.8 points?

PART 3 · SENSITIVITY

Inputs or chance?

$$\mathrm{Var}(Q) = \mathrm{Var}_X\big(\mu(X)\big) + \mathbb{E}_X\big(\sigma^2(X)\big), \quad \mu(X) = \mathbb{E}[Q \mid X]$$

- Q: the output (here the share infected). X: uncertain inputs. The first term: variance from inputs; the second: run-to-run noise at fixed inputs.
- Our design: 30 settings of school and town transmission (each 70%–130% of our values) × 8 runs. Inputs explain 44.5% of the variance, chance 55.5%.
- 90% of all 240 runs ended between 34% and 51% infected: report that range, not one run.

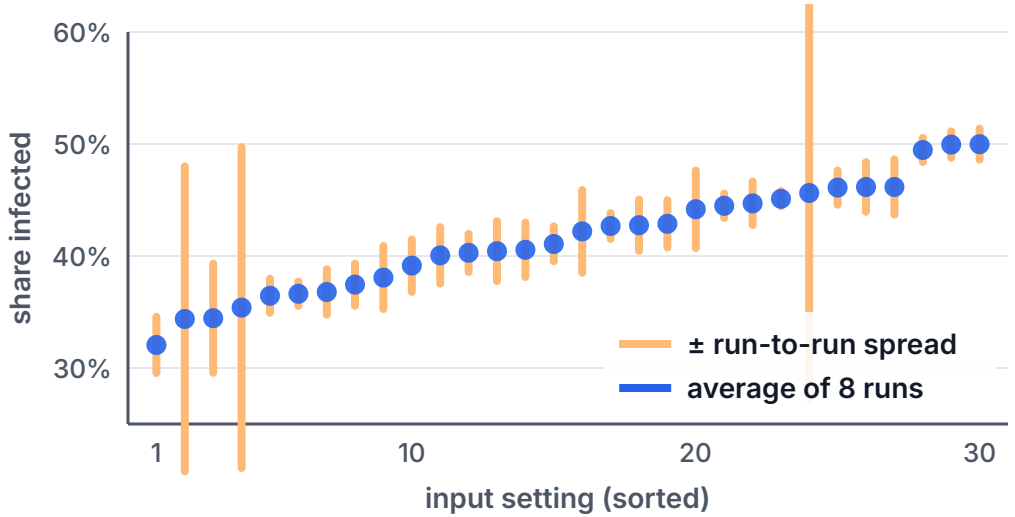

Our runs: each setting's average (dots) with its run-to-run spread (bars), sorted.

IN PLAIN WORDS

If chance dominates, more runs help; if inputs dominate, better data help. The split tells you where to spend effort.

Local sensitivity needs coupling

$$\frac{\partial \mu}{\partial x} \approx \frac{Q(x+h;\,\omega_1) - Q(x;\,\omega_2)}{h}, \qquad \omega_2 = \omega_1 \text{ (coupled) or independent}$$

- Nudge one input by h and measure the change. With independent runs the noise of two runs is divided by a small h: estimates swing wildly.
- Our runs, h = 5% of town transmission, 30 pairs: independent runs gave 11 negative estimates and one far off the scale; common random numbers (same ω) gave 0, average 0.26.
- Inputs that are not continuous (schools open or closed, a household size cap) are compared as scenarios. Global screening and variance-based methods (Morris, Sobol) come in Module 11.

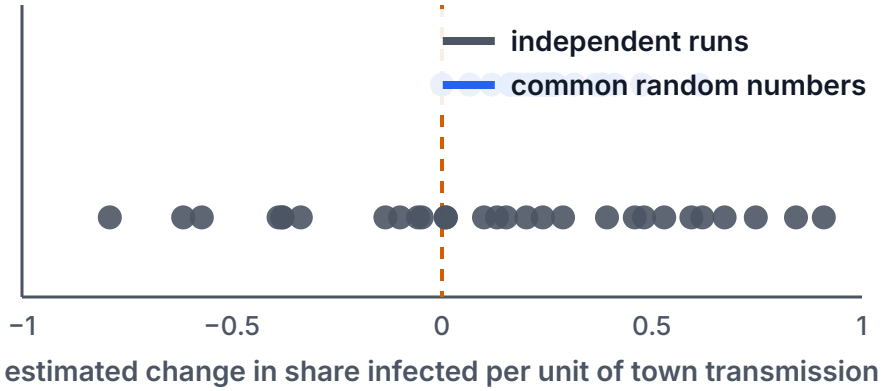

Each dot: one estimate from a pair of runs. Independent runs: grey (one value beyond the axis is not shown); common random numbers: blue.

TRY IT

Why does halving h make the independent estimate noisier, but leave the coupled estimate about as stable?

PART 3 · RESPONSIBILITY

Privacy, equity and policy risk

FACTS KNOWN ABOUT A PERSON	SHARE OF OUR TOWN SINGLED OUT
exact age, household size, neighbourhood	24.6%
age in 10-year bands, household size, neighbourhood	0.4%

Our made-up synthetic town of 4,578 people in 10 neighbourhoods.

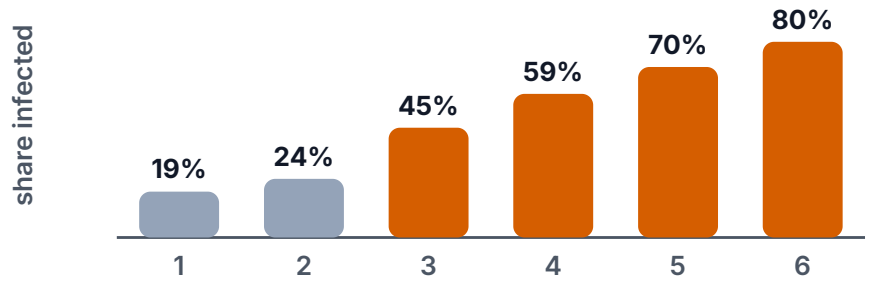

Flu-like runs: share infected by household size (people per household).

- **Privacy**: realistic agents are built from data about real people. Coarsen and aggregate, and treat synthetic populations as sensitive data.
- **Equity**: the burden is uneven (from 19% of people living alone to 80% in six-person households here), and so are the costs of interventions such as school closures. Report by group.
- **Policy risk**: a detailed model looks precise. State the question it was built for, the range of outcomes, the assumptions behind behaviour, and what it cannot tell you.

IN PLAIN WORDS

Detail is not certainty. The more realistic a model looks, the more carefully its limits must be stated.

Check yourself: Part 3

CHECK YOURSELF

7. Name two outputs that a layered agent model gives and a single-box SIR model cannot.

8. Inputs explain 30% of the output variance. Will 4 times as many runs per setting shrink the total spread of outcomes much? Why?

9. Why should results of an agent model of school closures be reported by household type or income group, not only in total?

TIP

Answers are on the next slides.

ANSWERS

Answers: warm-up and Parts 1–2

- **(a)** False: on average they give back the equations, plus noise. **(b)** False: weekly steps gave a 35% peak instead of 23%. **(c)** False: a fixed order sends infection across the whole line in one step. **(d)** True.
- **1.** $\beta = 8 \times 0.03 = 0.24$, $\gamma = 1/6$: $R_0 = 1.44$. **2.** Its global state space is finite, so it must eventually repeat; universality needs unbounded memory. **3.** Pairs grow with N^2; contact checks follow infectious agents, about proportional to N.
- **4.** Infection and recovery are decided once a week, so people infected in a step cannot recover within it and chains inside a week are missed; use steps of a day or less, or event-driven simulation.
- **5.** Sequential updates in a fixed order let changes run through the line in the update direction; shuffle the order every step (or update synchronously).
- **6.** Spread of the difference: $\sqrt{2 \times 9 \times (1 - 0.95)} = 0.95$ points; n = $(0.95 / 0.25)^2$, about 15 pairs.

TRY-IT ANSWERS

Try-it answers: 12 p × 8 = 3 gives p = 3.1%. λ = 0.2: Δt = 1: 0.20 vs 0.18; Δt = 7: 1.4 (not a probability) vs 0.75. Spreading infectiousness changes take-off chance (a few spreaders start most chains) more than final size, which depends on who can be infected. Right-to-left: each cell is updated before its left neighbour becomes infected. 10 steps: 221 and 441 cells. Edge checks: 300,000 × 8 × 20 = 48 million. CRN: $\sqrt{16 + 16 - 2 \times 0.9 \times 16} = 1.8$ points; independent: 5.7.

Answers: Part 3

- **7.** For example: who gets infected by age or group, and where infections happen (home, school, work); or the effect of closing one setting.
- **8.** No: more runs shrink only the chance part (70%), and only within each setting; the 30% from inputs stays. Better input data are needed to narrow it.
- **9.** Burden and costs are uneven: crowded households carry more infection, and closures cost some families far more. A total can hide who pays.

TRY-IT ANSWERS

Try-it answers: the presymptomatic share with fast isolation would change most, because people without symptoms are never isolated. Erlang with 4 stages of mean 1: variance $4 \times 1 = 4$, standard deviation 2 days. Runs for ±0.25 points: $(1.96 \times 1.8 / 0.25)^2$, about 200. Halving h doubles the noise term of independent runs; with common random numbers the two runs share their noise, which mostly cancels.

CREDITS

Sources and credits

- This deck accompanies Chapter 6 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 towns are our own runs (production/explains/sims/epi_m06_sims.py), with made-up values chosen to show mechanisms, not to estimate any real pathogen or place.
- Named works, described in our own words: Rule 110 universality (Cook, 2004); the ODD protocol (Grimm et al., 2006); FRED (Grefenstette et al., 2013); Covasim (Kerr et al., 2021); Gillespie's algorithm and tau leaping (Gillespie, 1977 and 2001).
- Companion videos: **Epidemic Modeling**, module 6, episodes 20–25. More from the lab: acml.teddylazebnik.com

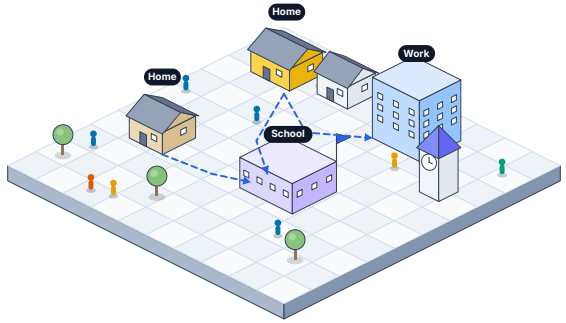

Applied Computational Mathematics Lab

END OF MODULE 6

Multi-agent simulation

Module 6 complete



YOU COVERED

- Why agents? Their power and their cost
- Building an agent-based model
- Randomness and reproducibility
- Two case studies: flu-like illness and COVID-19
- Designing your own simulation
- Sensitivity, ethics and policy risk

UP NEXT · MODULE 7

Data-driven models

From observation models to machine-learning forecasts

Based on Chapter 6 of T. Lazebnik, Epidemic Modeling: Principles, Algorithms, and Applications (Academic Press, 2026)