Applied Computational Mathematics Lab

EPIDEMIC MODELING · MODULE 10

Intervention policies

Interventions, optimal control and vaccination

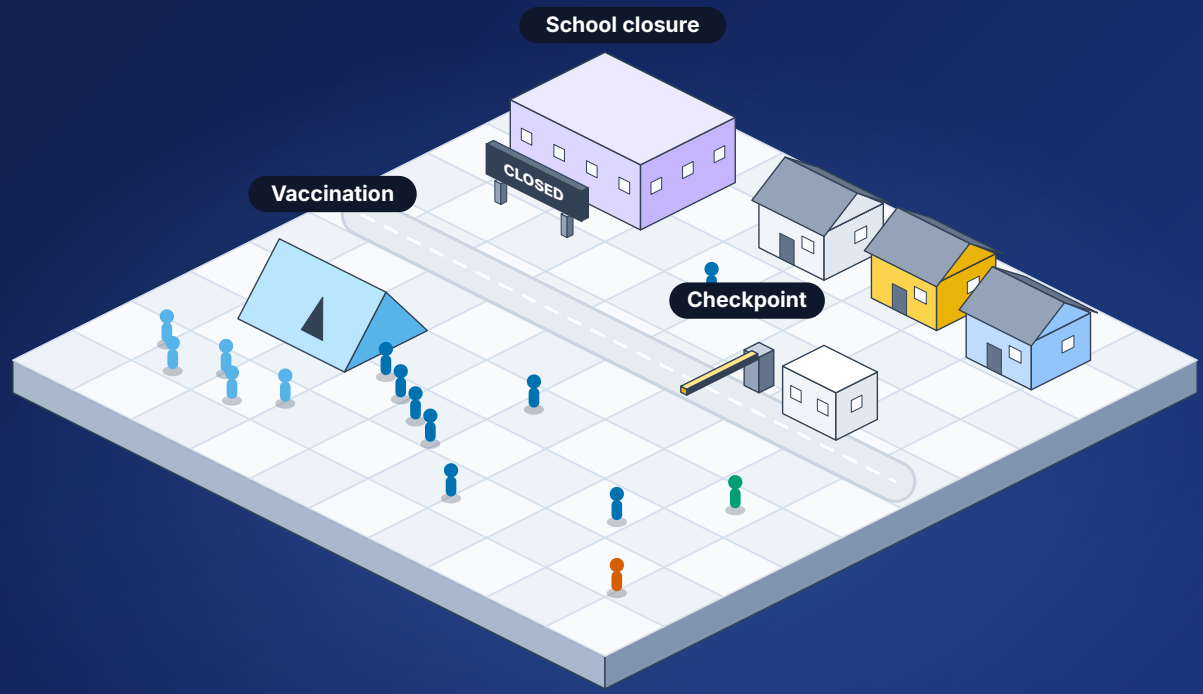

4 EPISODES

1. Putting interventions into the model
2. Optimal control for epidemics
3. Health and economy; vaccination strategy
4. Space, networks, strains and real use cases

Based on Chapter 10 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) A lockdown helps more the earlier it starts.

(b) Shielding a vulnerable group can save many lives while barely changing R_0.

(c) When the cost of a measure grows with the square of its strength, the best policy is one strict lockdown.

(d) Vaccinating the people with the most contacts first minimizes deaths.

HOW TO USE THIS

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

WHY THIS MODULE

Models that only predict cannot tell a mayor what to do. This module puts decisions inside the equations, chooses them with optimal control, and weighs them against other costs.

BEFORE WE START

This module's map

PART 1

Interventions in the model

A dial on transmission, settings in the contact matrix, vaccination as a flow, isolation as a box.

PART 2

Optimal control

The objective, Pontryagin's principle, two ways to solve it, and re-planning as data arrive.

PART 3

Trade-offs and vaccines

Health against the economy on a Pareto front, and who gets a scarce vaccine first.

PART 4

Space, data and evolution

Travel cuts, policies triggered by lagged data, drug resistance, and real epidemics.

GOALS

By the end you should be able to write an intervention into a model, set up and solve a small optimal-control problem, read a Pareto front without pretending it picks the answer, compare vaccine orders for different goals, and judge a trigger policy by the lag of its signal.

ONE RULE

Every number in this deck comes from made-up towns of one million people (R0 2.5, a thousand hospital beds) that we simulate. The values show mechanisms, not any real disease.

Applied Computational Mathematics Lab

PART 1

Interventions in the model

Decisions as inputs to the equations.

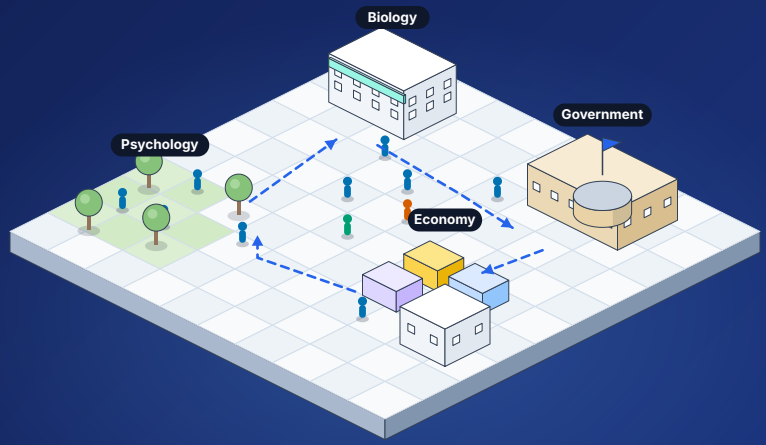

PART 1 · CONTROLS

Four ways an intervention enters the equations

INTERVENTION	IN THE MODEL	EXAMPLE
distancing, masks	$\beta(t) = \beta_0\,(1 - u(t))$	a lockdown of strength u
targeted measures	$C(t) = \sum_k w_k(t)\,C^{(k)}$	schools closed: the school matrix scaled down
vaccination	$\dot{S} = -\lambda S - v(t)\,S$	a flow from S straight to immune
testing and isolation	$\dot{I} = \lambda S - (\gamma + \delta(t))I$	found people move to an isolation box Q

Each control has a range (u between 0 and 1, doses limited by supply) and a cost.

- A control is a time-varying input chosen by a decision maker; the rest of the model is unchanged.
- Keep the mechanism readable: a parameter you can explain is a policy you can discuss.
- Feasibility matters as much as the formula: capacity limits, compliance, delays and supply all bound what a control can do.

IN PLAIN WORDS

Once decisions are inputs, the model can answer “what if we act, and when?”, not just “what will happen?”.

Strength alone is not enough

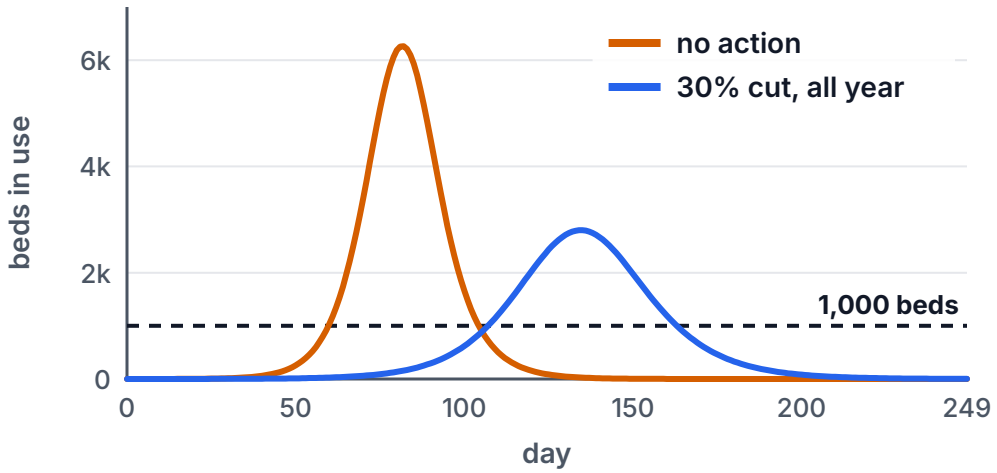

Our town (SEIR, R0 2.5, latent 3 days, infectious 5 days): 2% of infections need a bed for 10 days.

	NO ACTION	30% CUT, ALL YEAR
infected	89%	71%
peak new infections a day	32,794	14,190
peak beds	6,264	2,799
days over capacity	45	56
deaths	9,114	5,852

Deaths: 0.5% of infections, plus 40% of patients who find no bed (made-up values).

- Flattening lowers the peak by more than half, but the longer wave keeps beds over capacity for more days.
- A constant cut below the herd-immunity level only stretches the epidemic; to stay under capacity it must be much stronger, or well timed.

PART 1 · TIMING

When matters as much as how hard

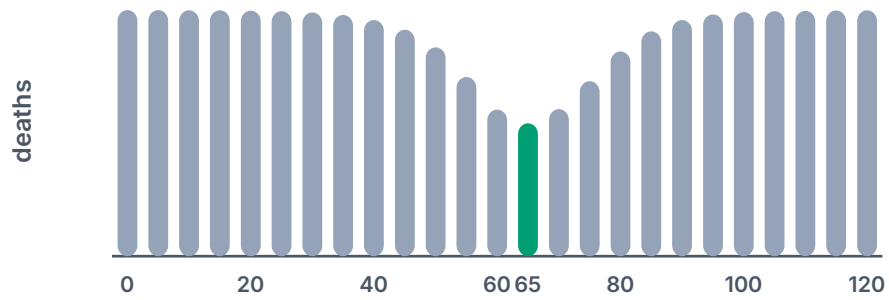

Deaths by the start day of a 60-day lockdown that cuts transmission by 60% (R held at 1). Green: the best start, day 65.

START DAY	DEATHS	DAYS OVER CAPACITY
no lockdown	9,114	45
10	9,109	45
65	4,921	33
100	9,039	45

- Too early: with few infections and nobody immune, the epidemic simply waits and resumes when the measure ends.
- Too late: the peak has passed and immunity is already doing the work.
- Well timed: the measure covers the peak and lets immunity build without a large overshoot.

TRY IT

Why can a lockdown that starts on day 0 and holds R exactly at 1 for 60 days leave the final size unchanged?

Targeted measures act on the contact matrix

CONTACTS A DAY	WITH CHILDREN	WITH ADULTS	WITH SENIORS
children	10.0	4.3	0.6
adults	1.6	9.3	1.0
seniors	0.7	3.3	2.7

Our age-structured town: 22% children, 60% adults, 18% seniors; home, school, work and other places (made-up numbers; reciprocal in totals).

$$K_{ij} = \frac{q}{\gamma} C_{ij}, \qquad R_0 = \rho(K)$$

Next-generation matrix and its spectral radius (q chosen so that R0 = 2.5).

POLICY	R0	INFECTED	DEATHS
usual contacts	2.50	85%	4,031
schools closed	2.09	78%	3,901
half of work from home	2.31	77%	3,650
half of other places	2.15	75%	3,097
all three	1.29	39%	1,781
shield seniors (their outside contacts 30%)	2.46	79%	2,787

Deaths per infection by age: 0.01%, 0.2%, 2.5% (made up); no hospital limit in this version.

- Closing schools lowers R_0 the most of the single measures, yet barely changes deaths: children rarely die here.
- Shielding seniors hardly moves R_0 but cuts deaths by 31%: the goal decides which number matters.

A flow out of S, and a box for people who are found

$$\dot{S} = -\lambda S - v(t)\,S, \qquad \dot{R} = \gamma I + v(t)\,S$$

0.5% OF SUSCEPTIBLE PEOPLE PROTECTED A DAY, FROM DAY	INFECTED	DEATHS
0	46%	3,318
30	65%	5,964
60	82%	8,221
90	89%	9,092
no vaccine	89%	9,114

Waning, partial or delayed protection need extra compartments, but the idea is the same: a controlled flow.

$$R = R_0\left[(1-f) + f\,\frac{\gamma}{\gamma+\delta}\right]$$

SHARE FOUND F	MEAN TIME TO ISOLATION	R	INFECTED
60%	never	2.50	89%
60%	4 days	1.67	68%
60%	2 days	1.43	53%
60%	1 day	1.25	37%
90%	4 days	1.25	37%
90%	2 days	0.89	0%

Our runs: SEIR town; found infectious people stop transmitting once isolated.

IN PLAIN WORDS

Speed and coverage both matter: with 40% never found, even instant isolation cannot bring R below 1 here.

Check yourself: Part 1

CHECK YOURSELF

1. Write the SIR equations with a transmission control u(t) and a vaccination rate v(t).

2. In our town a lockdown holding R at 1 for 60 days helped most when started on day 65. Explain in two sentences why day 10 was worse.

3. With R0 2.5 and an infectious period of 5 days, 80% of infections are found and isolated after 1 day on average. What is R?

TIP

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

Applied
Computational
Mathematics
Lab

PART 2

Optimal control

Letting the mathematics choose the policy.

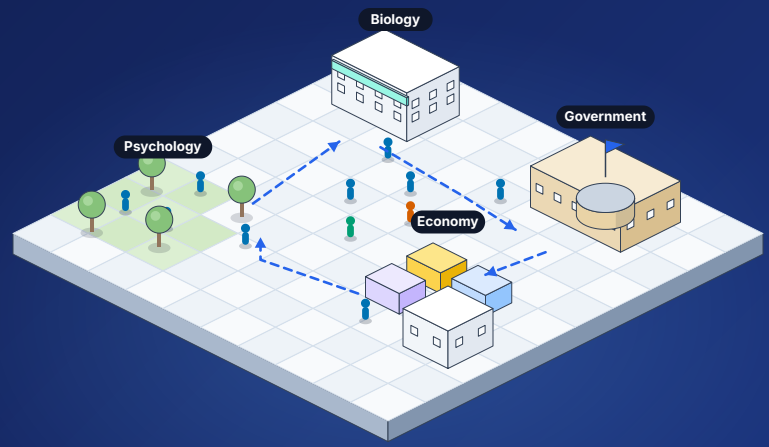

What we want to make small

$$\dot{x} = f\big(x(t), u(t), \theta\big), \qquad J(u) = \int_0^T L\big(x(t), u(t)\big)\, dt + \Phi\big(x(T)\big)$$

$$L = A\, I(t) + \tfrac{B}{2}\, u(t)^2, \qquad 0 \le u(t) \le u_{\max}$$

- x: the state (S, I); u: the control; L: the running cost; Φ: the cost left at the horizon.
- Our town: a simple SIR model, $\beta_0 = 0.25$, $\gamma = 0.1$ (R0 2.5), a cut of up to 70%, a horizon of 400 days.
- Φ = A times the infection-days still to come if all control stops at the horizon (from the final-size relation). Without it, suppressing until the last day looks free.

WEIGHTS

A is the weight on one person infected for one day; B scales the cost of a cut. Only their ratio matters for the shape of the policy, and choosing it is a value judgement.

WHY A SQUARE

The squared cost makes a small cut cheap and a large one expensive, which favours spreading effort over time. A cost in proportion to u does not, as Part 2 will show.

The maximum principle and its adjoints

$$\mathcal{H} = AI + \tfrac{B}{2}u^2 + \lambda_S\big(-\beta(1-u)SI\big) + \lambda_I\big(\beta(1-u)SI - \gamma I\big)$$

$$\dot{\lambda} = -\frac{\partial \mathcal{H}}{\partial x}, \quad \lambda(T) = \nabla\Phi\big(x(T)\big), \qquad \frac{\partial \mathcal{H}}{\partial u} = Bu + \beta SI(\lambda_S - \lambda_I) = 0$$

$$u^*(t) = \min\left\{u_{\max},\ \max\left\{0,\ \frac{\beta SI\,(\lambda_I - \lambda_S)}{B}\right\}\right\}$$

- The states run forwards from day 0; the adjoints run backwards from the horizon.
- Each adjoint is a shadow price: the future cost of one more person in that state. $\lambda_I - \lambda_S$ is the future cost of one more infection.
- The rule: cut harder when many infections are happening and when an infection costs much more in the future than a susceptible person.
- These are necessary conditions: a solution must satisfy them, and we check it numerically.

PART 2 · SOLVING IT

Two solvers, one answer

INDIRECT (ADJOINT) METHOD

1. Guess u(t); integrate S and I forwards.
2. Integrate the adjoints backwards from the horizon.
3. The gradient of J is $Bu + \beta SI(\lambda_S - \lambda_I)$; step against it, keep u in its range, accept only steps that lower J.
4. Stop when u matches the maximum-principle rule (here to within 0.0002).

DIRECT TRANSCRIPTION

1. Make u piecewise constant (one value a week, 58 values).
2. Evaluate J by integrating the model.
3. Hand the 58 numbers to a quasi-Newton optimiser with bounds.

METHOD	J	MODEL RUNS
adjoint method	73.23	365 iterations
direct, weekly steps	73.24	1,829 evaluations

Our runs (A = 10, B = 1): the two answers differ by 0.02%.

IN PLAIN WORDS

Agreement between methods built on different ideas is one of the best checks you can run on an optimal-control result.

A PRACTICAL NOTE

Forward–backward sweeps with a plain update can oscillate on epidemic problems; a damped or line-searched gradient step is safer.

The optimal policy in our town

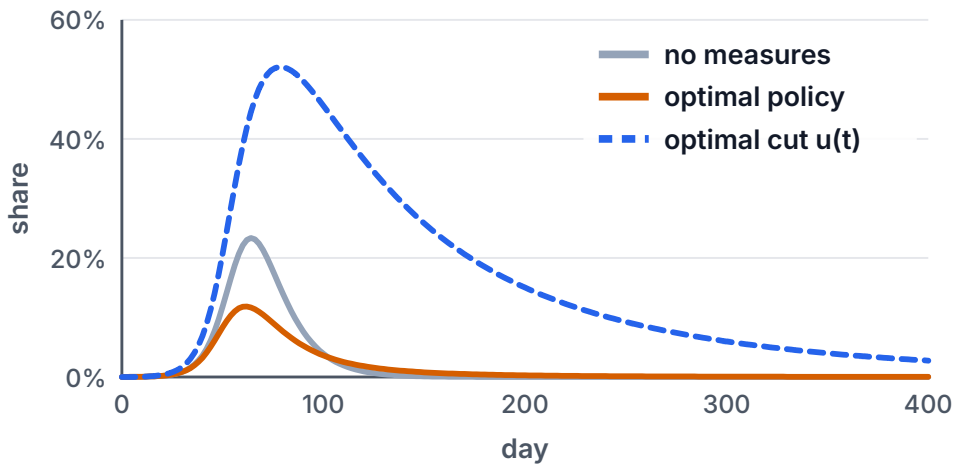

Share infectious with no measures and with the optimal policy, and the optimal cut (A = 10, B = 1).

POLICY	J	FINALLY INFECTED	PEAK INFECTIOUS
no measures	89.3	89%	23%
best constant cut (13%)	87.2	84%	18%
one strict 70% lockdown, days 60–110	79.5	67%	22%
optimal	73.2	63%	12%

- The optimum starts small, rises to a cut of about 52% after the peak of infections, then tapers slowly.
- It ends near herd immunity (60%): without a vaccine, the best a policy can do is avoid overshoot and spread the burden.
- The best single strict lockdown costs 8.6% more than the optimum.

PART 2 · SHAPES

Smooth or bang-bang

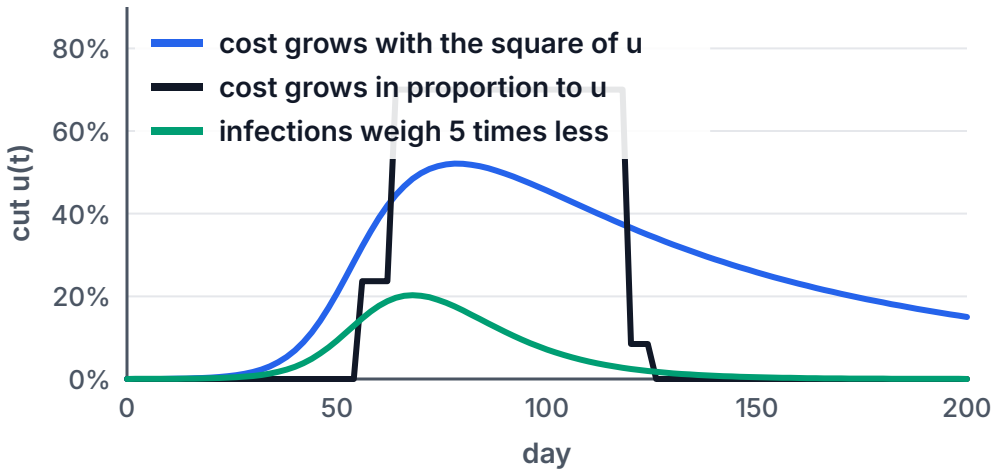

The optimal cut over the first 200 days under three cost assumptions.

- Cost in proportion to u (linear): the optimum is a single strict lockdown, here from about day 63 to day 118; the best block found by search starts on day 60 and lasts 60 days (cost 75.70 against 75.69).
- Cost with the square of u: a smooth rise and a slow taper.
- Infections weighted 5 times less: a gentle cut, never above 20%, for about 100 days.
- With a linear cost, $\mathcal{H}$ is linear in u, so the minimum sits at a bound: either no measure or the full one.

IN PLAIN WORDS

The shape of the best policy follows from what we count and how we weigh it. Changing the weights is a political choice, not a numerical one.

PART 2 · MODEL PREDICTIVE CONTROL

Plan, act, look, and plan again

THE PLANNER BELIEVES R0 = 2 (TRUTH: 2.5)	J IN THE TRUE TOWN
no measures	89.3
plan once, follow it to the end	89.7
re-plan every 14 days from the observed state	75.6
…and re-estimate R0 from the last 14 days each time	73.2
optimum with the true R0	73.2

Our runs: the planner sees prevalence; re-estimation uses its growth over the last block.

- An optimum is only as good as the model. A plan made with the wrong R0 and followed blindly did no better than doing nothing.
- Re-planning from the observed state corrects most of the error, even with the wrong parameter.
- Updating the parameter as well closes the gap. This is model predictive control: optimise over a horizon, apply the first part, observe, repeat.

TRY IT

Why does re-planning help even when the planner never learns the true R0?

Check yourself: Part 2

CHECK YOURSELF

4. From $\mathcal{H}$ for the controlled SIR model with $L = AI + \frac{B}{2}u^2$, derive $\partial\mathcal{H}/\partial u$ and the formula for u^*.

5. Why does a running cost in proportion to u lead to a bang-bang policy?

6. Without a terminal cost, what would an optimiser do near the end of the horizon, and why is that misleading?

TIP

Answers are at the end of the deck.

Applied Computational Mathematics Lab

PART 3

Trade-offs and vaccines

Two bills, and who goes first.

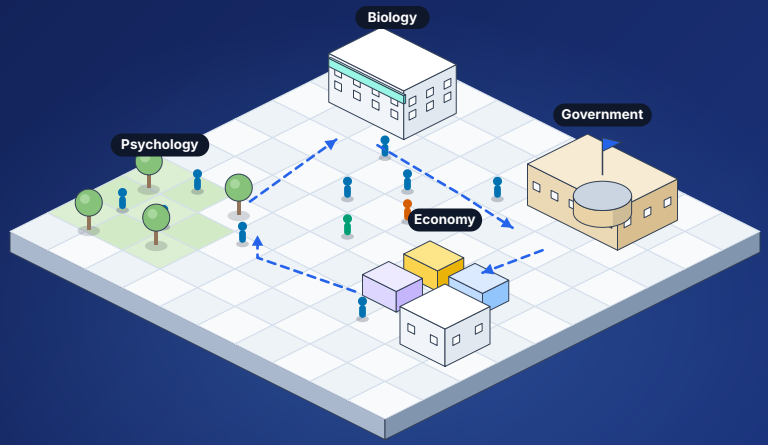

Putting the economy in the model

$$\text{output lost each day} = 0.5\,u(t) + 0.3\,b(t), \qquad b(t) = 0.5\,\frac{H(t)}{H(t) + 1000}$$

- u: lockdown strength; b: people's own contact cut, growing as hospital beds H fill (made-up rules).
- Transmission is scaled by $(1-u)(1-b)$: caution is endogenous, a feedback from the epidemic to behaviour (Module 2).

NO LOCKDOWN	WITHOUT CAUTION	WITH CAUTION
infected	89%	72%
deaths	9,114	5,871
output lost (share of a year)	0%	2.4%

Our runs: SEIR town, 500 days.

IN PLAIN WORDS

Doing nothing is not free: people protect themselves, and that costs output too. A lockdown should be compared with this baseline, not with an untouched economy.

ECONOMIC–EPIDEMIOLOGICAL MODELS

Coupled models put welfare, behaviour or resource limits inside the dynamics, so that health and economic outcomes come from the same run.

PART 3 · PARETO FRONT

Most policies are wasteful

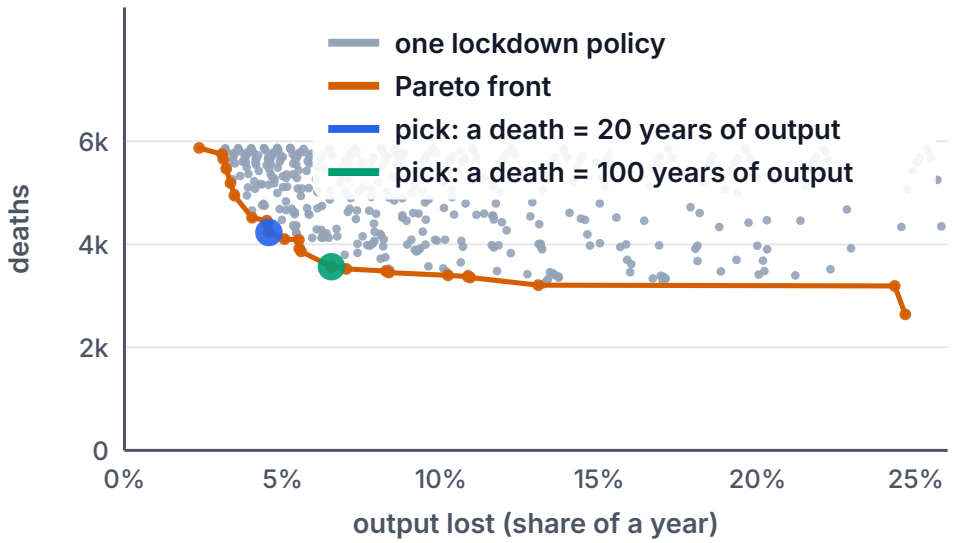

Our 468 lockdowns (strength 20–70%, start day 10–130, 30–240 days) and the front they leave.

- A policy is dominated when another one has both fewer deaths and a smaller loss. Only 24 points survive: the Pareto front.
- Example of waste: a long lockdown started late, after the peak: about the same deaths as doing nothing, at a high cost.
- Moving along the front trades output for lives. Which point to choose depends on how a death is weighed against output: a value judgement, not a calculation.

WEIGHT ON ONE DEATH	CHOSEN POLICY	DEATHS	OUTPUT LOST
20 years of output	40% cut, days 70–100	4,231	4.6%
100 years of output	40% cut, days 60–120	3,565	6.5%

When help is coming, waiting is worth more

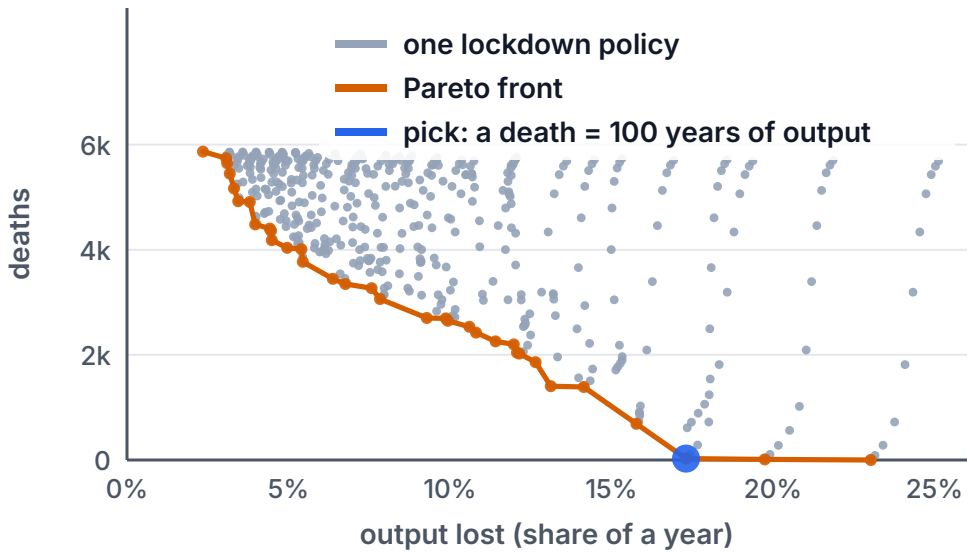

The same policies when a vaccine protects 1% of susceptible people a day from day 200 (90% effective).

- With a vaccine coming, a long, strong suppression no longer just delays the wave: it skips most of it. At the higher weight, the chosen policy is a 70% cut from day 10 to day 190: 26 deaths, 17.3% of a year's output.
- Without a vaccine, the fewest deaths any policy here achieved was 2,640.
- What lies beyond the horizon (vaccines, treatments, new variants) can change the best policy completely; it belongs in the model.

TRY IT

Why can the same lockdown be wasteful without a vaccine and wise with one?

Herd immunity and the next-generation matrix

$$R_v = R_0\,(1-\varepsilon p) < 1 \iff p > \frac{1 - 1/R_0}{\varepsilon}$$

Well mixed, R0 2.5: 60% immune; with a 90% effective vaccine, 67% coverage.

$$K(p) = \operatorname{diag}(1-\varepsilon p_1, \ldots, 1-\varepsilon p_m)\,K, \qquad \rho\big(K(p)\big) < 1$$

- With groups, coverage is a vector and the threshold is a condition on the spectral radius.
- Vaccinating a group helps through its own protection and through how strongly it is coupled to the others.

400,000 DOSES BEFORE ANY SPREAD, BY ORDER	COVERAGE (CHILDREN, ADULTS, SENIORS)	R
seniors first	0%, 37%, 100%	2.18
children first	100%, 30%, 0%	1.47
adults first	0%, 67%, 0%	2.09
everyone equally	40%, 40%, 40%	1.60

Our age-structured town, a vaccine that blocks 90% of infections.

IN PLAIN WORDS

Children have the most contacts in our town, so covering them lowers R the most. Lowering R is not the same as saving the most lives.

PART 3 · PRIORITY

Who should go first?

ORDER	BLOCKS INFECTION: INFECTIONS	DEATHS	PREVENTS SEVERE DISEASE: INFECTIONS	DEATHS
seniors first	645,020	1,266	789,961	1,339
children first	533,693	3,194	793,807	3,714
adults first	557,539	2,823	788,059	3,399
everyone equally	566,430	2,365	791,760	2,805
no vaccine	847,655	4,031		

Our runs: 5,000 doses a day from day 30, 400,000 in all. Both vaccines prevent 90% of deaths among the vaccinated; one blocks 90% of infections, the other 30%.

SENIORS FIRST, STARTING ON DAY	0	30	60	90
deaths	853	1,266	2,579	3,860

- To minimise infections, vaccinate the people with the most contacts (indirect protection).
- To minimise deaths, protect those most at risk directly; here that held for both vaccines.
- If the vaccine barely blocks infection, indirect protection fades: every order gives about the same number of infections.
- Starting 30 days later doubled the deaths.

IN SHORT

The best order depends on the vaccine, the contacts, the timing and the goal: it is not a universal ranking.

Check yourself: Part 3

CHECK YOURSELF

7. Policy X: 3,000 deaths and 6% of output lost. Policy Y: 3,200 deaths and 7%. Is Y on the Pareto front? Why?

8. With $R_0 = 4$ and a vaccine that is 80% effective, what coverage reaches herd immunity in a well-mixed population?

9. Give one reason the order that minimises infections may not minimise deaths.

TIP

Answers are at the end of the deck.

Applied
Computational
Mathematics
Lab

PART 4

Space, data and evolution

Where to act, when to act, and how the pathogen answers.

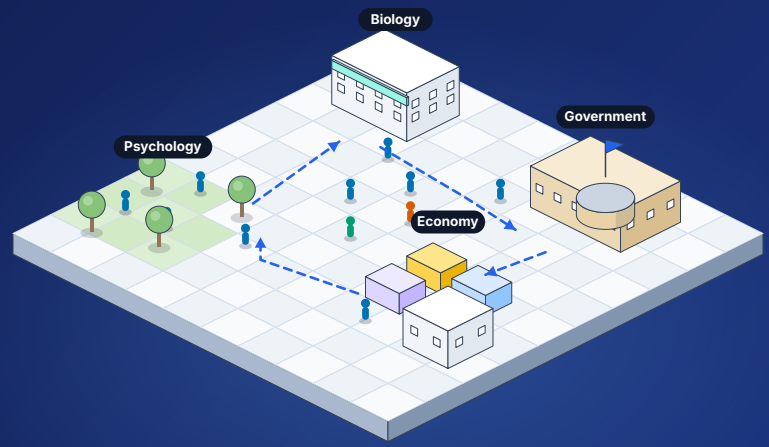

PART 4 · TRAVEL

Travel cuts buy time

$$M_{ij}(t) = M_{ij}^{0}\left(1 - u_{ij}(t)\right)$$

POLICY	HUB	NEIGHBOR TOWNS	FAR TOWNS	FINALLY INFECTED
no cut	38	38	51	89%
all travel −50%	42	42	58	89%
all travel −90%	51	51	77	89%
all travel −99%	66	66	105	89%
hub links only −99%	66	38	59	89%

Day each city passes 1 infectious person in 10,000. Our six cities: a hub of 2 million and five towns of 500,000 on a ring; the outbreak starts in town 1; 0.4% of town residents travel to the hub each day, 0.1% to each neighboring town.

- A tenfold cut bought about two weeks at the hub; a hundredfold cut about four. The final size did not change.
- Infection follows every path left open: cutting only the hub's links delays the hub, but the neighboring towns still catch it along the ring.
- On networks the same logic favors targeting: links and people that connect the most are worth more than an equal cut spread evenly (Module 4).

IN PLAIN WORDS

Time bought is only worth something if it is used: to prepare beds, tests, treatments and vaccines.

PART 4 · TRIGGERS

Policies that react to data, and their lag

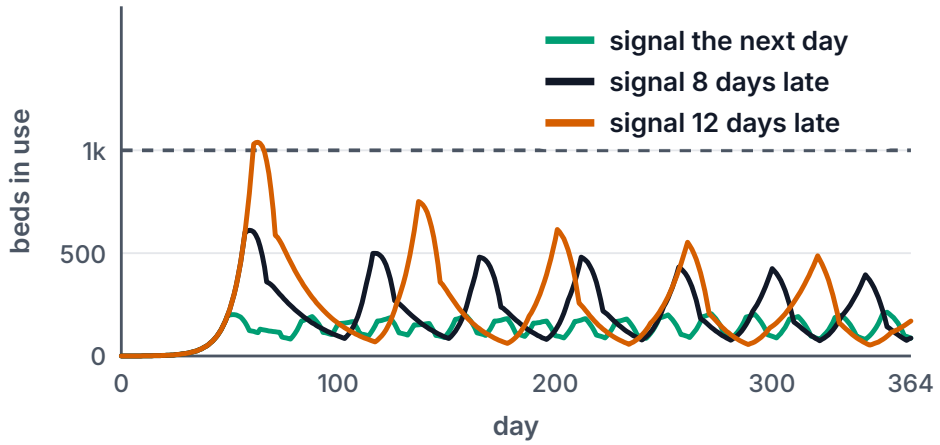

A 70% cut goes on when the signal passes 1,500 new infections a day and off below 500 (at least 14 days on); dashed: 1,000 beds.

SIGNAL	LAG (DAYS)	PEAK BEDS	DAYS OVER	LOCKDOWNS	DEATHS (YEAR 1)
infections, the next day	0	213	0	17	1,175
wastewater	4	352	0	11	1,618
reported cases	8	611	0	7	1,916
hospital admissions	12	1,039	5	5	2,229

Lags as in Module 9 (made-up values).

- The later the signal, the larger the overshoot past the trigger: with a 12-day lag the beds overflowed.
- Lags also make the response cruder: fewer, longer lockdowns and almost twice the deaths.
- Choose thresholds with the lag in mind, or trigger on faster streams.

PART 4 · EVOLUTION

Policies shape selection

$$r_i(t) = \beta_i(t)\, S(t) - \gamma_i, \qquad \frac{df_i}{dt} \approx f_i \left(r_i - \bar{r}\right)$$

SHARE OF INFECTIONS TREATED	R WILD	R RESISTANT	INFECTIONS	RESISTANT SHARE	RESISTANT MAJORITY FROM DAY
0%	2.50	2.25	894,612	0%	–
30%	1.92	2.25	777,242	1%	263
60%	1.56	2.25	730,585	33%	142
90%	1.32	2.25	846,205	95%	90

Our runs: an antiviral halves the infectious period of treated people; a resistant strain, 10% less transmissible, is unaffected; treatment raises the rate at which resistance appears.

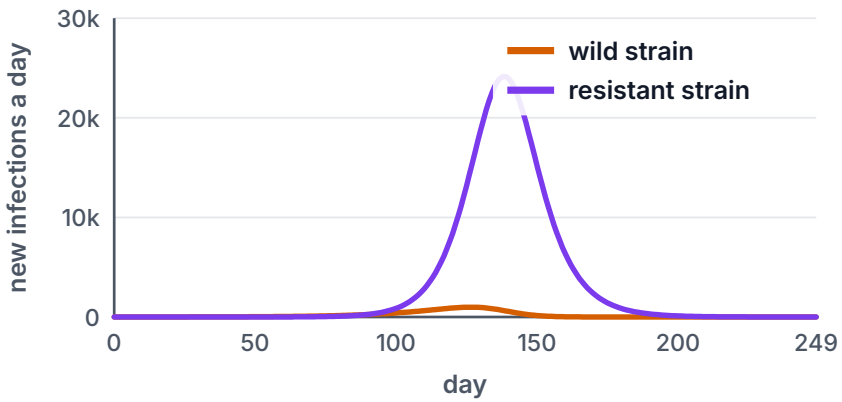

Treating 90% of infections: the resistant strain takes over.

- Treating 60% cut infections by almost a fifth, but a third were resistant; treating 90% handed 95% of infections to the resistant strain.
- Vaccines that prevent disease but not infection can relax selection against virulent strains, as seen with Marek's disease in poultry (Read et al., 2015).

What real epidemics taught

1918

Influenza in US cities

Cities that acted early and kept layered measures in place generally saw lower peaks of death (Hatchett et al., 2007; Bootsma and Ferguson, 2007).

PANDEMIC PLANNING

Containing a new influenza

Containment at the source depends on transmissibility, detection delays and how fast targeted antivirals reach contacts (Ferguson et al., 2005; Longini et al., 2005).

2003

SARS

Case finding and faster isolation after symptoms helped bring transmission down (Riley et al., 2003).

2014

Ebola in West Africa

Isolation capacity and safe burials had to scale up fast enough to keep pace with growth (WHO Ebola Response Team, 2014).

2020

COVID-19

Mitigation versus suppression, triggers, vaccine priorities and variants became daily decisions (Ferguson et al., 2020).

IN PLAIN WORDS

Across all of them: act on the true goal (peaks, deaths, time to a vaccine), expect delays in the data, and keep the response able to change course.

Check yourself: Part 4

CHECK YOURSELF

10. Early growth doubles every 5 days. Roughly how many days does a 90% cut in imported cases delay an epidemic’s arrival?

11. A trigger uses reported cases, which lag infections by 8 days. Name two ways to reduce the overshoot.

12. Use the replicator equation to explain why heavy treatment can make a resistant strain win even if it transmits less well.

TIP

Answers are on the next slides.

ANSWERS

Answers: warm-up and Parts 1–2

- **(a)** False: started on day 10, a 60-day lockdown changed almost nothing; day 65 was best. **(b)** True: shielding seniors moved R0 from 2.50 to 2.46 and cut deaths by 31%. **(c)** False: that gives a smooth policy; a cost in proportion to the strength gives one strict lockdown. **(d)** False: children first gave the fewest infections, seniors first the fewest deaths.
- **1.** $\dot{S} = -\beta(1-u)SI - vS,\ \dot{I} = \beta(1-u)SI - \gamma I,\ \dot{R} = \gamma I + vS.$
- **2.** On day 10 almost nobody was immune; holding R at 1 only paused growth, and the epidemic resumed at full speed when the lockdown ended.
- **3.** $R = 2.5\,[0.2 + 0.8 \times 0.2/(0.2+1)] = 2.5 \times 0.333 = 0.83.$
- **4.** $\partial\mathcal{H}/\partial u = Bu + \beta SI(\lambda_S - \lambda_I) = 0 \Rightarrow u^* = \beta SI(\lambda_I - \lambda_S)/B$, clipped to $[0, u_{\max}]$.
- **5.** $\mathcal{H}$ is then linear in u, so its minimum over an interval is at an end point: 0 or $u_{\max}$ (except on singular arcs).
- **6.** It would keep suppressing (or stop caring) at the end, because infections after the horizon cost nothing; the epidemic would simply happen later.

TRY-IT ANSWERS

Try-it answers: with R exactly 1 and almost everyone susceptible, the epidemic neither grows nor shrinks, then resumes. Re-planning helps because each new plan starts from the true state, so errors stop accumulating.

Answers: Parts 3–4

- **7.** No: X has fewer deaths and a smaller loss, so Y is dominated.
- **8.** $(1 - 1/4)/0.8 = 0.9375$: about 94%.
- **9.** Those with the most contacts (children here) are rarely those most at risk of dying (seniors); the vaccine may also block disease better than infection.
- **10.** Growth rate $r = \ln 2/5 \approx 0.14$ per day; a tenfold cut costs $\ln 10/r \approx 17$ days. Our six cities gave 13 days at the hub.
- **11.** Lower the on-threshold, or use a faster stream (wastewater, nowcasts that correct for delay).
- **12.** $df_i/dt \approx f_i(r_i - \bar{r})$: treatment lowers the wild strain's growth rate, so the resistant one has the higher r even with a transmission cost, and its share grows.

TRY-IT ANSWER

Try-it answer: without a vaccine, suppression only postpones the wave, paying twice; with a vaccine coming, the postponed infections are largely prevented.

CREDITS

Sources and credits

- This deck accompanies Chapter 10 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 and curves come from our own made-up towns (production/explains/sims/epi_m10_sims.py), with values chosen to show mechanisms, not to describe any real disease, place or policy.
- Named works, described in our own words: Pontryagin's maximum principle (Pontryagin et al., 1962); Read et al., 2015 (Marek's disease); Hatchett et al., 2007; Bootsma and Ferguson, 2007; Ferguson et al., 2005; Longini et al., 2005; Riley et al., 2003; WHO Ebola Response Team, 2014; Ferguson et al., 2020.
- Companion videos: **Epidemic Modeling**, module 10, episodes 41–44. More from the lab: acml.teddylazebnik.com

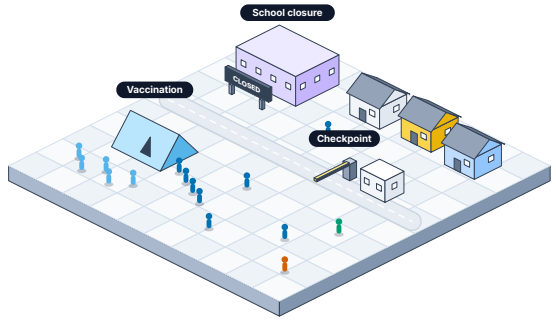

Applied Computational Mathematics Lab

END OF MODULE 10

Intervention policies

YOU COVERED

- Putting interventions into the model
- Optimal control for epidemics
- Health and economy; vaccination strategy
- Space, networks, strains and real use cases

Module 10 complete

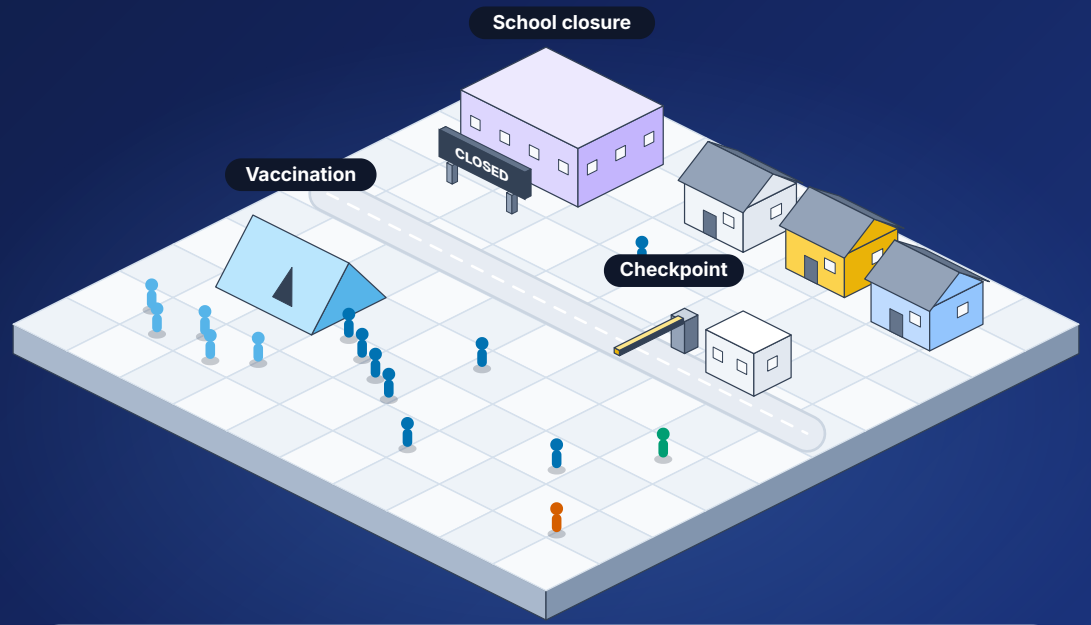

UP NEXT · MODULE 11

Uncertainty, sensitivity and model comparison

Where uncertainty comes from, and how to measure it

Based on Chapter 10 of T. Lazebnik, Epidemic Modeling: Principles, Algorithms, and Applications (Academic Press, 2026)