An elimination strategy for COVID-19 is the cheapest option

Lennon Ó Náraigh [1], Áine Byrne [1]

[1] School of Mathematics and Statistics, University College Dublin, Belfield, Dublin 4, Ireland

30th September 2020

Aim of Talk

In this talk, we formulate an optimal-control problem to determine what is in some sense the 'best' way to control the outbreak of COVID-19 in Ireland:

- Formulate a basic ODE model, fit it to data from the 'first wave' (March-May 2020)
- Formulate an optimal-control problem, including **state constraints**
- Solve the optimal-control problem numerically.

The numerical solutions show that elimination is the cheapest way to deal with the COVID-19 outbreak.

Structure of talk

- Crash course in Compartmental Models of Mathematical Epidemiology (SIR)
- Compartmental Model for COVID-19
- Optimal Control Theory

Caveats

- A lot of discussion / speculation about COVID-19 by ill-informed people.
- I don't want to add to it.
- I will speak about mathematical models only.
- Its applicability can be debated by real experts...

... but here is what an elimination strategy looks like

https://www.bbc.com/news/world-asia-china-53816511

Home | US Election | Coronavirus | Video | World | UK | Business | Tech | Science | Stories | Entertainment & Arts | More

Asia | China | India

Wuhan coronavirus: From silent streets to packed pools

18 August 2020

Share

Coronavirus pandemic

SIR model

Part 1: A brief review of compartmental models, starting with the simplest **SIR** model.

Key assumption: homogeneous population.

Other assumptions: no natural births/deaths on the timescale of the epidemic.
The model is then just 'conservation of people':

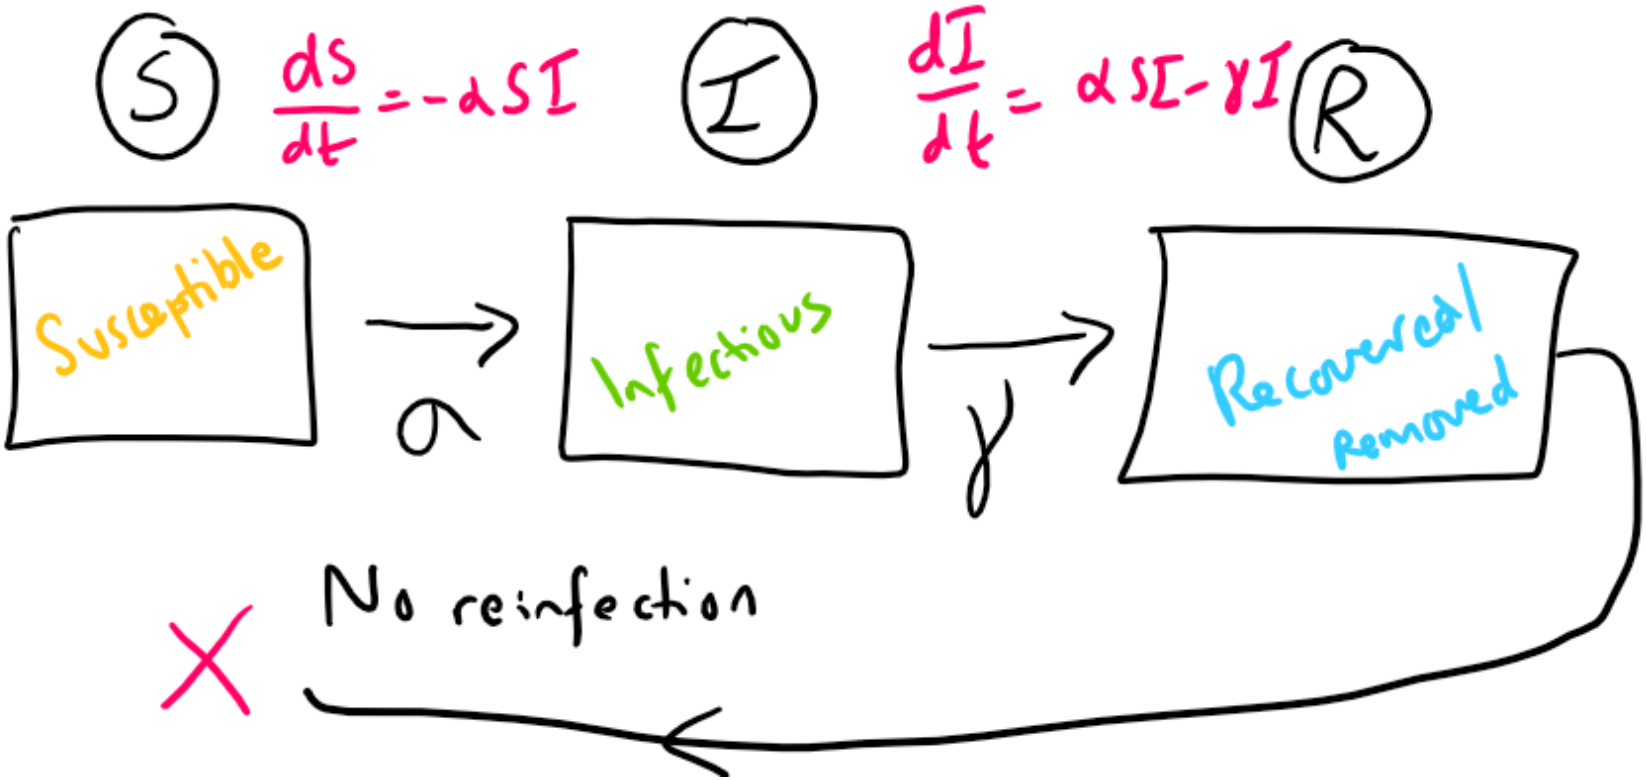

SIR – Equations

With these assumptions, the model equations read:

$$\begin{aligned} \frac{\mathrm{d}S}{\mathrm{d}t} &= -\alpha SI, \\ \frac{\mathrm{d}I}{\mathrm{d}t} &= \alpha SI - \gamma I, \\ \frac{\mathrm{d}R}{\mathrm{d}t} &= \gamma I. \end{aligned}$$

The equations conserve the total number of people, $S + I + R = N = \mathrm{Const.}$.
The initial condition for an outbreak is:

$$S(0) = S_0, \qquad I(0) = N - S_0, \qquad R(0) = 0.$$

With $S_0 = N - 1$ and $I(0) = 1$ we have the idea of 'patient zero'.

SIR – Equations

The constant γ is a rate (1/time). We write $\alpha = \beta/N$ so that β is also a genuine rate:

$$\begin{aligned} \frac{\mathrm{d}S}{\mathrm{d}t} &= -\frac{\beta}{N} SI, \\ \frac{\mathrm{d}I}{\mathrm{d}t} &= \frac{\beta}{N} SI - \gamma I, \\ \frac{\mathrm{d}R}{\mathrm{d}t} &= \gamma I. \end{aligned}$$

These equations can be reduced down to a single ODE:

$$\frac{1}{\gamma}\frac{\mathrm{d}u}{\mathrm{d}t} = \frac{\beta}{\gamma} - \frac{\beta}{\gamma}\frac{N}{S_0}\mathrm{e}^{-u} - u, \qquad u = \frac{\beta}{\gamma}\frac{R}{N}.$$

SIR – Equations

The model ODE can again be rescaled to give:

$$\frac{\mathrm{d}u}{\mathrm{d}\tau} = \frac{N}{S_0} - \frac{1}{R_0}u - \mathrm{e}^{-u}, \qquad \tau = \beta S_0 t/N.$$

This is an autonomous ODE, $\mathrm{d}u/\mathrm{d}\tau = f(u; S_0/N, R_0)$. Notice, $S_0/N < 1$. A fixed-point analysis yields two scenarios:

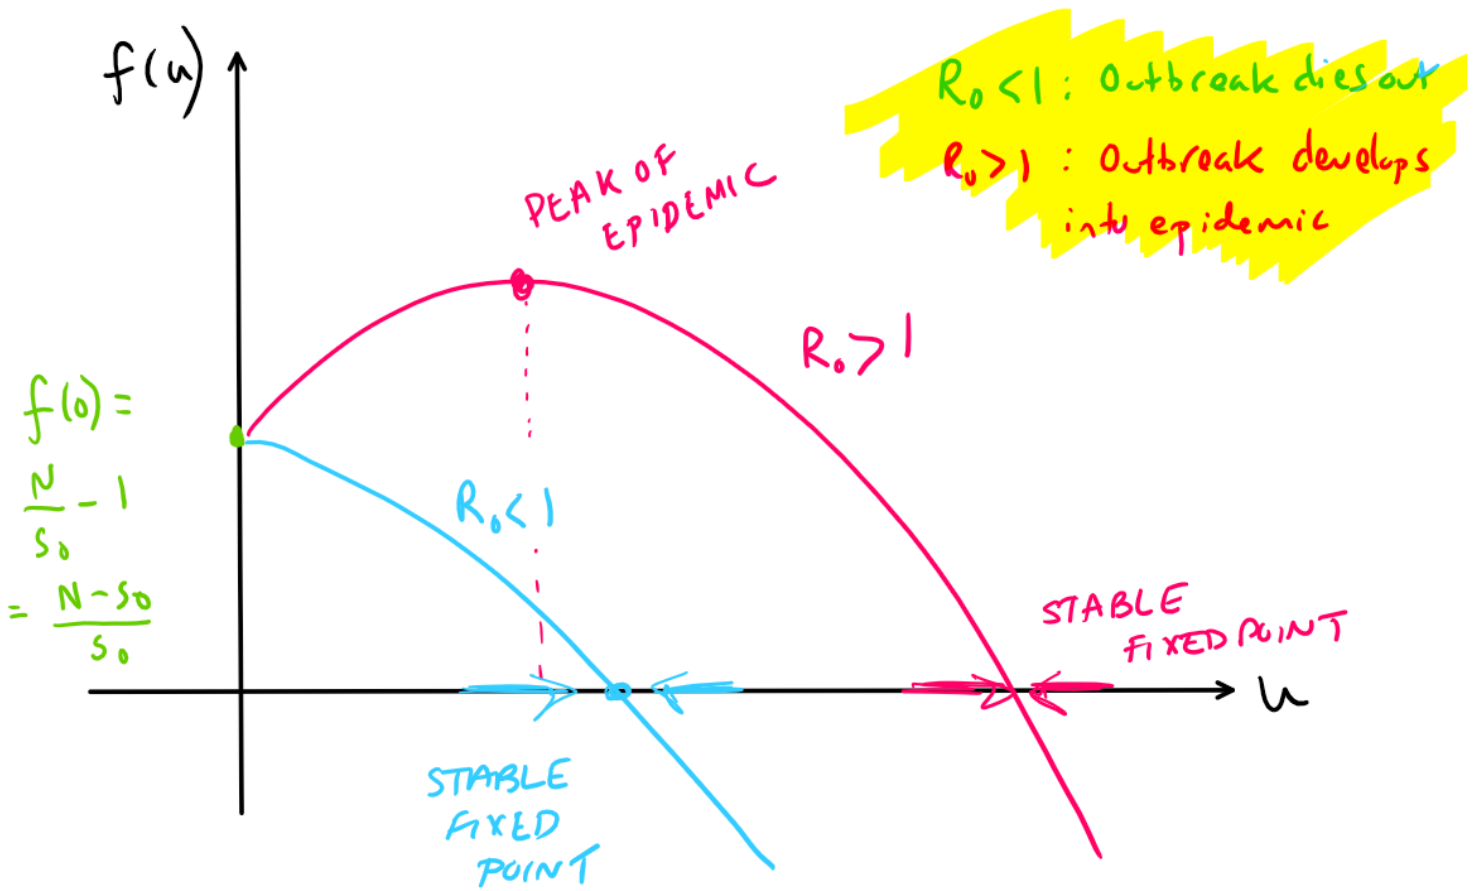

Basic Reproductive Number – Formula

The parameter R_0 is the dreaded **basic reproduction number:**

$$R_0 = \frac{\beta S_0}{\gamma N}.$$

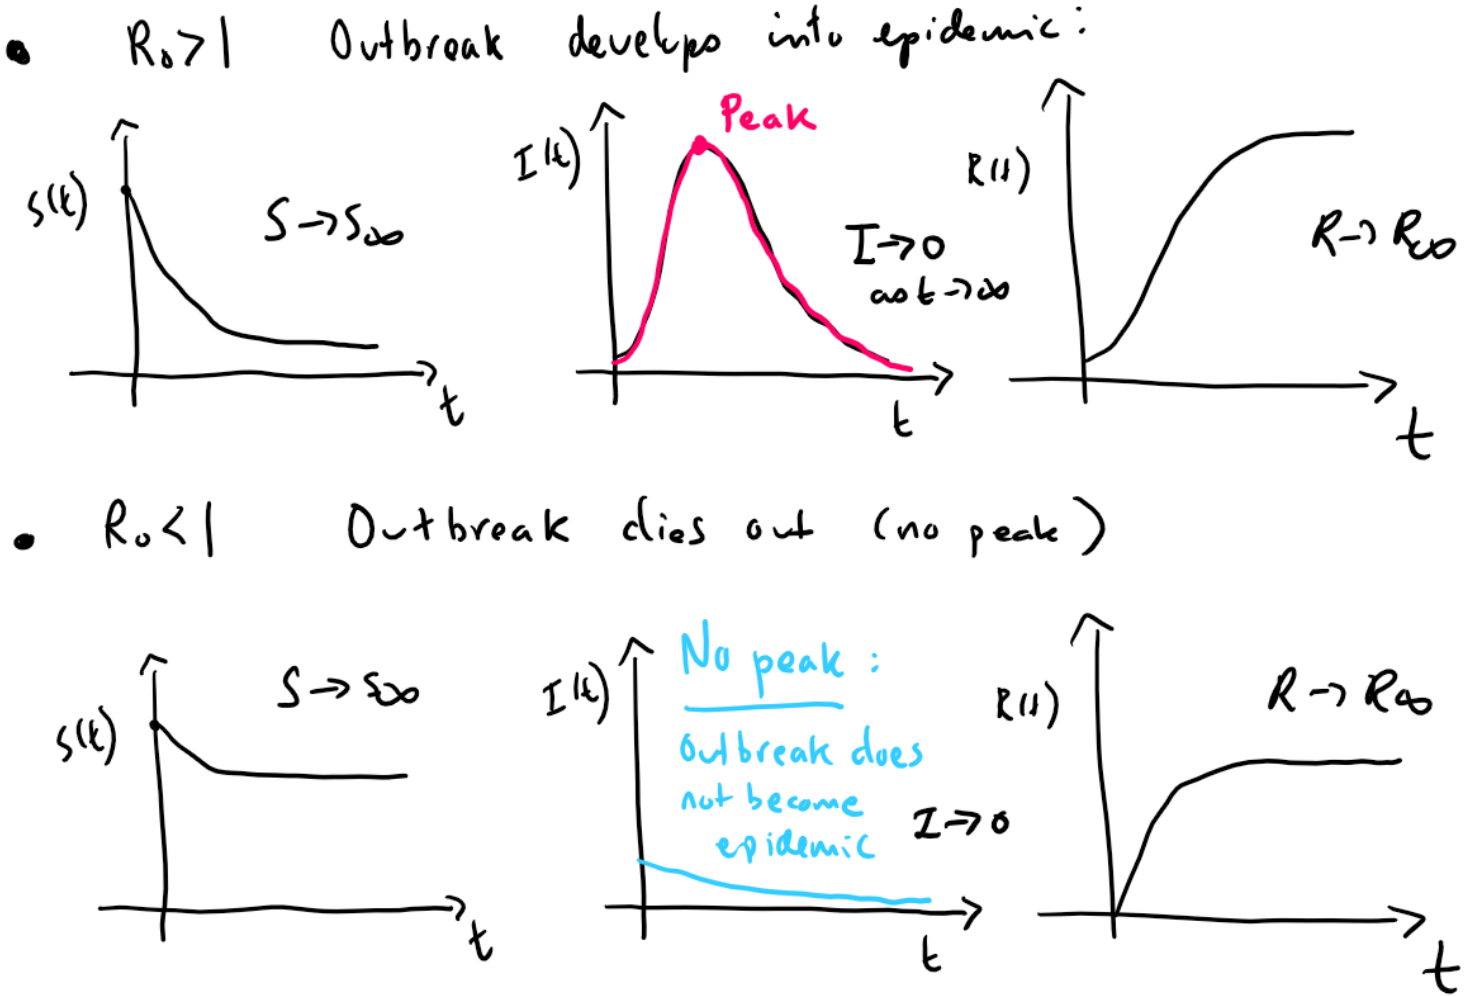

Basic Reproductive Number – Intuitive understanding

The familiar meaning of R_0 as '**the number of individuals that an infected person will go on to infect subsequently**' can by obtained integrating the I-compartment across a time interval Δt:

$$I(t+\Delta t) - I(t) = \gamma \Delta t \left[\frac{\beta}{\gamma} \frac{S(t)}{N} - 1 \right] I(t).$$

In the early stages of the outbreak, when the susceptible population is not depleted, $S(t) \approx S_0$, hence

$$I(t+\Delta t) - I(t) = \gamma \Delta t \left[\frac{\beta}{\gamma} \frac{S_0}{N} - 1 \right] I(t).$$

An infections person remains infectious for time γ^{-1}. Hence, take $\gamma \Delta t = 1$ to get

$$I(t+\Delta t) - I(t) \approx \left[\frac{\beta}{\gamma} \frac{S_0}{N} - 1 \right] I(t) = (R_0 - 1) I(t),$$

hence $I(t+\Delta t) = R_0 I(t)$. This gives the required interpretation of R_0: $I(t)$ individuals infect $R_0 I(t)$ further individuals.

Early-stage Exponential Growth, Late-stage burn-out

From the previous slide, in the early stage of the outbreak when $S(t) \approx S_0$, we have

$$I(t_n) = R_0^n I(0), \qquad t_n = n\Delta t.$$

In the late stages of the outbreak, there are not very many people left to infect, and the epidemic burns out. **Not everybody catches the disease:** $S_\infty/N = 1 - (R_\infty/N)$, where

$$\frac{R_\infty}{N} = \frac{u_*}{R_0}\frac{S_0}{N},$$

and where u_* is the fixed point of $\mathrm{d}u/\mathrm{d}\tau = f(u; S_0/N, R_0)$.

The fact that not everybody catches the disease can be thought of as 'burn-out'.

Late-stage burn-out, continued

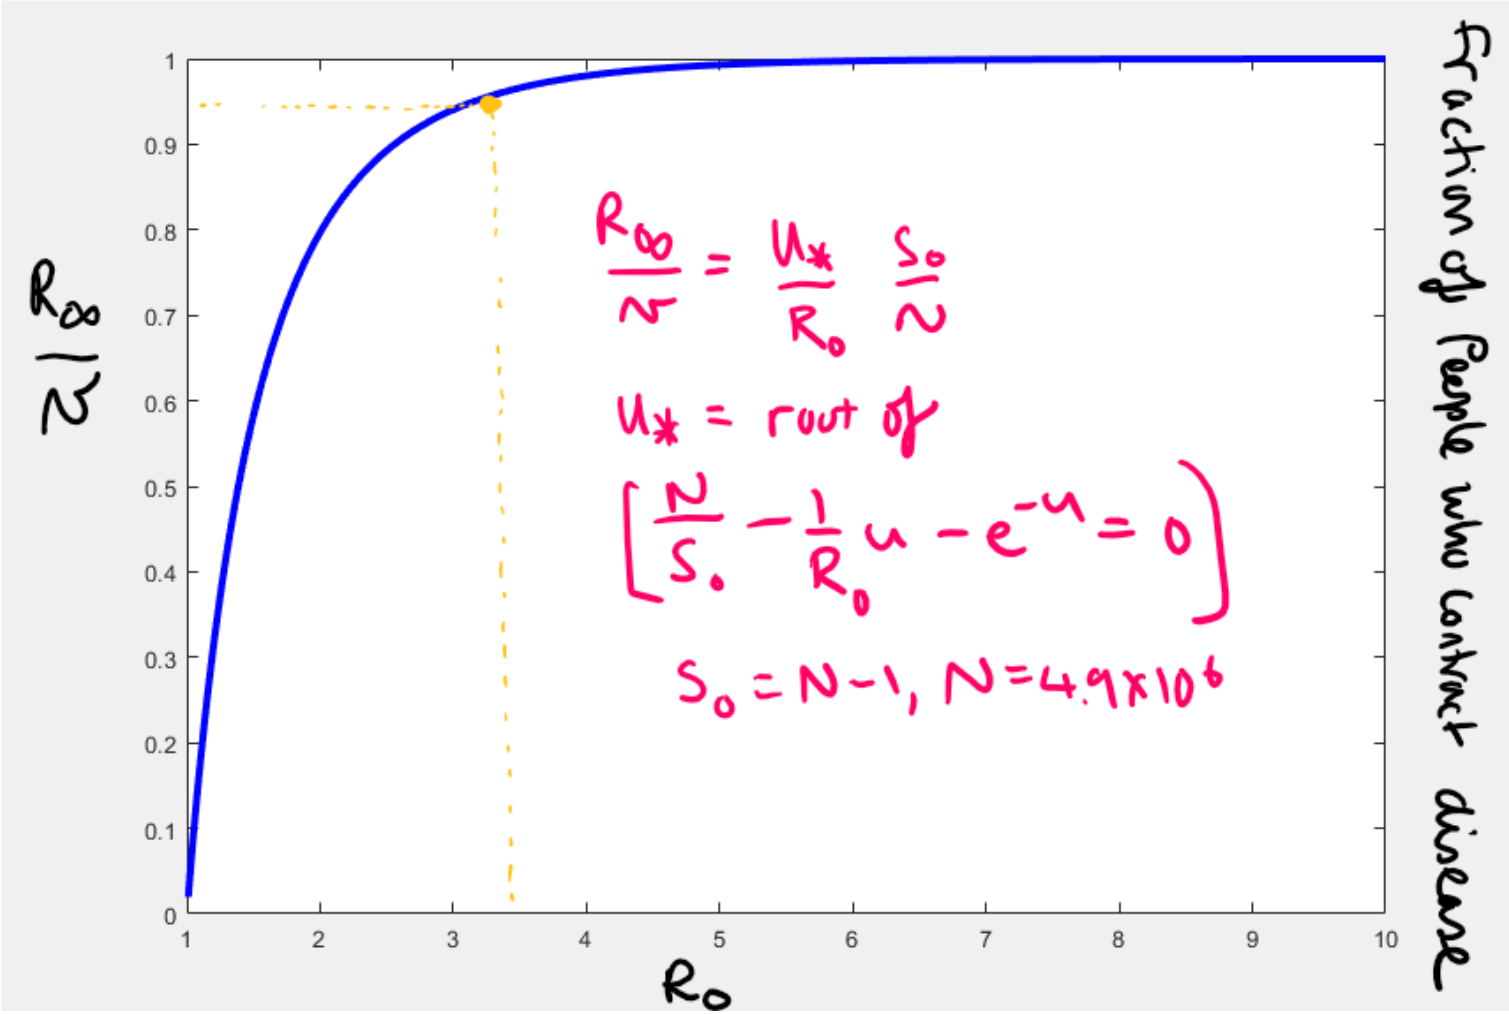

Vaccination Strategy

Suppose a vaccine for the disease exists. To control a future outbreak, we reduce the number of susceptible individuals from $S_0 = N - 1$ to

$$S_0 = N - \underbrace{1}_{\text{Patient Zero}} - \underbrace{fN}_{\text{Fraction of population vaccinated}}$$

Here, $0 < f < 1$. Starting-value for R_0 is:

$$R_0 = \frac{\beta}{\gamma}\frac{S_0}{N} = \frac{\beta}{\gamma}\frac{N-1}{N} \approx \beta/\gamma.$$

After mass vaccination, the new value is:

$$R_{eff} = \frac{\beta}{\gamma}(1-f).$$

To control the disease, we require a threshold value $R_{eff} = 1$ (ideally $R_{eff} < 1$). Therefore, the fraction of the population that needs to be vaccinated is $f = 1 - (1/R_0)$. When this number of people has been vaccinated we say **herd immunity** has been reached.

No Vaccine

In the absence of a vaccine, and with $S_0 = N - 1$, the way to control the spread of the epdiemic is to reduce $R_0 \approx \beta/\gamma$.

The parameter γ is fixed by biology, so we can only hope to control β.

Recall, $\mathrm{d}S/\mathrm{d}t = -(\beta/N)SI$. Hence, for a contagious disease β can be decomposed as

$$\beta = c \times p = (\text{Number of contacts of an individual per unit time}) \times (\text{Probability that a contact leads to infection})$$

Thus, the epidemic can be controlled by:

- Reducing contacts ($c \searrow$),
- Making contacts safer ($p \searrow$)

Extended Model for COVID-19

Part 2: We introduce an extended compartmental model for SARS-CoV-2 (virus) / COVID-19 (disease). The virus has multiple timescales, which means the basic SIR model needs more compartments:

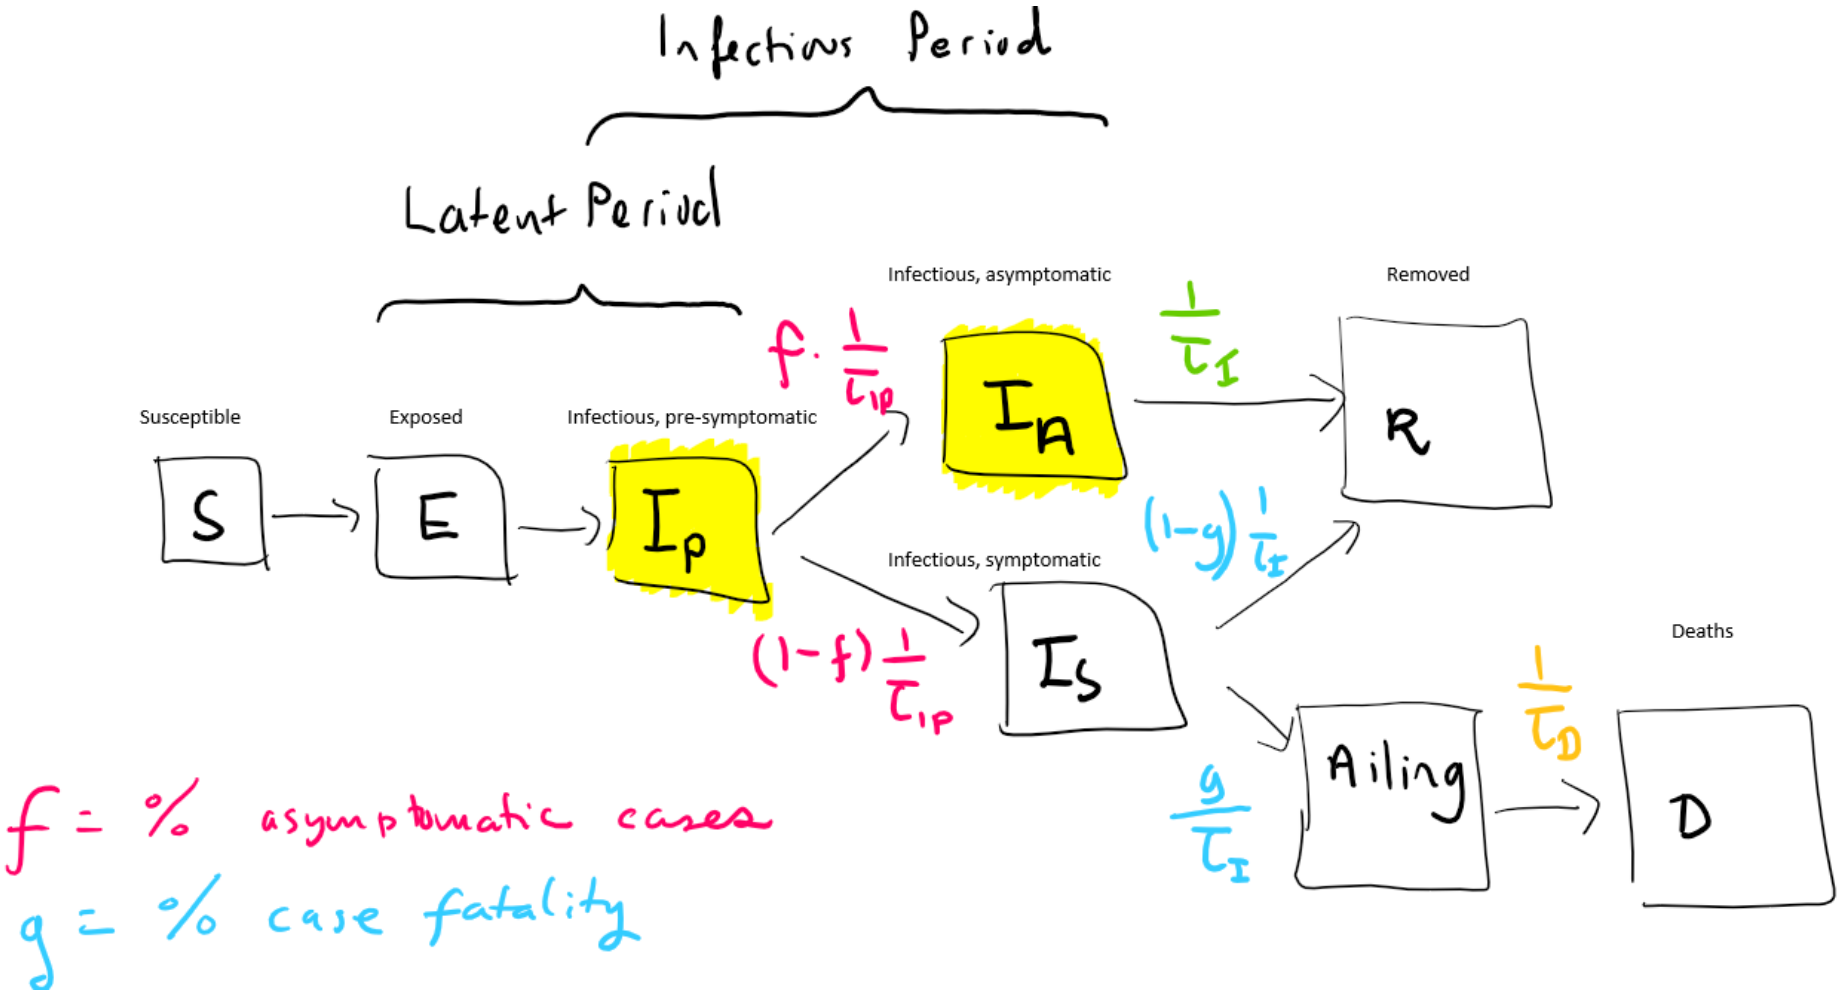

Diagnostic Compartments

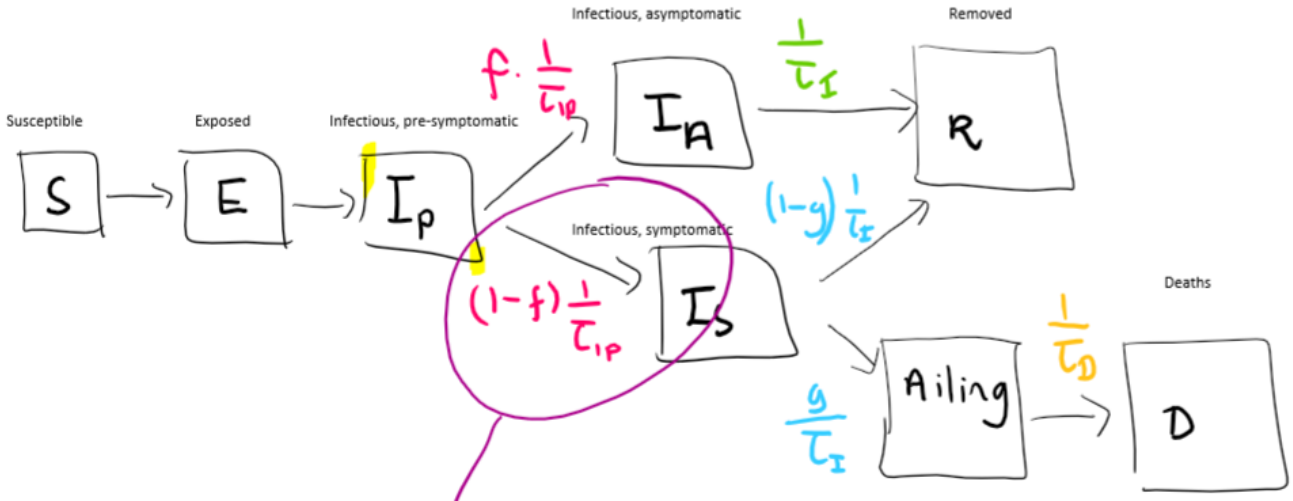

$$\frac{d}{dt}(\#\text{ Waiting For Test}) = \left(\begin{array}{l}\text{Rate at which people}\\ \text{go into the } I_S\\ \text{compartment}\end{array}\right) - \left(\#\text{ Waiting for Test}\right) / \text{Waiting Time}$$

$$\frac{d}{dt}\left(\#\text{Confirmed}\right) = \frac{1}{\text{Waiting Time}}\left(\#\text{ Waiting for test}\right)$$

Equations

$$F = \frac{c\,(q_P I_P + q_A I_A + q_S I_S)}{N} = \text{force of infection}$$

$$\begin{aligned}
\frac{dS}{dt} &= -\frac{cS\,(q_P I_P + q_A I_A + q_S I_S)}{N}, \\
\frac{dE}{dt} &= \frac{cS\,(q_P I_P + q_A I_A + q_S I_S)}{N} - \frac{1}{\tau_E} E, \\
\frac{dI_P}{dt} &= \frac{1}{\tau_E} E - \frac{1}{\tau_{IP}} I_P, \\
\frac{dI_A}{dt} &= \frac{f}{\tau_{IP}} I_P - \frac{1}{\tau_I} I_A, \\
\frac{dI_S}{dt} &= \frac{1-f}{\tau_{IP}} I_P - \frac{1}{\tau_I} I_S, \\
\frac{dR}{dt} &= \frac{1}{\tau_I} I_A + \frac{1-g}{\tau_I} I_S.
\end{aligned}$$

Dynamic Model
6 eqns

$$\begin{aligned}
\frac{d}{dt}(\text{Ailing}) &= \frac{g}{\tau_I} I_S - \frac{1}{\tau_D}(\text{Ailing}), \\
\frac{dD}{dt} &= \frac{1}{\tau_D}(\text{Ailing}), \\
\frac{d}{dt}(\text{Awaiting Test}) &= \frac{1-f}{\tau_{IP}} I_P - \frac{1}{\tau_T}(\text{Awaiting Test}), \\
\frac{dC}{dt} &= \frac{1}{\tau_T}(\text{Awaiting Test}).
\end{aligned}$$

Diagnostic Compartments

Set-up

'Patient zero' introduced at $t = -t_{offset}$:

$$S(t = -t_{offset}) = N - 1 = (4.9 \times 10^6 - 1) - 1, \qquad I_P(t = -t_{offset}) = 1.$$

Then, $t = 0$ is the day of the first recorded case of COVID-19 in Ireland (Feb 29th).

- Model fitted to the 'first wave': from February 29th 2020 ($t_j = 0$) and ending on May 17th 2020 ($t_j = 76$)
- Transmission probabilities re-scaled: $\beta_0 = cq_p$, $i_A = q_A/q_P$, $i_S = q_S/q_P$.
- Effective value of β ($=\beta_j$) introduced to take account of public health-interventions:

$$\beta_j = \begin{cases} \beta_0, & j < 13 \\ \beta_1, & 13 \leq j < 28, \\ \beta_2, & j \geq 28. \end{cases}$$

Optimization

Penalty function:

$$\Delta = \sum_{j=0}^{n} \left\{ \ln\left[C_{model}(t_j)\right] - \ln\left[C_{data}(t_j)\right] \right\}^2 + \sum_{j=0}^{n} \left\{ \ln\left(\frac{dC_{model}}{dt}\right)_{t_j} - \ln\left(\frac{dC_{data}}{dt}\right)_{t_j} \right\}^2$$
$$+ \sum_{j=j_{D1}}^{n} \left\{ \ln\left[D_{model}(t_j)\right] - \ln\left[D_{data}(t_j)\right] \right\}^2 + \sum_{j=j_{D2}}^{n} \left\{ \ln\left(\frac{dD_{model}}{dt}\right)_{t_j} - \ln\left(\frac{dD_{data}}{dt}\right)_{t_j} \right\}^2.$$

Optimization over parameters:

$$\Delta_0 = \min\left[\Delta\left(t_{offset}, \beta_0, \beta_1, \beta_2, f, g, \tau_E, \tau_{IP}, \tau_I, \tau_D, i_A, i_S\right)\right]$$

- Bounds on the search space are obtained from medical literature on COVID-19
- Optimization performed using MATLAB's simulated annealing method:
  - Initial guess
  - New draws; each draw involves solving the ODE system for selected parameters
  - A new draw is accepted if it lowers Delta,
  - A new draw is accepted with a certain probability if it increases Delta, this is to stop the algorithm getting stuck in local minima
  - Standard convergence criteria, these are checked for robustness

Results

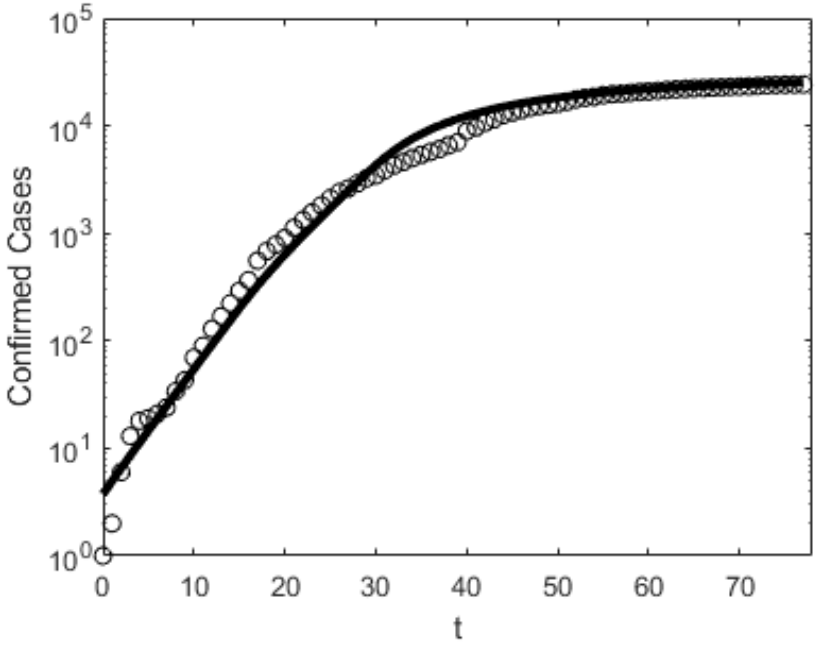
Confirmed Cases
t

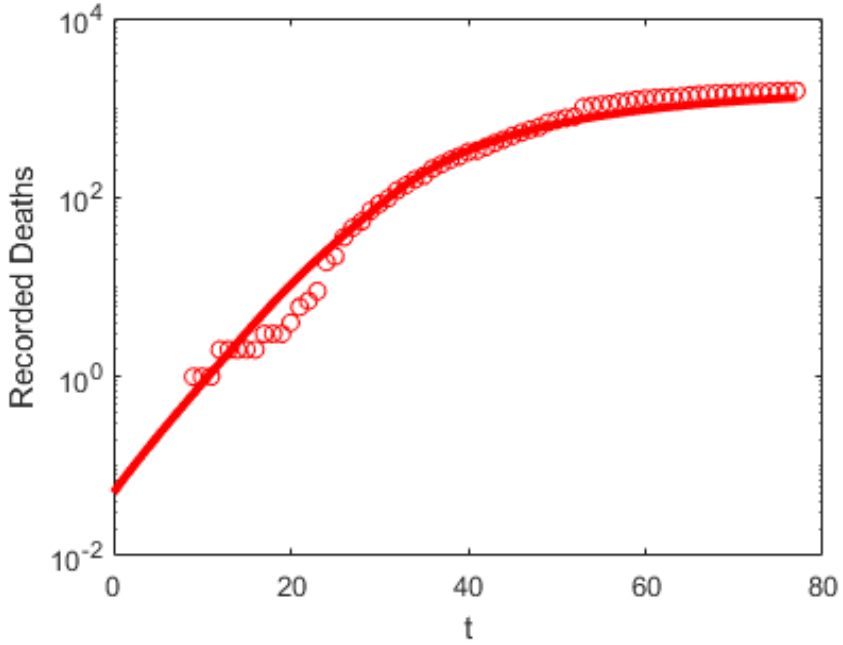
Recorded Deaths
t

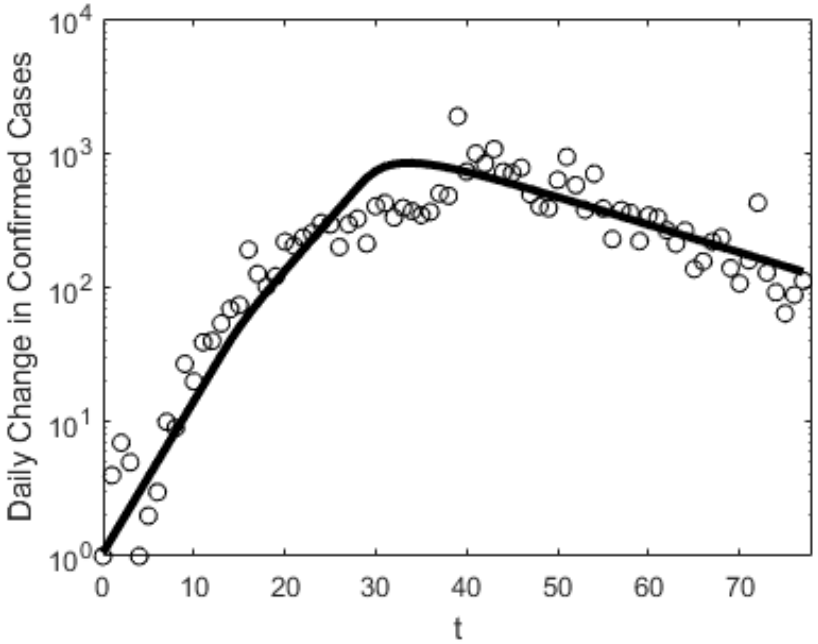
Daily Change in Confirmed Cases
t

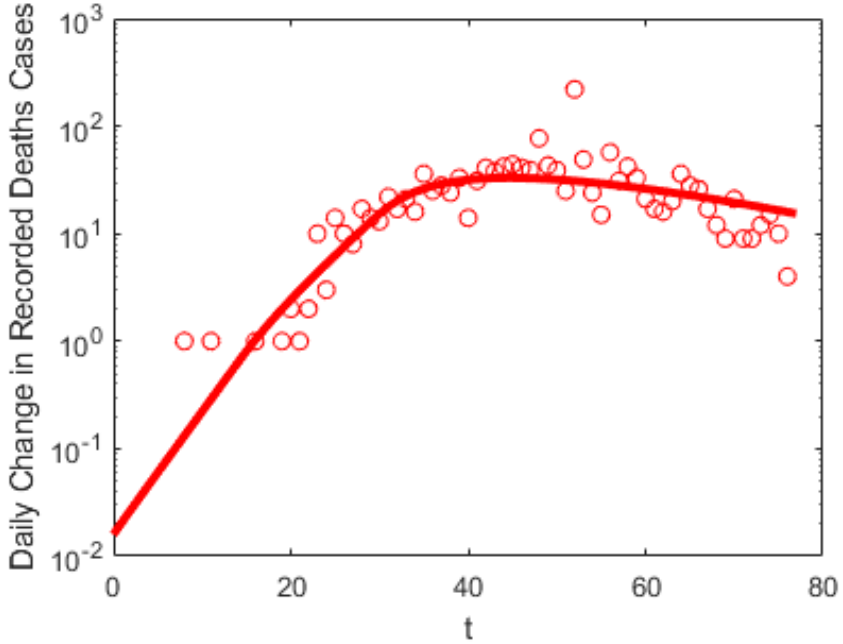
Daily Change in Recorded Deaths Cases
t

Confidence Intervals

Confidence intervals on the results are generated by bootstrapping. Significant spread in values, especially for fraction asymptomatic.

Table A.3: Optimal parameter values corresponding to the optimization problem (6). The reported confidence intervals use $n_{bootstrap} = 1000$.

	t_{offset} (days)	β_0 $(\text{days})^{-1}$	β_1 $(\text{days})^{-1}$	β_2 $(\text{days})^{-1}$	f	g	τ_E (days)
Lower Bound	7.9057	1.4679	0.9524	0.3205	0.2245	0.0500	3.6257
Best Estimate	9.9831	1.4695	1.1009	0.3576	0.3084	0.0599	3.7486
Upper Bound	9.9936	1.6961	1.2201	0.3908	0.4445	0.0614	3.7976

	τ_{IP}	τ_I	τ_D	τ_T	i_A	i_S
Lower Bound	1.1394	1.9160	14.9812	3.3146	0.3303	0.2741
Best Estimate	1.2938	2.1738	15.2210	3.8911	0.4096	0.3405
Upper Bound	1.3841	2.4337	16.5007	5.2055	0.4758	0.4367

R_0 for a multi-compartment model

R_0 can be computed for a multi-compartment model, in analogy with a simple SIR model. But now it is more difficult:

$$\frac{d}{dt}\delta\underline{x} = J\,\delta\underline{x}$$

around $\underline{x}_0 = (N, 0, 0, \dots)$

$$\frac{d}{dt}\begin{pmatrix}\delta E\\ \delta I_1\\ \vdots\\ \delta I_n\end{pmatrix} = \begin{pmatrix} & & \\ & & \\ & & \end{pmatrix}\begin{pmatrix}\delta E\\ \delta I_1\\ \vdots\\ \delta I_n\end{pmatrix}$$

$$= (F - V)\begin{pmatrix}\delta E\\ \delta I_1\\ \vdots\\ \delta I_n\end{pmatrix}$$

Notice

$R_0 \propto \beta_0$ .. $\frac{dS}{dt} = -\beta_0 \frac{S}{N} \sum_j I_j$.

$R_0 = \max \text{spec}\,(FV^{-1})$

For the model, the matrices F and V are given by:

$$F = \begin{pmatrix} 0 & \beta_0 & \beta_0 i_A & \beta_0 i_S \\ 0 & & 0 & 0 \\ 0 & & 0 & 0 \\ 0 & & 0 & 0 \end{pmatrix} \qquad V = \begin{pmatrix} \tau_E^{-1} & 0 & 0 & \\ -\tau_L^{-1} & \tau_{IP}^{-1} & 0 & 0 \\ 0 & -f\tau_{IP}^{-1} & \tau_I^{-1} & 0 \\ 0 & 0 & -(1-f)\tau_{IP}^{-1} & \tau_I^{-1} \end{pmatrix},$$

hence

$$R_0(\beta_0, i_S) = \beta_0 \left(\tau_{IP} + i_A f \tau_I + f(1-f) i_S \frac{\tau_I^2}{\tau_{IP}} \right).$$

R_0 for a multi-compartment model – Results

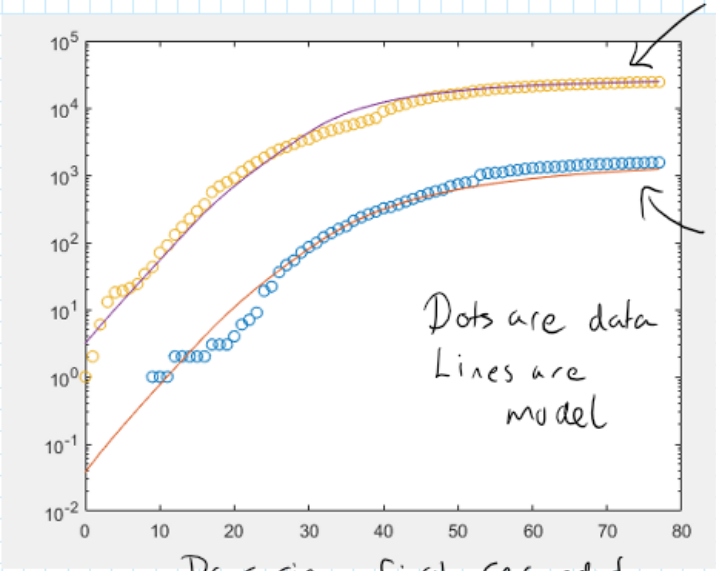

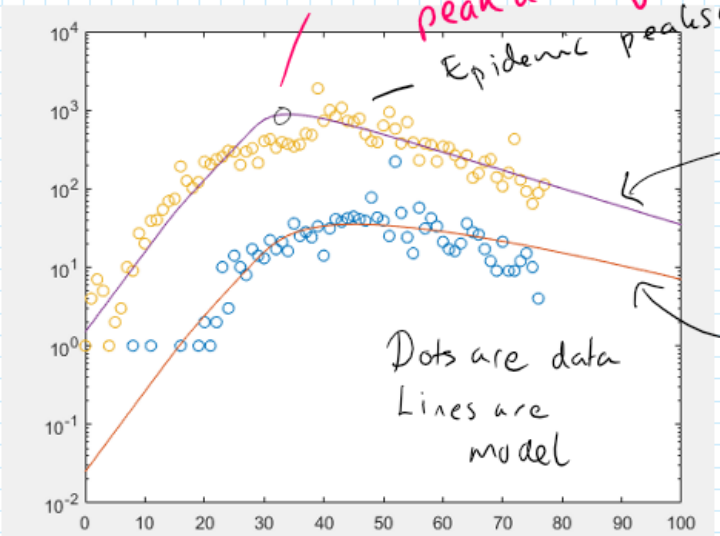

Parameters of epidemic

$R_0 = 3.6$ ... without case isolation

$R_0 = 2.6$ ... with case isolation

$R_0 = 1.8$ ... with case isolation + schools closed

$R_0 = 0.67$ .. with case isolation + lockdown

f

Fraction asymptomatic: 0.2

(prior assumption is $0.1 < f < 0.5$)

Confidence Intervals

Table A.4: Estimates of the basic reproduction number and the effective reproduction number under various NPIs, with lower and upper bounds generated from the confidence intervals in Table A.3.

Intervention	β (Best Fit)	i_S (Best Fit)	$R_0(\beta, i_S)$ (Lower Bound)	$R_0(\beta, i_S)$ (Best Fit)	$R_0(\beta, i_S)$ (Upper Bound)
No interventions, basic reproduction number	1.4695	1	2.7664	3.4495	4.9079
Case isolation, no other interventions	1.4695	0.3405	2.0798	2.6945	3.7549
Case isolation, schools and universities are closed, mass gatherings are banned	1.1009	0.3405	1.2880	2.0187	2.7204
Case isolation, all non-essential services and industries closed, travel restrictions	0.3576	0.3405	0.4847	0.6557	0.8324

Extremes

Vaccination strategy: $f \simeq 1 - \frac{1}{R_0} \simeq 71\%$

Optimal Control Theory

Part 3: We look out Optimal Control Theory. Given that NPIs start on Day 13, we ask, what would have been the optimal sequence of NPIs ('**counterfactual scenarios**').

Setup I

We solve the model with the estimated parameters. But now we take:

$$\beta(t) = \begin{cases} \beta_0, & t \in (-t_{offset}, t_0], \\ \beta_0[1 - u(t)], & t \in (t_0, T]. \end{cases}$$

- $u(t)$ is the effect of introducing **non-pharmaceutical interventions** to control the epidemic.
- For example, $u(t) = 0.5$ corresponds to a 50% reduction in daily contacts at a population level.
- This can be expected to carry a commensurate **economic cost**.

Setup II

We are going to specialize to piecewise-constant functions $u(t)$:

The piecewise-constant function $u(t)$ is characterized as follows:

$$u(t) = \begin{cases} u_1, & t_0 < t \le t_{s1}, \\ u_2, & t_{s1} < t \le t_{s2}, \\ \vdots & \\ u_n, & t_{s,n-1} < t \le T. \end{cases}$$

Here, n is an integer (one or greater), $u_1, \cdots, u_n$ are real numbers between zero and one, and $t_{s1}, \cdots, t_{s,n-1}$ are switching times (real numbers between t_0 T, such that $t_{s1} < \cdots t_{s,n-1} < T$).

The time T is taken to be a fixed **time horizon** of one year. Justification: expected development time for a vaccine.

Cost Function

We quantify the cost to the economy of implementing a sequence $\{u_1, \cdots, u_n\}$ of controls by the cost function

$$J = \sum_{i=1}^{n} u_i(t_{s,i} - t_{s,i-1}),$$

with the convention that $t_{s,0} = t_0$ and $t_{s,n} = T$. As such, J is a function of the real variables $\{u_1, \cdots, u_n, t_{s1}, \cdots, t_{s,n-1}\}$.

Idea:

- Non-pharmaceutical interventions (NPIs) reduce daily contacts
- Assume that economic activity is proportional to these contacts.
- Hence, a linear relationship between the level of the NPIs and the cost to the economy
- Intense NPIs may carry a disproportionate cost to the economy, linearity may break down as $u_i \to 1$. So we will look at a quadratic cost function as well.

State Constraints

We furthermore insist on an optimal control problem where human life is put on a very high footing. As such, we propose the following additional constraints on the optimal control problem:

$$kI_S(t) \leq B, \qquad t \in (-t_{offset}, T]$$

In this context, the positive constant k may be thought of as the percentage of symptomatic cases who require a hospital bed (or a bed in ICU) at time t, and the positive constant B represents a corresponding capacity limit (number of hospital beds, number of ICU beds, etc.). Finally, we impose the additional constraint

$$\frac{dI_S}{dt} \leq 0, \qquad t = T,$$

which is true if and only if $\tau_{IP}^{-1}(1-f)I_P - \tau_I^{-1}I_S \leq 0$ **at** $t = T$. This rules out any 'optimal' strategy in which a large epidemic peak would occur just beyond the horizon at $t = T$. The exclusion of such strategies is desirable for public-health reasons.

Optimal Control Problem

Under the epidemic modelled by the SEIR model, over a fixed time horizon T, compute the minimizer of

$$\min_{n\geq 1}\left\{\min_{\substack{u_1,\cdots,u_n,\\ t_{s1},\cdots,t_{s,n-1}}}\left[\sum_{i=1}^{n} u_i(t_{s,i}-t_{s,i-1})\right]\right\},$$

subject to the given state constraints.

Theory:

- Optimal Control problems are usually solved with the Pontryagain Maximum Principle (PMP).
- Difficult to apply with state constraints.
- Solution with state constraints typically involves 'bounardy arcs', where the control $u(t)$ guides the solution along the boundary of the state constraint.
- Unimaginable in a public-health context.
- So the piecewise constraint function $u(t)$ is preferred, combined with a numerical optimization approach.

Methodology

- We solve the SEIR model for an arbitrary sequence of controls $\{u_1, \cdots, u_n, t_{s1}, \cdots, t_{s,n-1}\}$.
- In Matlab / Octave programming we execute the following command:

```
penalty=ode_solve_seir(u);
```

  Here, `ode_solve_seir` is a Matlab / Octave ODE45 solver which takes in the input $\mathtt{u} = \{u_1, \cdots, u_n, t_{s1}, \cdots, t_{s,n-1}\}$, solves the SEIR model out to the final time $t = T$, and returns the penalty function J.
- The state constraints are taken into account by adjoining to the penalty function an additional term

$$\left[\tanh\left(\frac{0.016\max(I_S) - 300}{0.1}\right) + 1\right]\left[0.016\max(I_S) - 300\right]^2 + \left[\tanh\left(\frac{(dI_S/dt)_{t=T}}{0.1}\right) + 1\right]\left[(dI_S/dt)_{t=T}\right]^2. \quad (1)$$

Methodology – Hierarchy of function calls

- We next call built-in optimization algorithms in Matlab / Octave which optimize the penalty function `penalty=ode_solve_seir(u);` over the set of all feasible controls $\{u_1, \cdots, u_n, t_{s1}, \cdots, t_{s,n-1}\}$.
- We use meta-heuristic algorithms for finding global optima: simulated annealing and particle-swarm optimization.
- Each has independent stopping criteria – we vary these to check for robustness.
- Convergence to the global optimum is guaranteed theoretically but we can't guarantee that the numerical method has reached the global optimum – but we have done our best!
- Program structure where functions can be called hierarchically is very powerful and makes coding easy. Repository:

`https://github.com/ainebyrne/ONaraighByrne_COVID19_optimal_control.`

Presentation of Results

- Presentation of results: systematically look at $n = 2$ controls (one switch), $n = 3$ controls (two switches), ...
- $n = 2$ is revealing because optimization algorithms can be compared to a brute-force approach (small parameter space).
- Key parameter of interest is u_{max}, $0 \leq u(t) \leq u_{max} \leq 1$ – we impose a cutoff $u(t) \leq u_{max}$ and see how the optimal strategy depends on that cutoff.
- This parametrizes the maximum intensity of NPIs that are socially acceptable.

Another cutoff

In passing – we impose another cutoff, namely $F \leftarrow 0$ if $I_j < 1$. This bridges the gap between the continuum SEIR model and the discrete nature of the population.

This introduces a jump discontinuity in the model, which might be important.

Results: $n = 2$

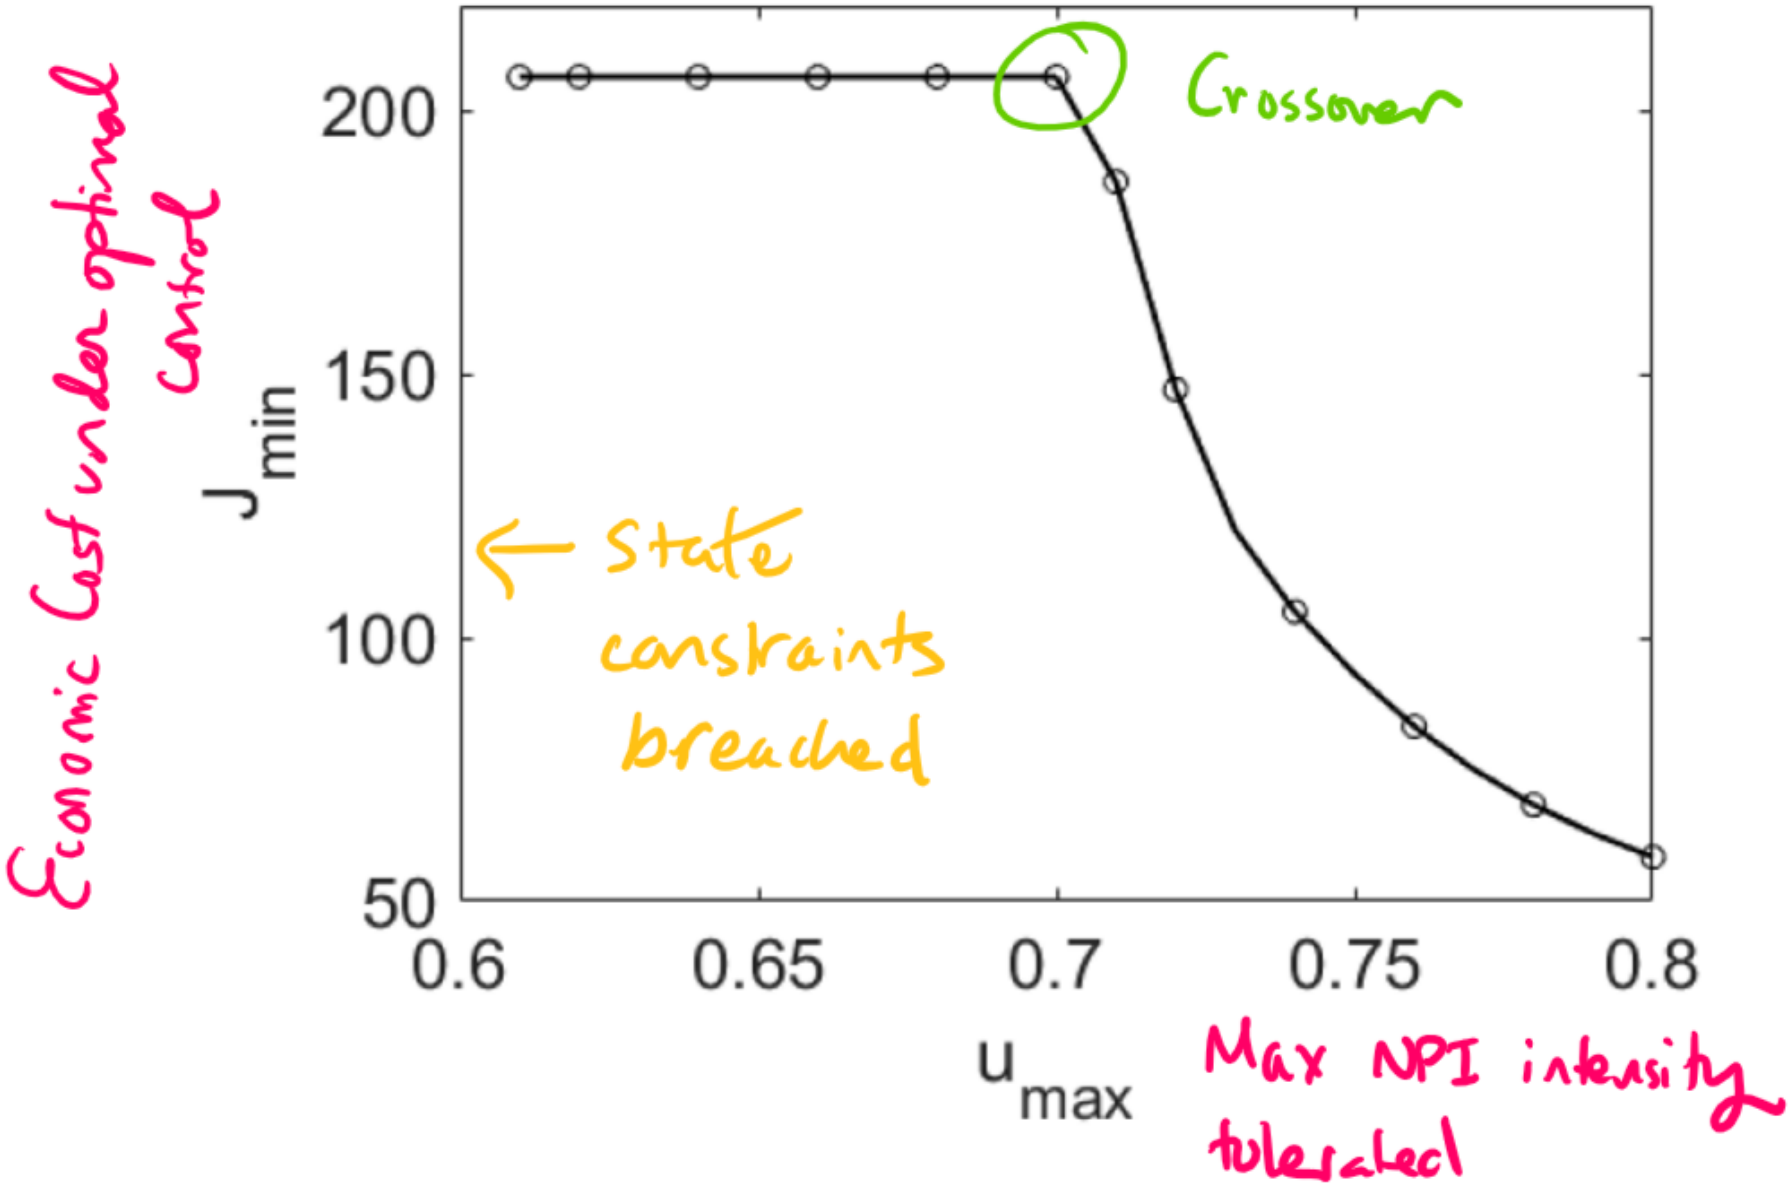

Plot of J_{min} as a function of u_{max}, with $n = 2$. The plot range is $u_{max} \in [0.61, 0.8]$.

Crossover

The crossover occurs as the optimal strategy switches from elimination to mitigation.

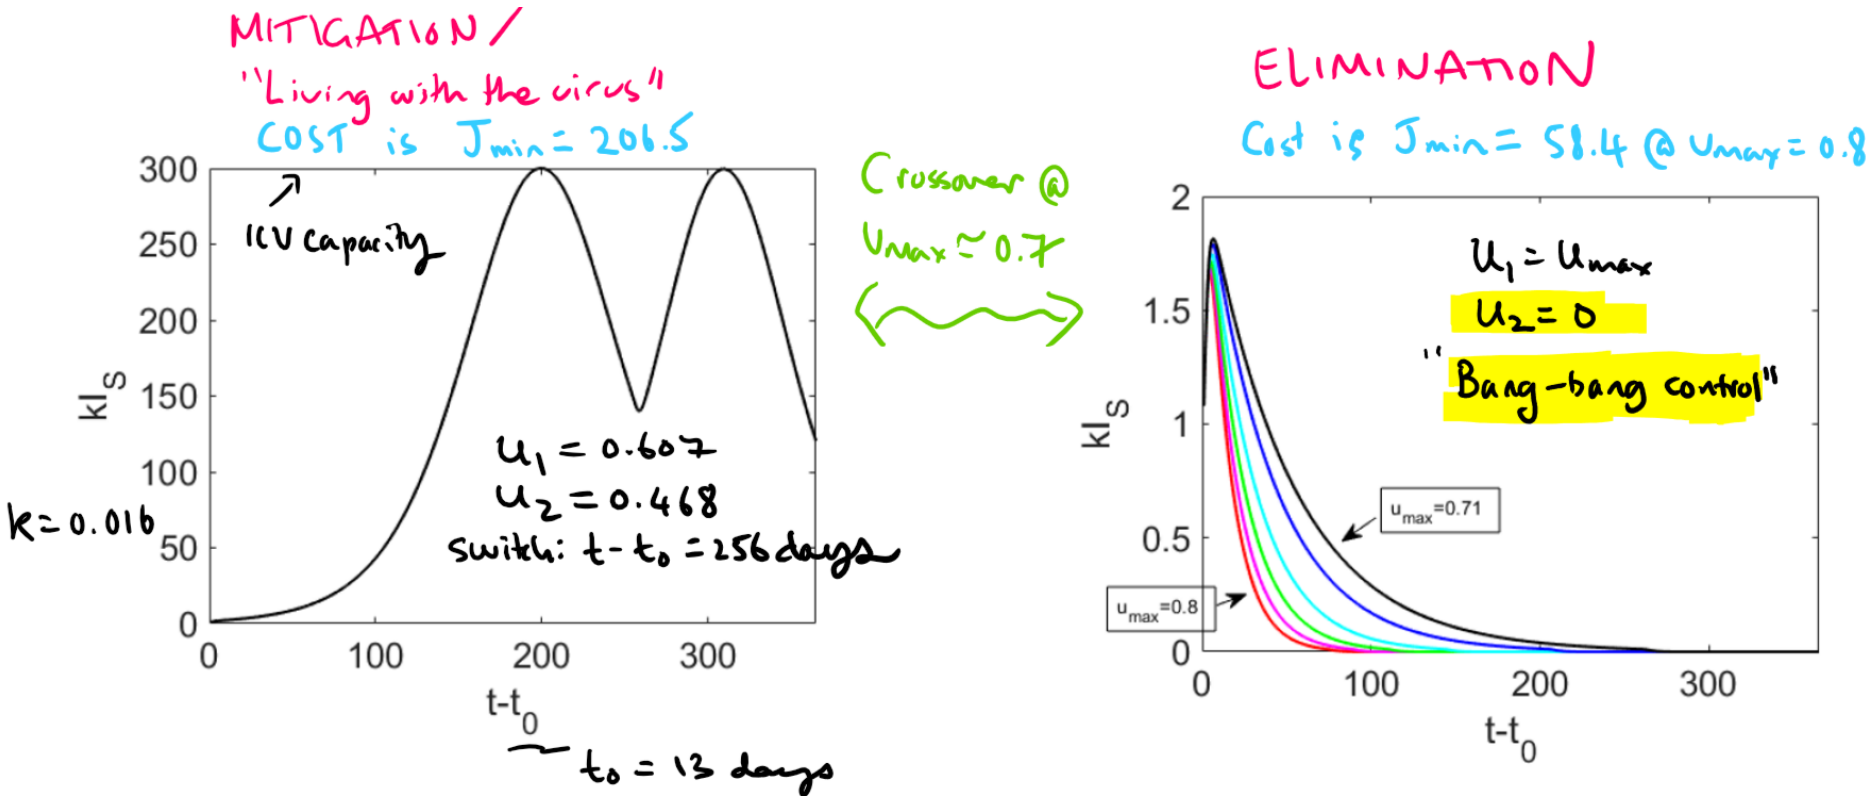

Elimination is almost 4 times cheaper than mitigation.

Discussion

- For $u_{max} \gtrsim 0.7$ the optimal control is 'bang-bang' – on at the max and then off. **This is common in optimal control problems with a linear cost function.**
- For $u_{max} \lesssim 0.7$ the optimal trajectory reaches the boundary of the state constraint – similar to a 'boundary arc' in the theoretical approach.
- The bang-bang control with $u_{max} \gtrsim 0.7$ is considerably cheaper.

Validation via brute-force approach

Fix t_{s1} but otherwise vary u_1 and u_2 to produce

$$\Phi(u_1, u_2) = \texttt{penalty}([u_1, u_2, t_{s1} = \texttt{fixed}]).$$

Plot $\Phi(u_1, u_2)$, marking in the **state constraints**:

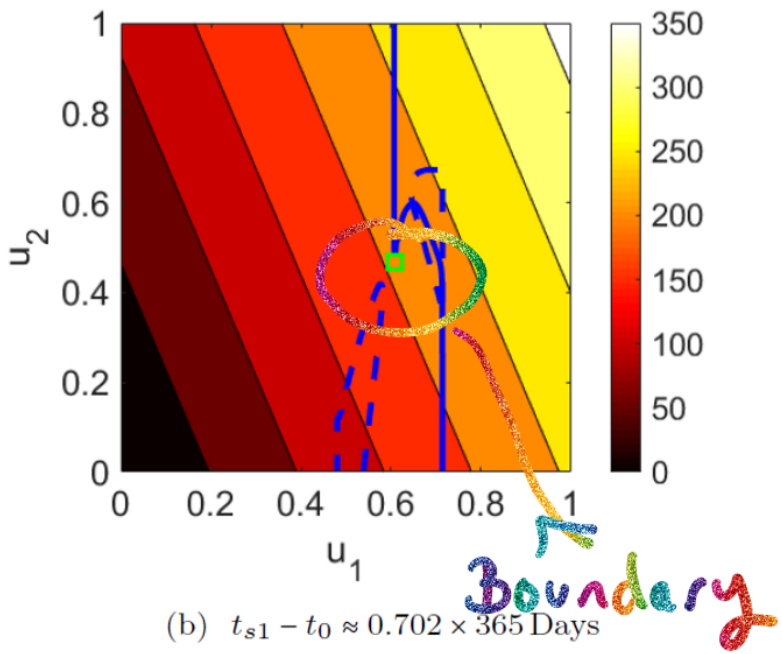

(b) $t_{s1} - t_0 \approx 0.702 \times 365$ Days

Optimal switching time for v_{max} between 0.61 and 0.7

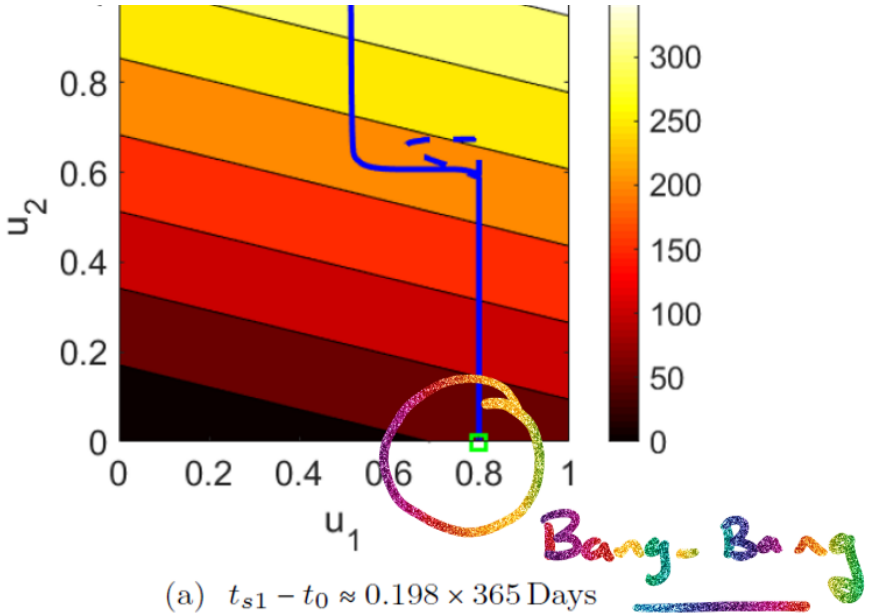

(a) $t_{s1} - t_0 \approx 0.198 \times 365$ Days

Optimal switching time for $v_{max} = 0.8$

Location of minima **agrees** with numerical simulated annealing.

Multiple Switches

With multiple switches, the cheapest option is still elimination, but the crossover occurs at much higher values of u_{max}, meaning both options become comparable:

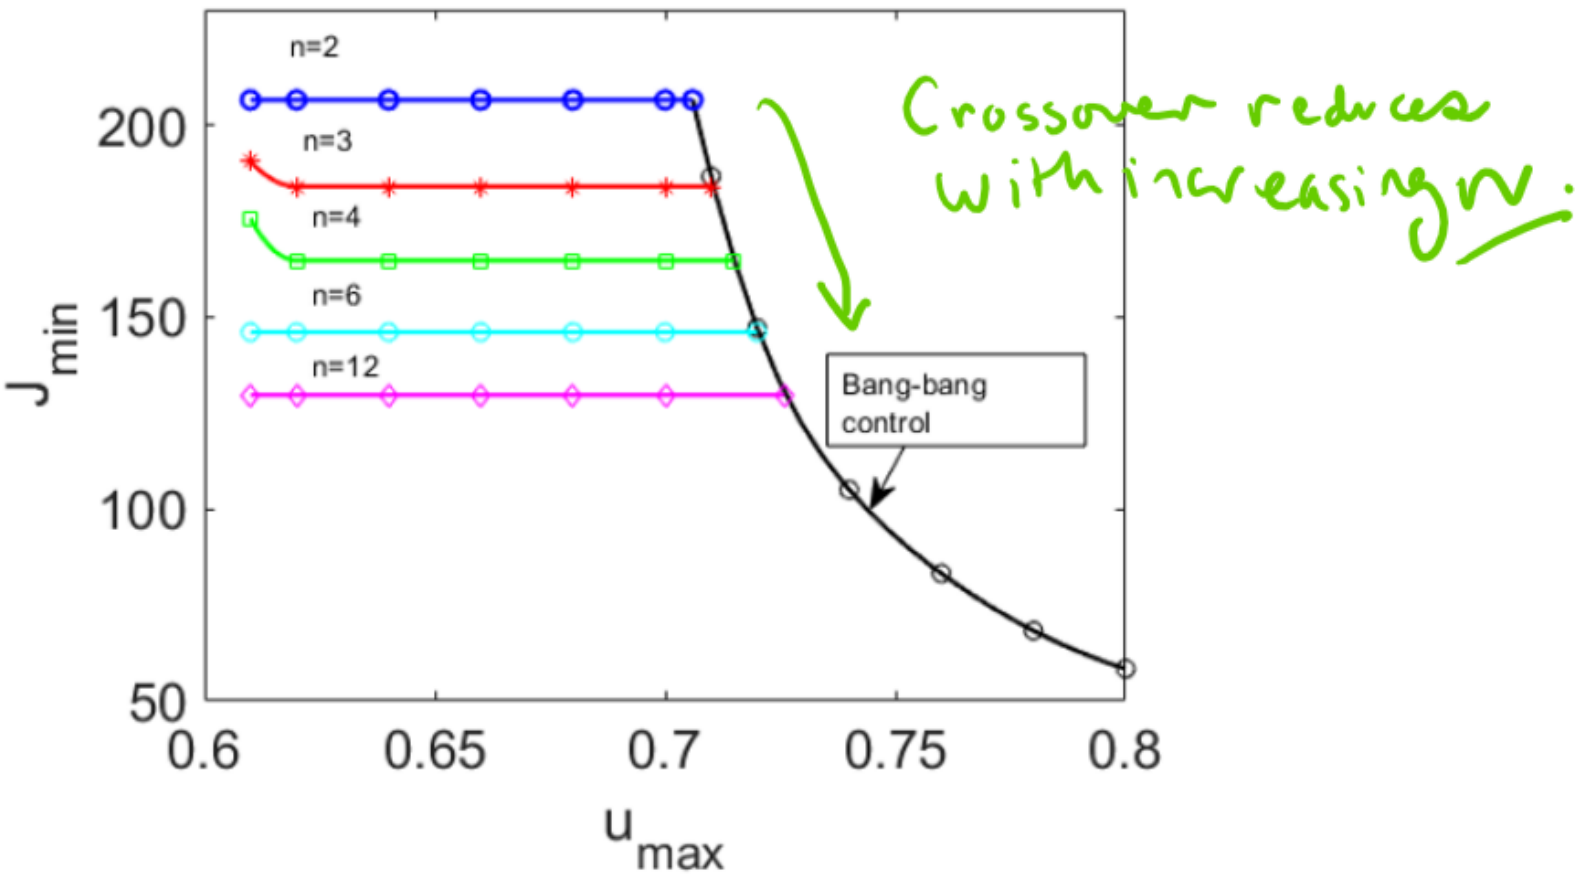

Figure 8: Plot of J_{min} as a function of u_{max}, with $n \geq 2$.

Quadratic Penalty Function

With a quadratic penalty function, elimination is still the cheaper option, but the crossover moves to even higher values of u_{max}:

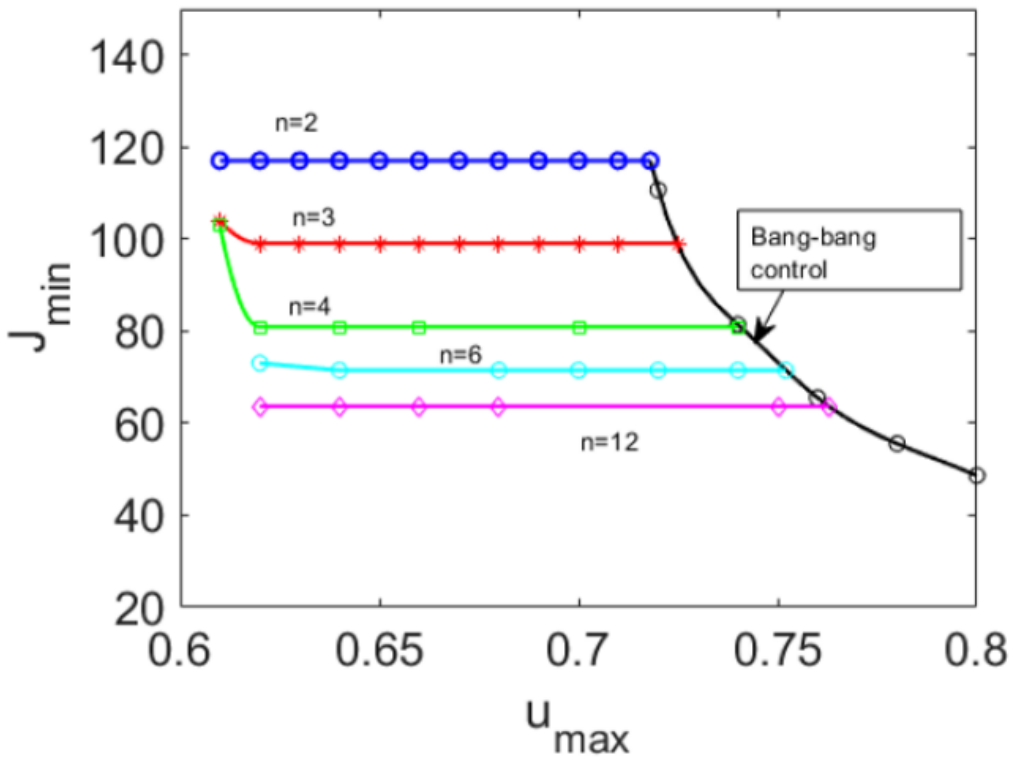

$$\gamma = \sum_{i=1}^{\hat{}} u_i^2 \left(t_{s,i} - t_{s,i-1} \right)$$

Figure 11: Plot of J_{min} as a function of u_{max}, with $n \geq 2$, quadratic cost function (Equation (21)).

Mitigation Strategies revisited

- The mitigation strategies are prone to 'overshoot': **a small deviation in the controls leads to large departures from the state constraints**.
- Due to the exponential growth away from the $I = 0$ unstable fixed point in the model $R_0 > 1$.
- Concerning from a public-health perspective.

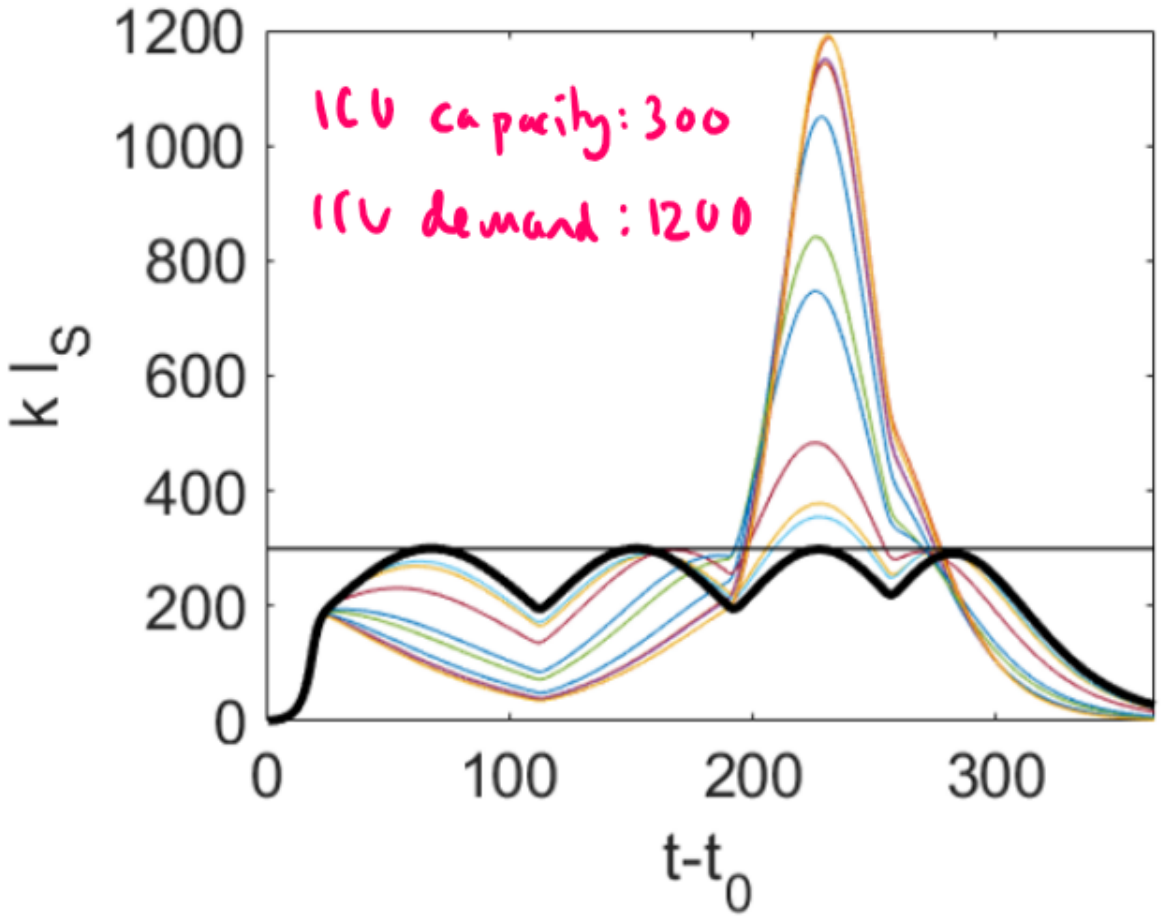

Comparison between counterfactual optimal strategy and reality

- In Ireland, NPIs ramped up in mid-march, starting on March 13th with school closures and culminating on March 29th with 'stay-at-home orders'.
- In reality, it would have been better to implement the maximum-intensity NPIs from the very beginning (feasibility?).

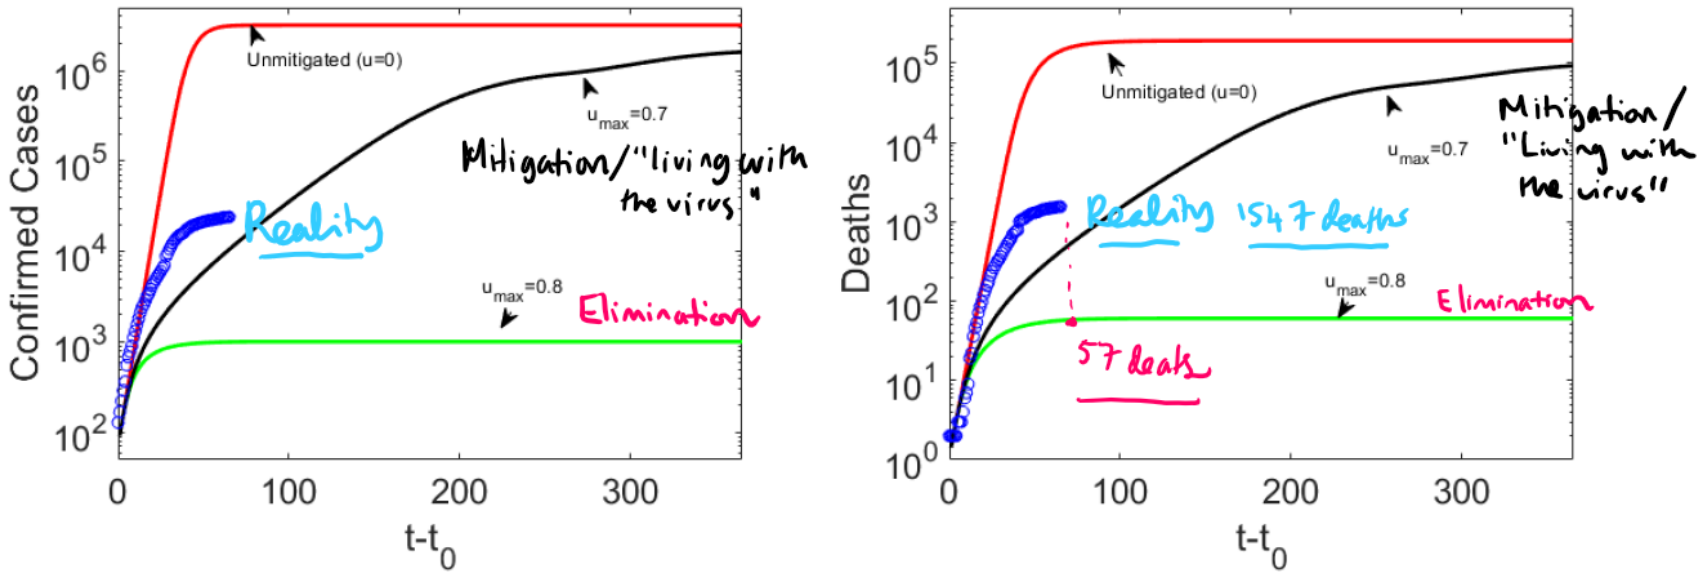

(a) Confirmed Cases (b) Deaths

Figure 6: **Time evolution of the outbreak for three scenarios: disease outbreak with controls $u_{max} = 0.7$ and $u_{max} = 0.8$ and also, disease outbreak with no controls. The $-\circ$ curves show the actual course of the outbreak over the time interval for which the model is fitted to the data.**

Continuation

- Reality: On day $t - t_0 = 65$ there were 1547 deaths.
- Elimination strategy: On the same day there would have been 57 deaths.
- The difference is the intensity of the NPIs at the beginning of the outbreak.
- Notice also, the reality is a very strange mixture of mitigation and elimination – understandable, as public-health doctors were only learning about the disease at the start of the year.

A very simple lesson

- From the data: Maximum-intensity NPIs at the start of an epidemic saves lives and is a cheaper option.
- Reason is exponential growth – early-stage NPIs for COVID-19 in Ireland in March 2020 still had $R_{eff} > 1$, hence exponential growth in cases and deaths.

A more subtle lesson

Suppose I eliminate the disease ($R_0 \to R_{eff}$) as soon as I hear about patient zero ($I = 1$). Suppression in an **SIR** model: the time taken to drive I from $I = 1$ to $I = 1/2$ is

$$t_1 = \frac{\ln 2}{\gamma(1 - R_{eff})}, \qquad R_{eff} < 1.$$

If instead I wait until there are n active cases before introducing NPIs, elimination time is:

$$t_n = \frac{\ln(2n)}{\gamma(1 - R_{eff})}, \qquad R_{eff} < 1.$$

If I wait until $n = 1000$ acrive cases, then the difference is $t_n/t_1 = 10.97$, i.e. 11 times longer.

The commensurate cost to the economy is 11 times greater.

And, going back to lives and not money, as the number of deaths is proportional to the number of infections,

The number of deaths is (at least) 1000 times greater..

Summary

- In the absence of therapies or vaccines, a population-level reduction in daily contacts is one of only a very limited number of available measures to control the outbreak of a contagious disease.
- Reduction in contacts costs money.
- We have used optimal-control theory to minimize this cost, while controlling the spread of the epidemic such that ICU demand never exceeds capacity.
- Model calibrated for COVID-19 epidemic in Ireland ('first wave': March-May 2020).
- Results of counterfactual scenarios: **elimination of the virus is the cheapest option**, i.e. cheaper than mitigation.
- Result makes intuitive sense – suppress exponential growth before the numbers get out of control. Easiest at the very beginning of the outbreak when case numbers are low.

Take-home message

Elimination is best but it may be infeasible (??). Instead, I would prefer the following take-home message:

- Exponential growth is difficult to intuit!
- But from the simple argument and the SIR model, postponing a 'lockdown' is costly in terms of blood and treasure.
  - The cost in treasure scales logarithmically with the number of active cases at lockdown.
  - The cost in blood scales linearly.