Optimal Control of an SIR Epidemic via Contact Controls

David I. Ketcheson

Go20 Conference

Malta

May 25th, 2023

The SIR model

- Assumes recovered patients are immune

- Simple, but effective in predicting basic pattern of an epidemic

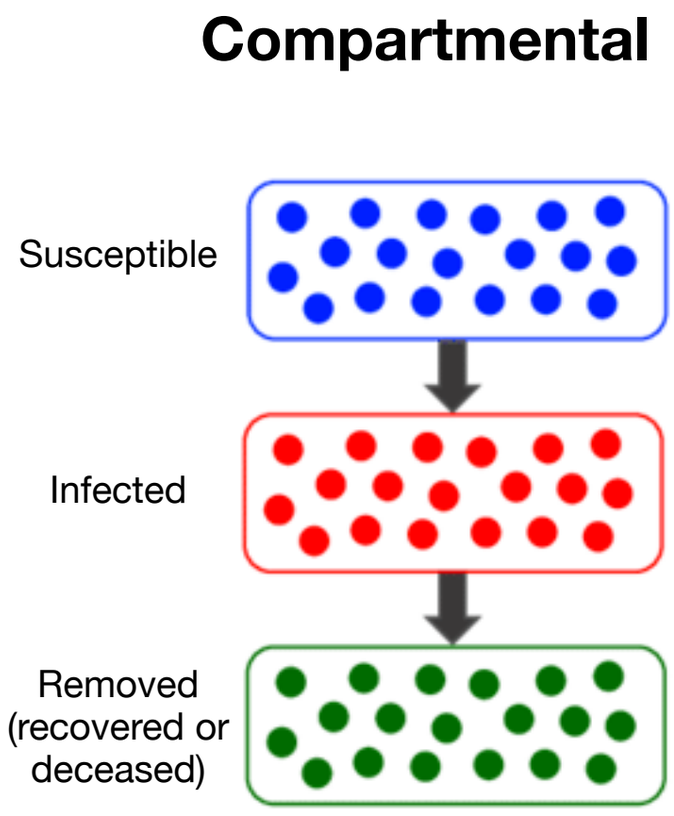

The SIR model

$$x'(t) = -\beta y x$$
$$y'(t) = \beta y x - \gamma y$$
$$z'(t) = \gamma y$$

$$x + y + z = 1$$

Kermack & McKendrick, 1927

x: Susceptible fraction
y: Infected fraction
z: Removed fraction

β : Contacts per person per day

γ^{-1} : Length of infectious period

$\sigma_0 = \beta/\gamma$: Average number of contacts by one infected individual

(# of contacts per day) x (# of days contagious)

Basic reproduction number

Dynamics

Can reduce system to 2 equations:

$$x'(t) = -\gamma\sigma_0 yx$$

$$y'(t) = \gamma\sigma_0 yx - \gamma y$$

- Equilibrium: $y = 0$
- Stable if $x \leq 1/\sigma_0$

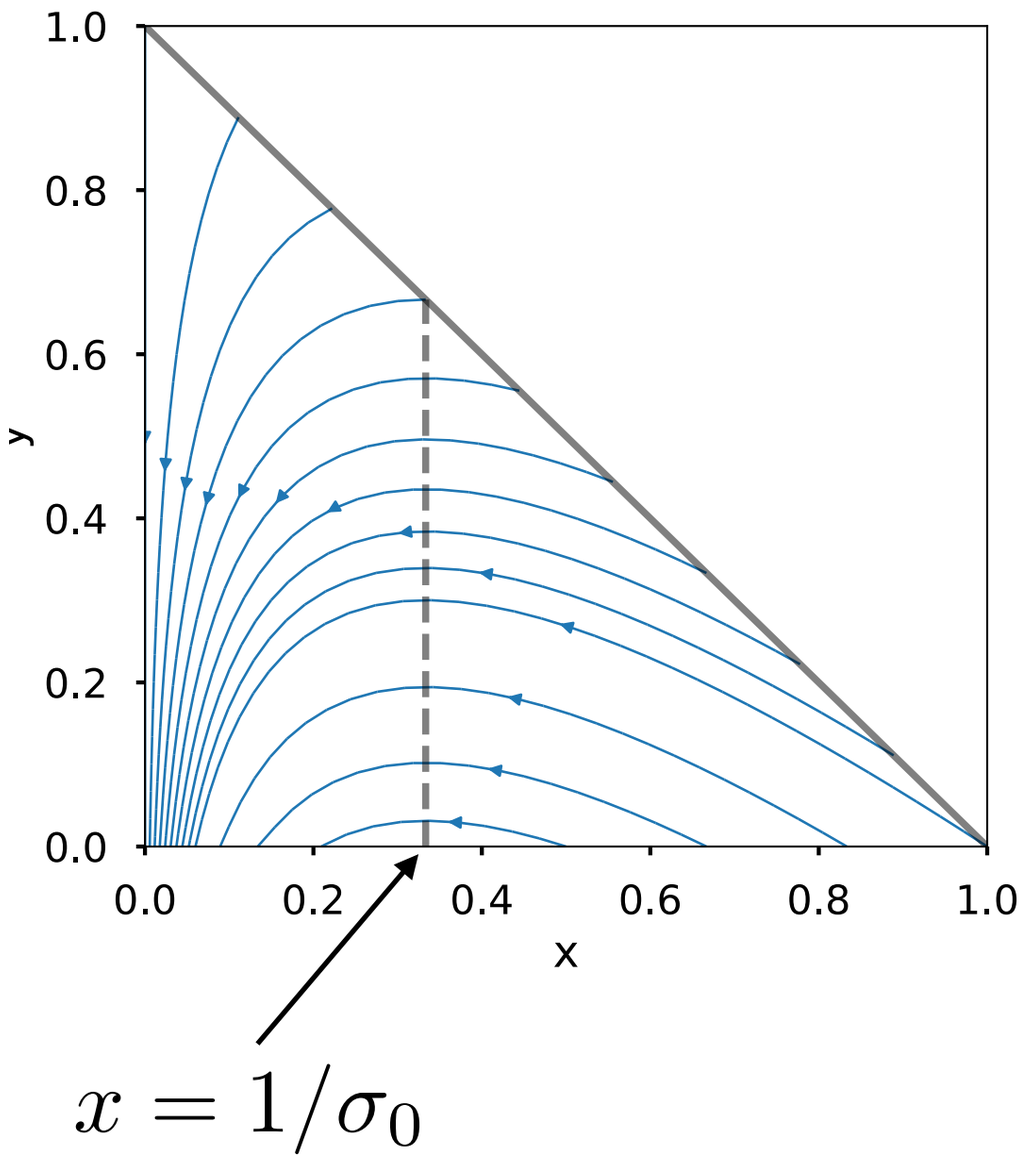

$$y'(t) = (\beta x - \gamma)\, y$$

$$\beta x = \gamma \iff x = 1/\sigma_0$$

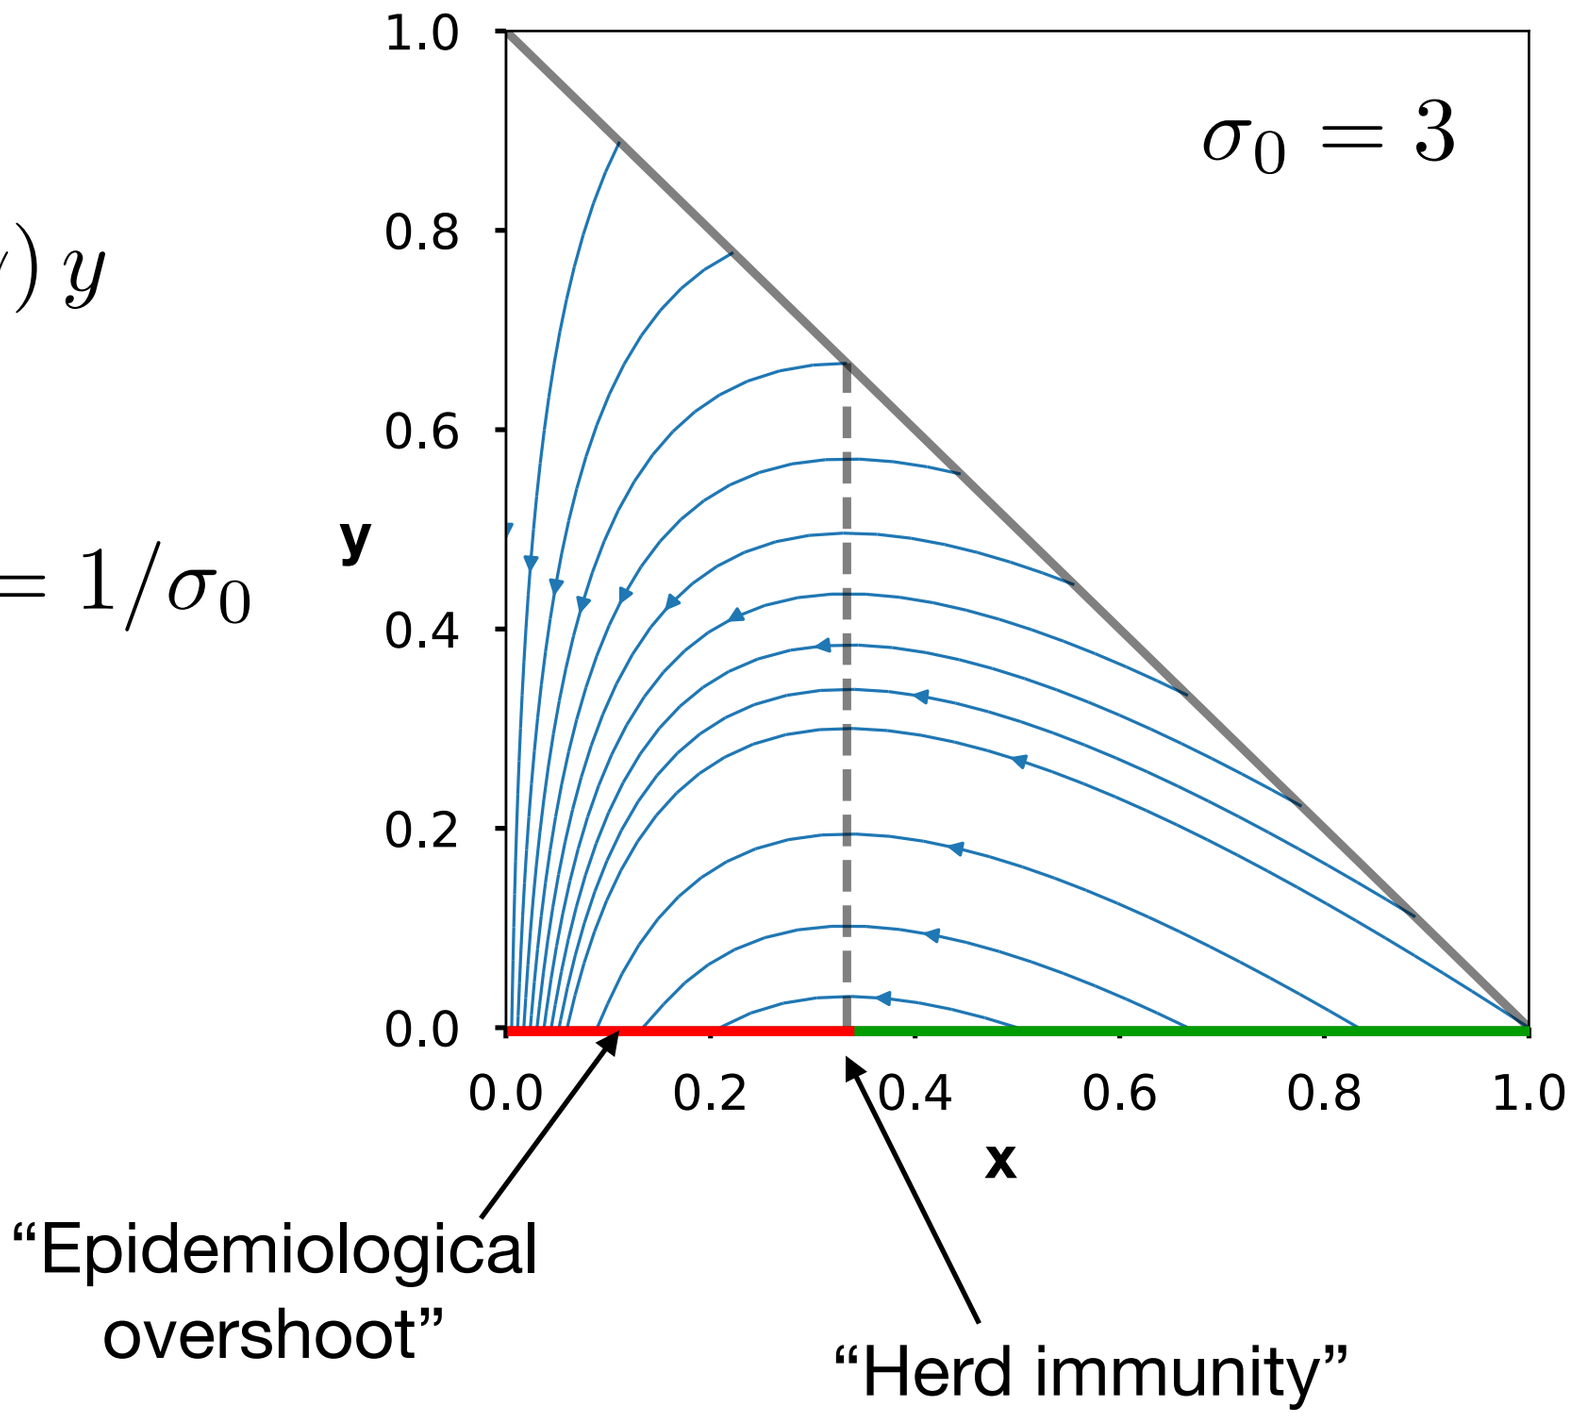

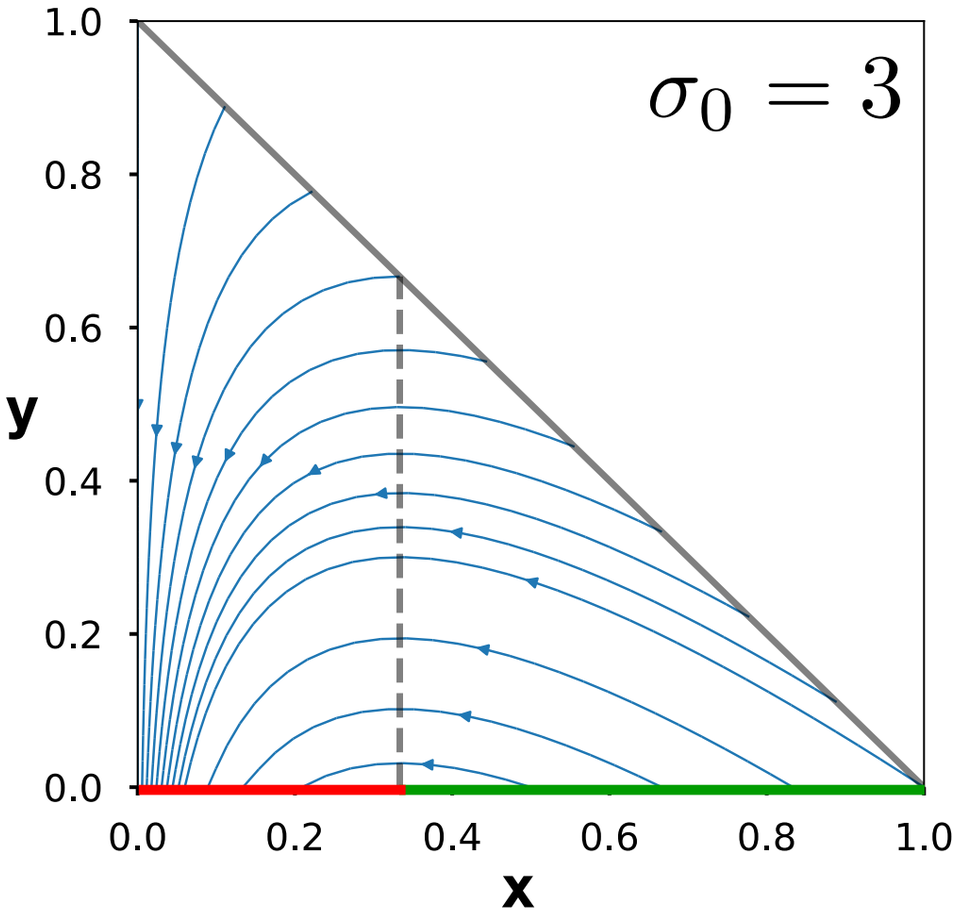
$\sigma_0 = 3$
y
x
0.0
0.2
0.4
0.6
0.8
1.0

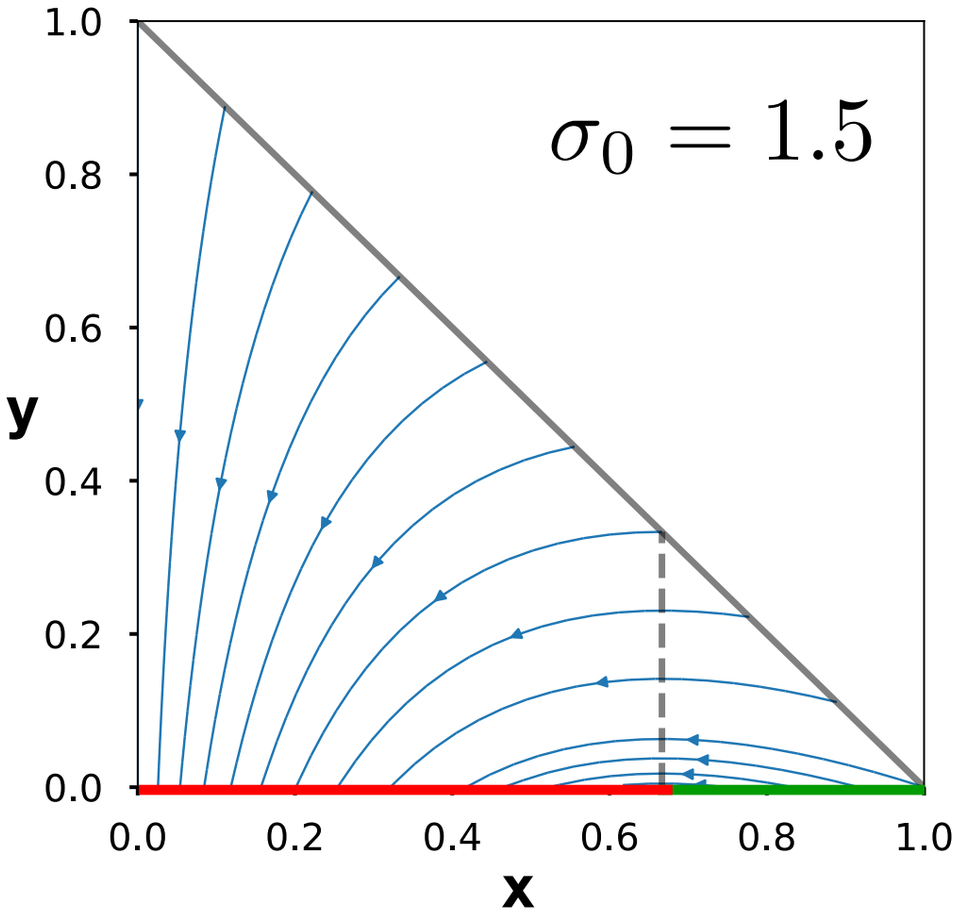
$\sigma_0 = 1.5$
y
x
0.0
0.2
0.4
0.6
0.8
1.0

Natural course of an epidemic

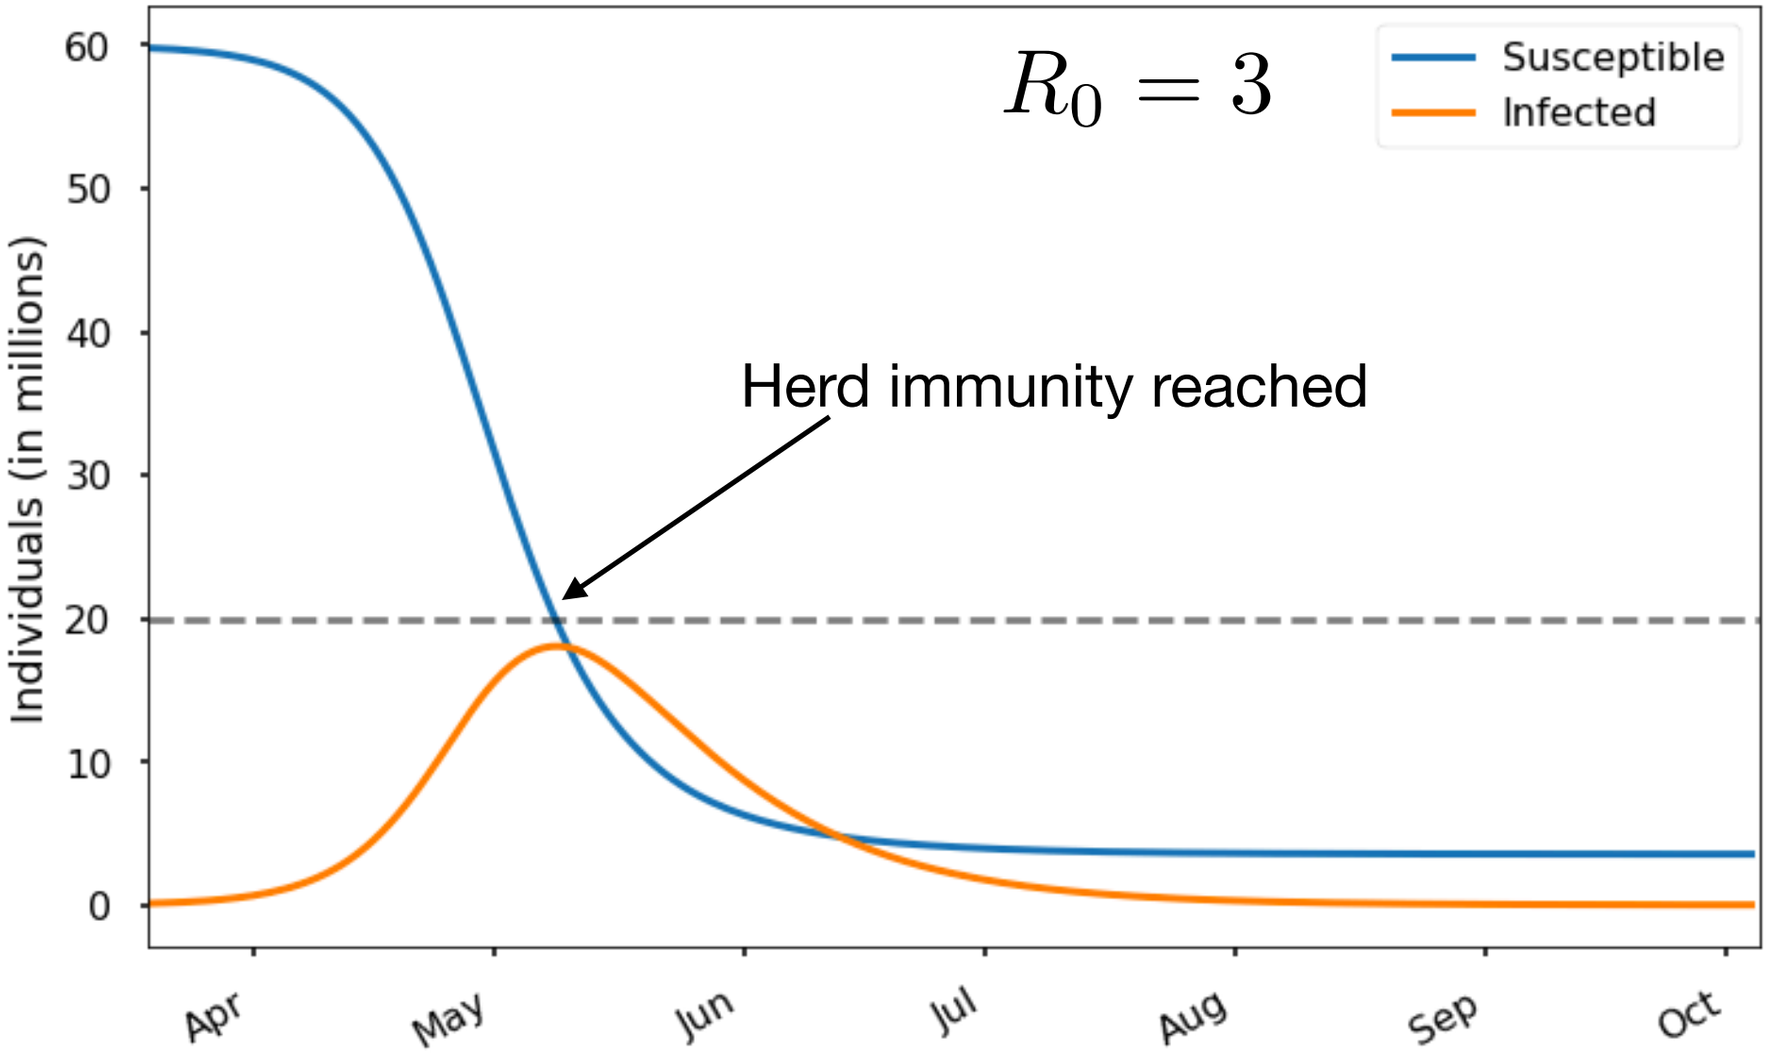

The SIR model with contact control (SIRq)

SIR:

$$x'(t) = -\gamma\sigma yx$$
$$y'(t) = \gamma\sigma yx - \gamma y$$

With social distancing (SIRq):

$$x'(t) = -\gamma\sigma(t) yx$$
$$y'(t) = \gamma\sigma(t) yx - \gamma y$$
$$0 \leq \sigma(t) \leq \sigma_0$$

$\sigma(t)$: Time-dependent contact rate

Assumptions

- Epidemic proceeds according to SIRq model
- No vaccination
- Intervention can be imposed only up to a **finite time T**, after which $\sigma(t) = \sigma_0$

Goal: minimize eventual total infections, or equivalently maximize

$$x_\infty = \lim_{t\to\infty} x(t) = 1 - \lim_{t\to\infty} z(t)$$

Objective function:

$$J = -\lim_{t\to\infty} x(t) + \int_0^T L(\sigma(t))dt$$

Penalty for infections

Penalty for control

Effect of constant contact reduction

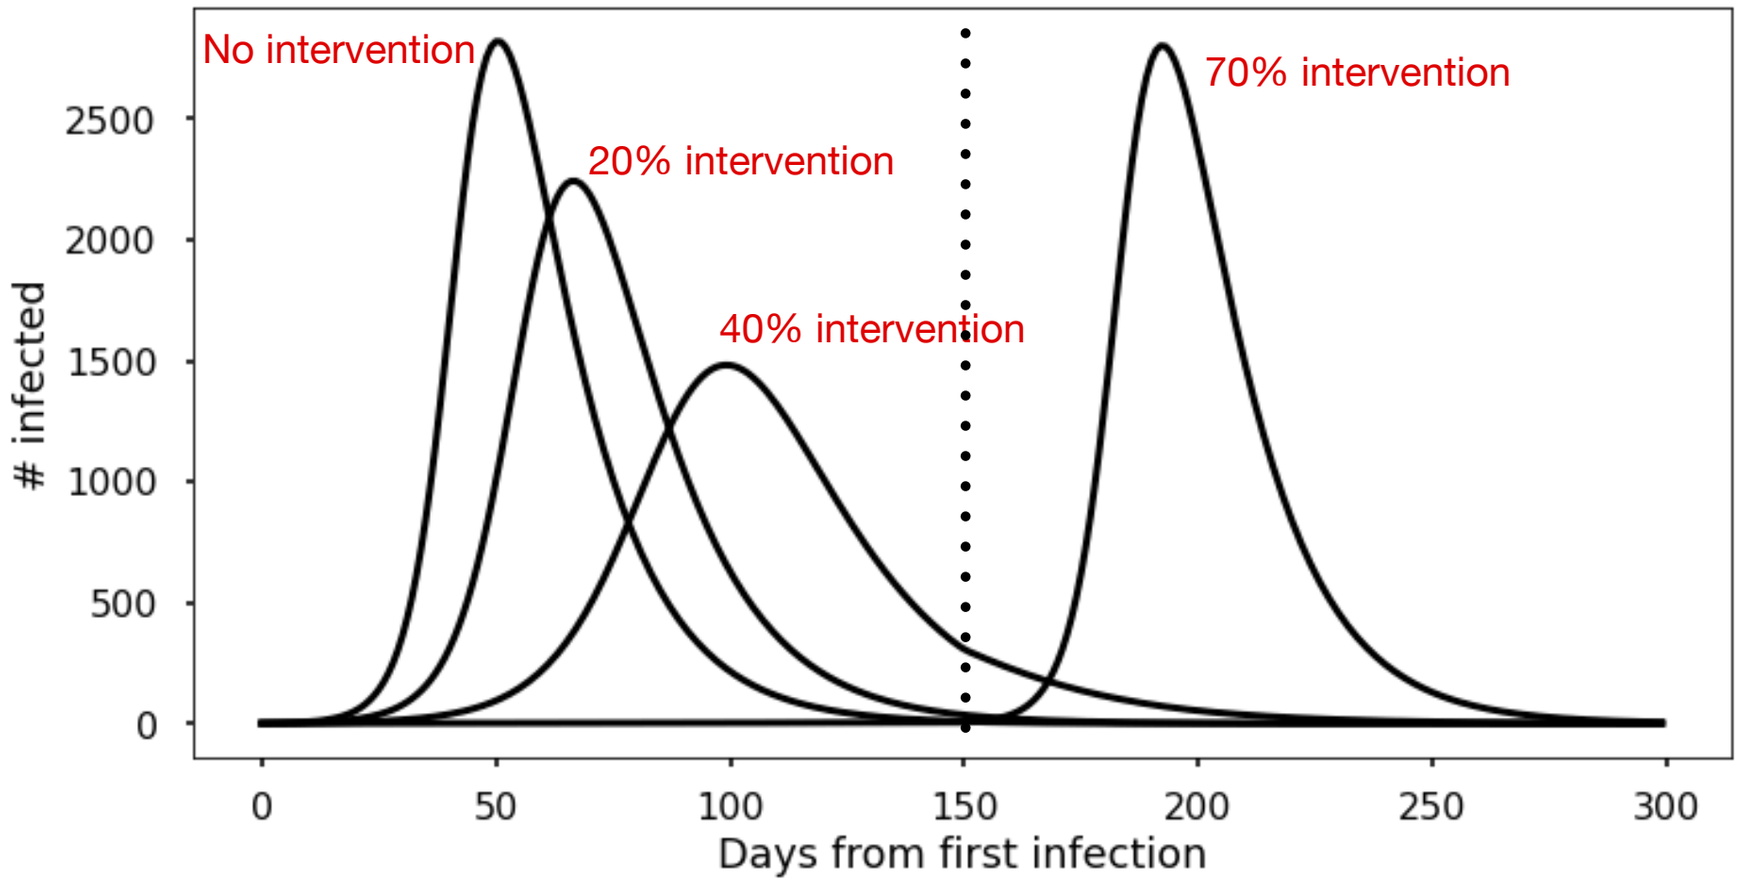

T = 150

For each scenario, σ(t) is held at a constant value up to t=T.

Total # infected is about the same in all cases.

Given $(x_0, y_0) \in \mathcal{D}, \sigma_0 > 0, T > 0,$

choose an admissible control $\sigma(t) : [0, T] \to [0, \sigma_0]$ to minimize

$$J = -\lim_{t \to \infty} x(t) + \int_0^T L(\sigma(t))dt$$

subject to SIRq:

$$x'(t) = -\gamma\sigma(t)yx$$
$$y'(t) = \gamma\sigma(t)yx - \gamma y$$

We know $x_\infty \leq 1/\sigma_0$

The neutrally stable equilibrium can only be reached after an infinite time interval

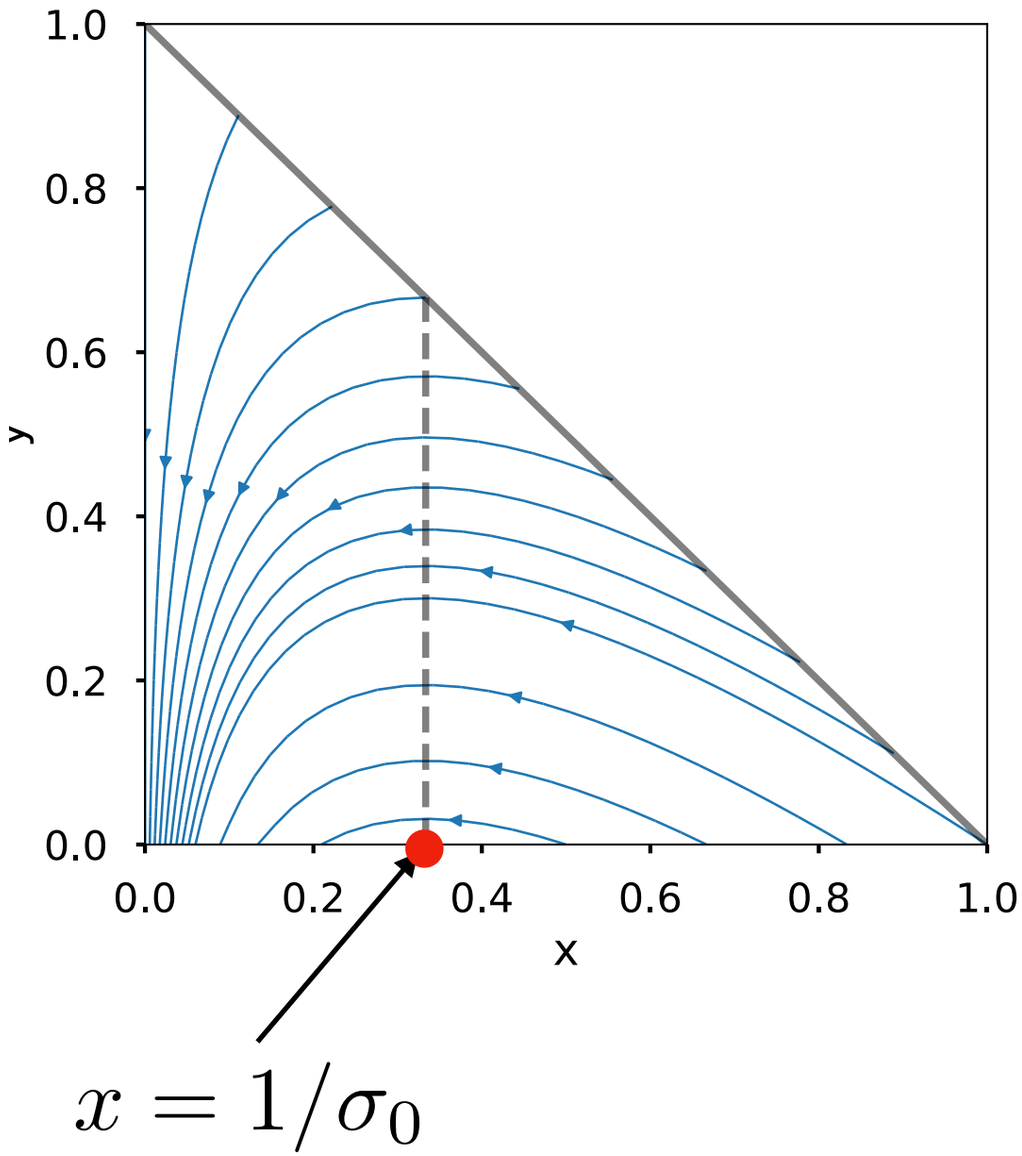

Controls that reach the neutrally-stable equilibrium

- Infinitely many such controls exist

- Green line: no control until herd immunity is reached, then maximum control

- Orange line: constant control

$$W_0(-\mu(x_0, y_0, \sigma_*)\sigma_*) = -\frac{\sigma_*}{\sigma_0}$$

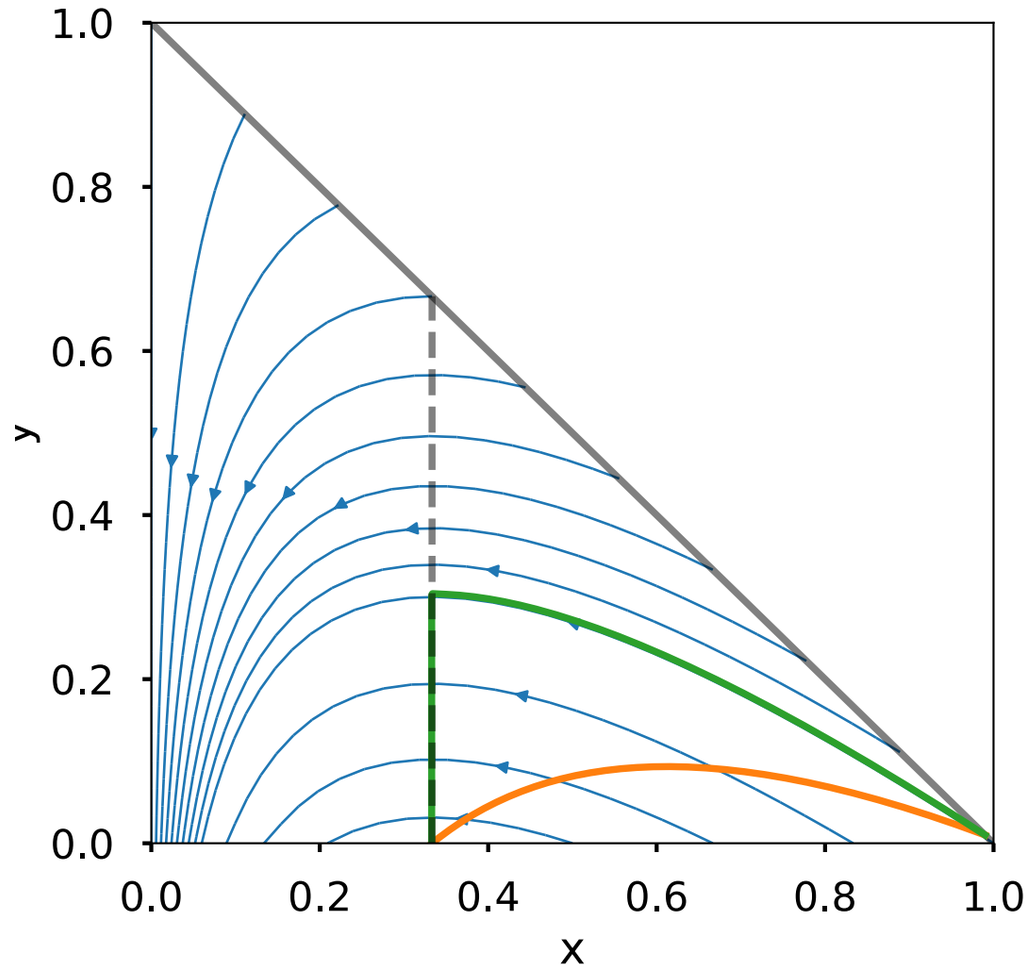

Zero-cost finite-time control

What is the optimal control if there is no cost for reducing contact?

$$\begin{aligned}
&\text{Given } (x_0, y_0) \in \mathcal{D}, \sigma_0 > 0, T > 0, \\
&\quad \text{choose an admissible control } \sigma(t) : [0, T] \to [0, \sigma_0] \\
&\text{to minimize} \quad J = -\lim_{t \to \infty} x(t)
\end{aligned}$$

Major difficulty with this problem: objective depends on state as $t \to \infty$

Finite-time reformulation

- Uncontrolled trajectories are level sets of

$$\mu(x, y, \sigma_0) := x(t)e^{-\sigma_0(x(t)+y(t))}$$

- Final state can be computed from state at time T:

$$x_\infty(x, y, \sigma_0) = -\frac{1}{\sigma_0} W_0(-\mu(x, y, \sigma_0)\sigma_0).$$

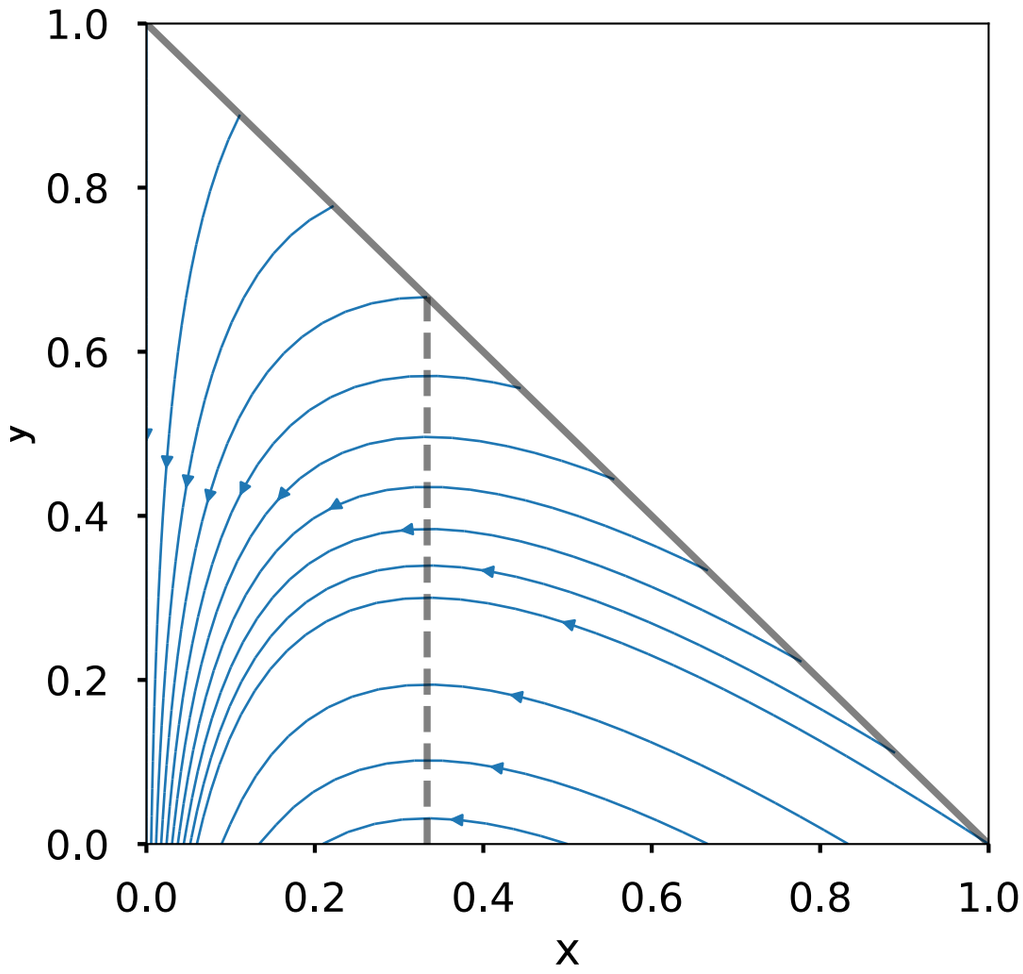

Finite-time reformulation

Given $(x_0, y_0) \in \mathcal{D}, \sigma_0 > 0, T > 0,$

choose an admissible control $\sigma(t) : [0, T] \to [0, \sigma_0]$ to minimize

$$J = -x_\infty(x(T), y(T), \sigma_0)$$

subject to SIRq.

Impact of control

$$\frac{\partial x_\infty}{\partial t} = \frac{\gamma y x_\infty}{1 - \sigma_0 x_\infty}(\sigma_0 - \sigma(t))$$

Always positive if $\sigma < \sigma_0$

$$x_\infty(x_0, y_0, \sigma_0) \leq x_\infty \leq 1/\sigma_0$$

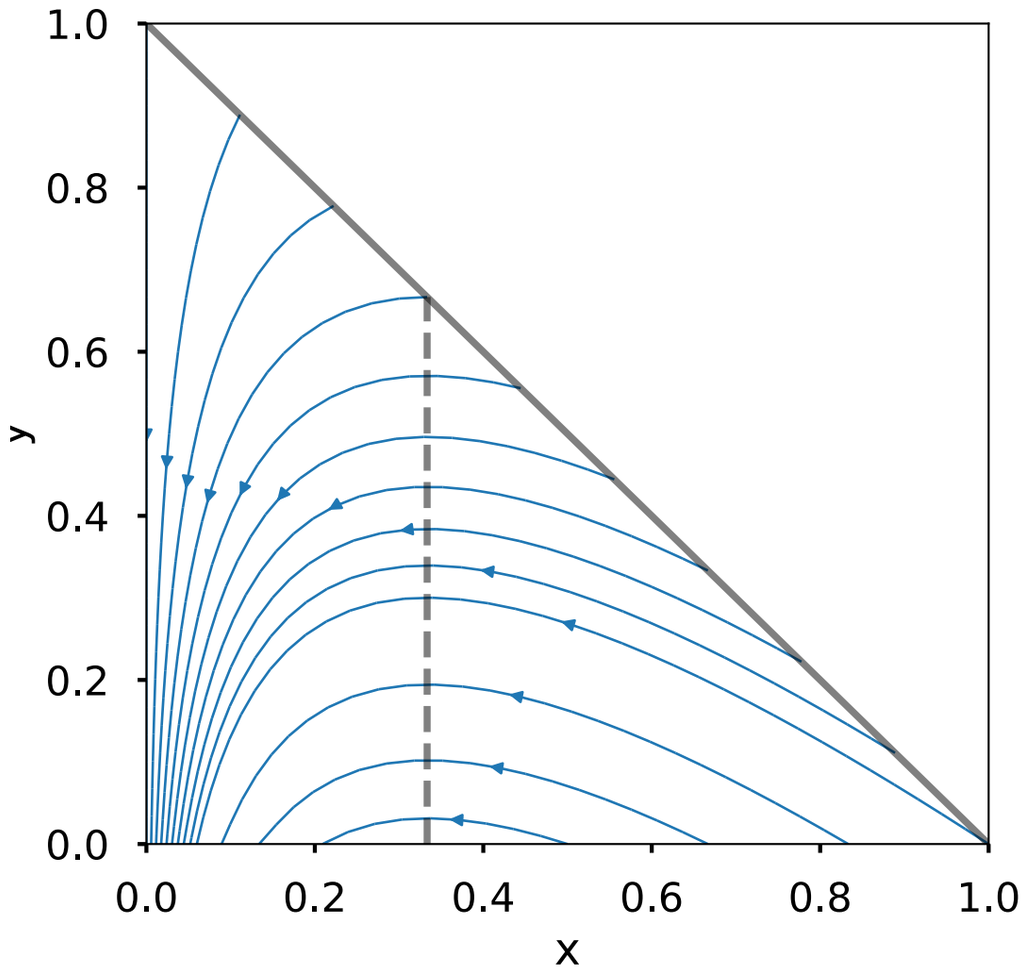

Pontryagin's Maximum Principle

- The PMP states that the optimal solution maximizes the Hamiltonian:

$$H(x(t), y(t), \sigma(t), \lambda_{1,2}(t), t) = -\lambda_1(t)\gamma\sigma(t)y(t)x(t) + \lambda_2(t)\gamma y(t)(\sigma(t)x(t) - 1) + L(x(t), y(t), \sigma(t))$$

The new variables satisfy a system of ODEs with conditions at the final time T. Thus the PMP leads to a two-point boundary value problem.

The optimal control is given implicitly by

$$\left.\frac{\partial L}{\partial \sigma}\right|_{\sigma(t)=\hat{\sigma}(t)} = -(\lambda_2(t) - \lambda_1(t))\gamma y x$$

The optimal control is given implicitly by

$$\left.\frac{\partial L}{\partial \sigma}\right|_{\sigma(t)=\hat{\sigma}(t)} = -(\lambda_2(t) - \lambda_1(t))\gamma y x$$

But also subject to the bounds:

$$\sigma^*(t) = \max\left(0, \min\left(\sigma_0, \hat{\sigma}(t)\right)\right)$$

Because the RHS is independent of σ,
the optimal control will always be the upper or lower bound.

This is known as a **bang-bang control**.

Bang-bang SIR controls

$$\sigma = 0 \implies \begin{aligned} x'(t) &= 0 \\ y'(t) &= -\gamma y \end{aligned}$$

- Alternate between following flow lines (control off) and moving straight downward (control on)
- The best control is the one that reaches a given point in the shortest time
- It can be shown that the green path is fastest
- This further implies that there is only one switching time

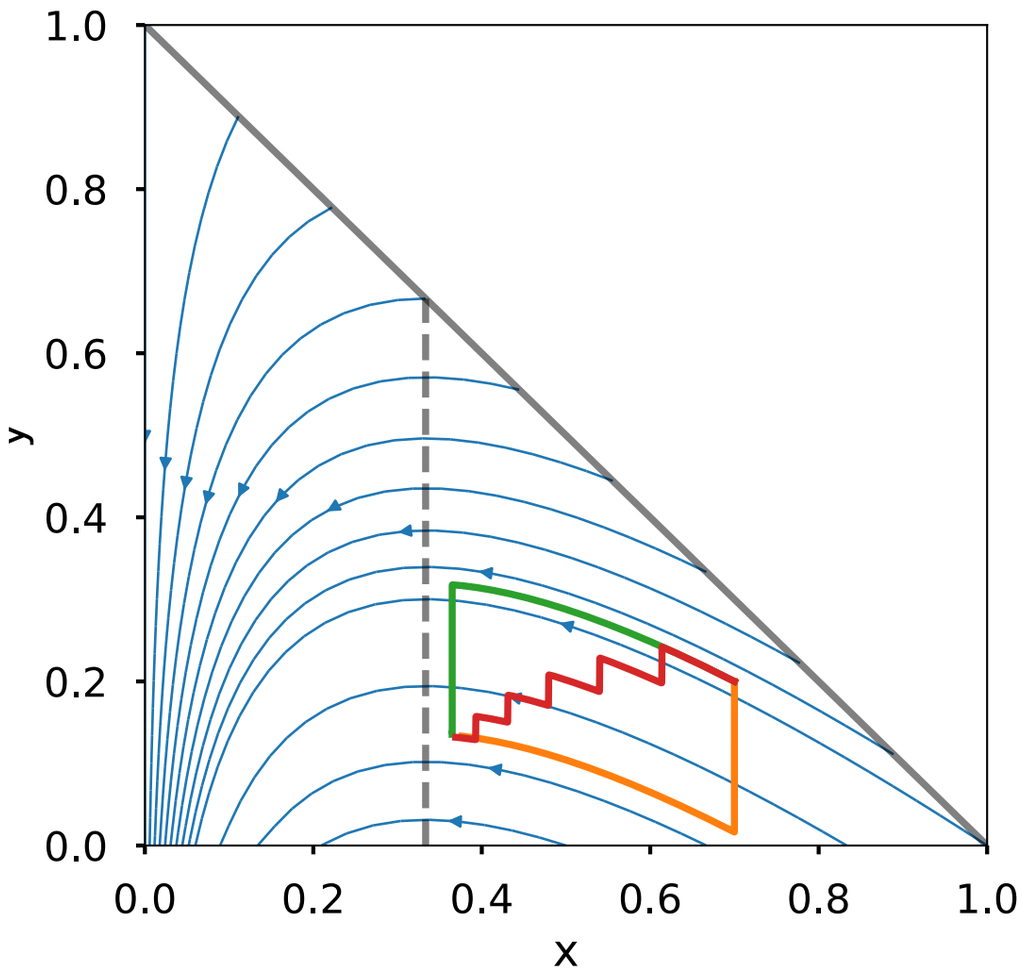

Optimal zero-cost control

Theorem 3. *The optimal control for* (15) *is unique and is given by*

$$\sigma(t) = \begin{cases} \sigma_0 & t < t^* \\ 0 & t^* \le t \le T, \end{cases}$$

where

$$t^* = 0 \text{ if } x_0 \le \frac{1}{\sigma_0(1 - e^{-\gamma T})},$$

and otherwise t^* *is the unique solution of*

$$x(t^*; \sigma_0, x_0, y_0) = \frac{1}{\sigma_0(1 - e^{-\gamma(T-t^*)})}$$

The proof is based on analysis specific to SIRq, since it doesn't satisfy the conditions usually required for uniqueness.

Optimal zero-cost control

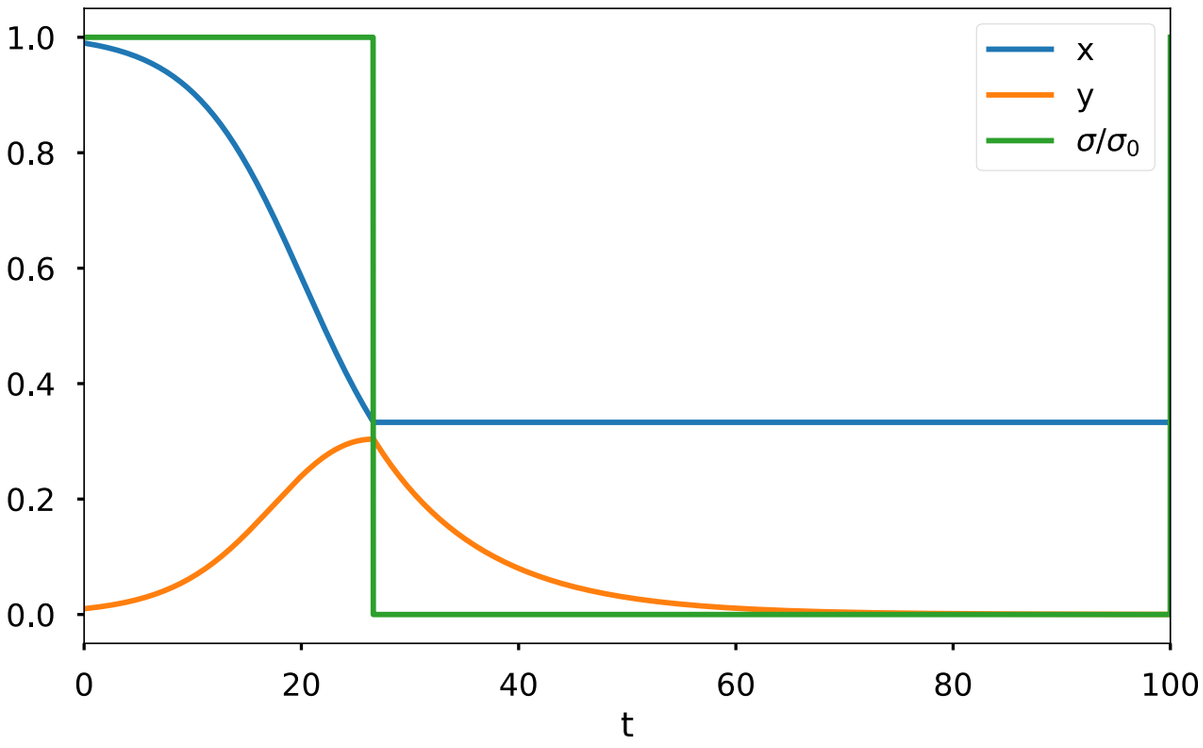

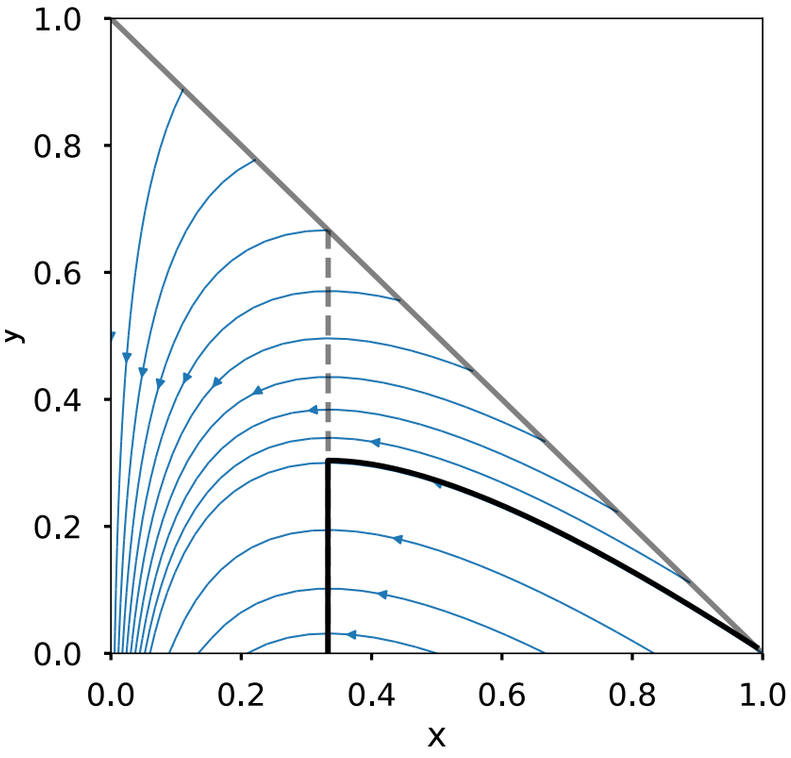

If you have enough time: wait for herd immunity, then impose full lockdown.

Optimal zero-cost control

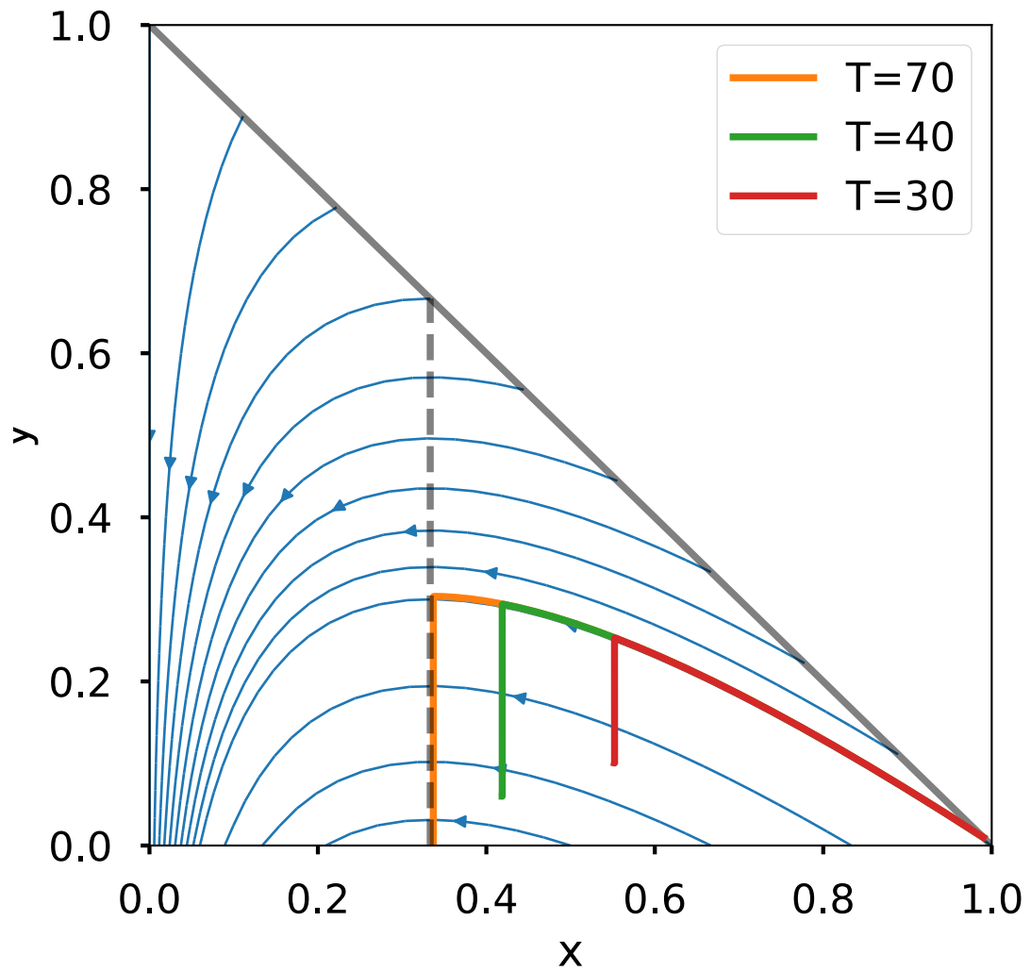

If you don’t have enough time: impose full lockdown starting at the optimal moment.

What if we can’t impose full lockdown?

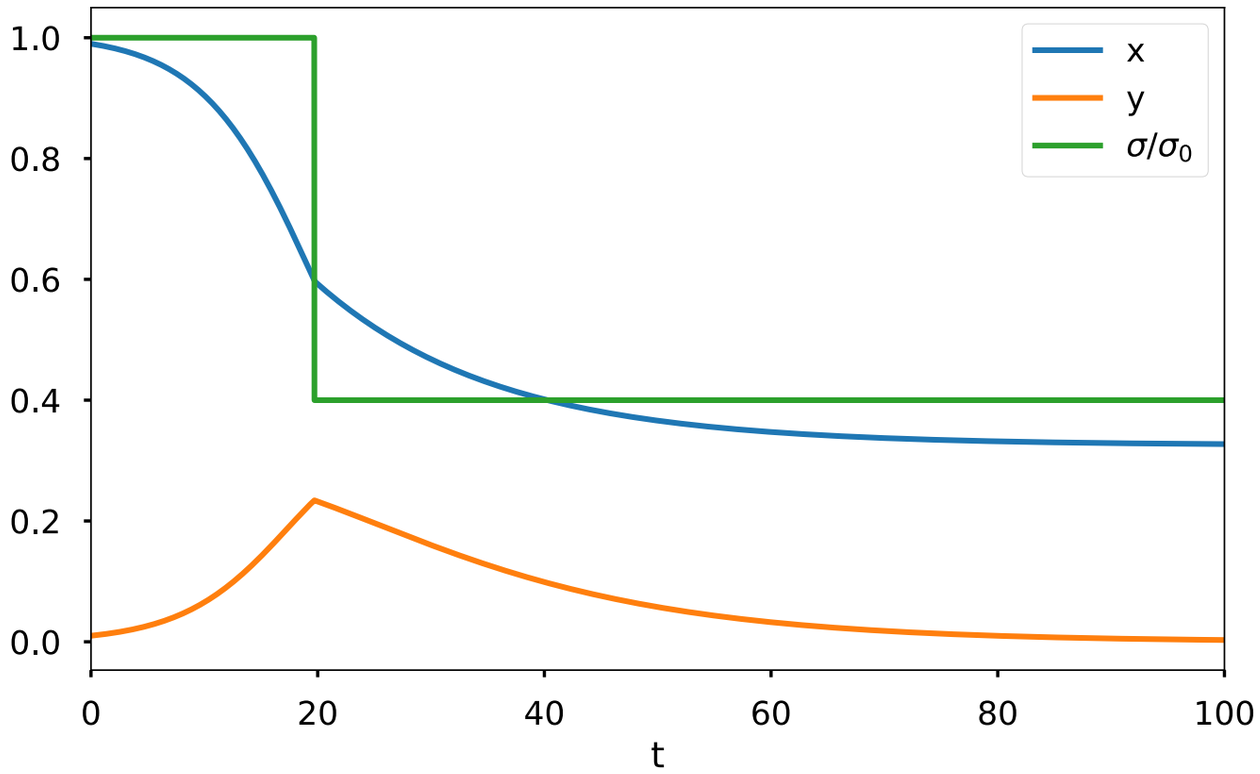

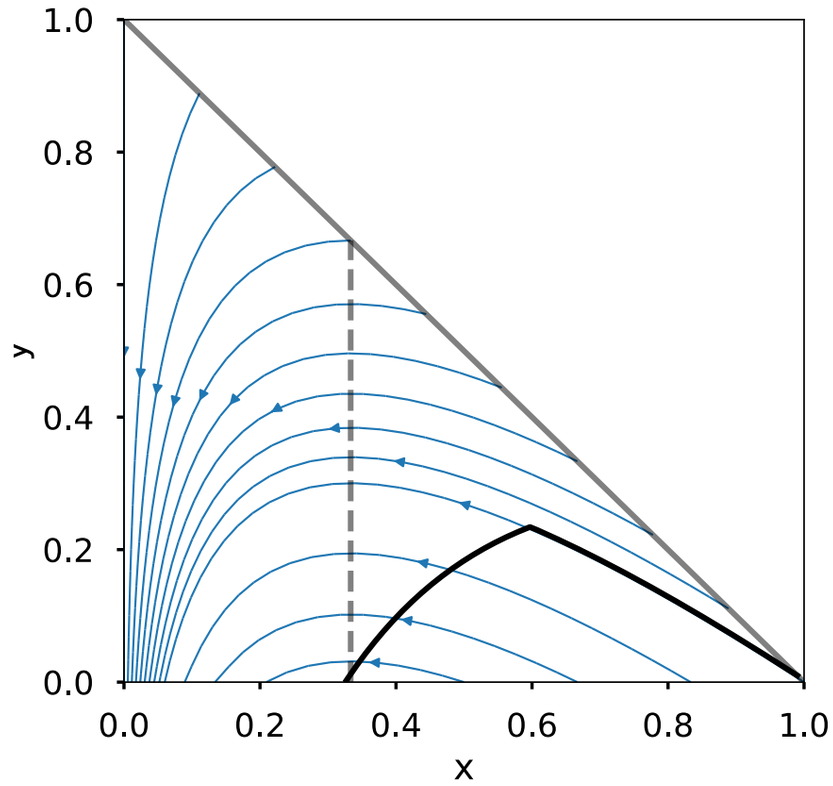

Hamilton-Jacobi-Bellman formulation

$$u_t = \gamma y u_y - \gamma x y \min_{\sigma} \left((u_y - u_x)\sigma \right)$$

- PDE with **final** condition

$$u(x, y, T) = -x_\infty(x, y, \sigma_0)$$

- x, y become **independent** variables

- u(x,y,t) is the optimal cost function

$$\sigma(t) = \begin{cases} 0 & u_y(x, y, t) > u_x(x, y, t) \\ \sigma_0 & u_y(x, y, t) < u_x(x, y, t) \end{cases}$$

Hamilton-Jacobi-Bellman formulation

- 2D advection equation $u_t = \gamma y u_y - \gamma x y \min_{\sigma} \left((u_y - u_x)\sigma \right)$
- At t=T, $u_y > u_x$, so $\sigma(T) = 0$.

$$u(x, y, T) = -x_\infty(x, y, \sigma_0)$$

- Values are advected along vertical lines ($\sigma(T) = 0$) or SIR trajectories ($\sigma(T) = 0$).

$$\sigma(t) = \begin{cases} 0 & u_y(x, y, t) > u_x(x, y, t) \\ \sigma_0 & u_y(x, y, t) < u_x(x, y, t) \end{cases}$$

Optimal control with non-zero cost

$$L = c_2 \left(1 - \frac{\sigma(t)}{\sigma_0}\right)^2$$

- The (weak) PMP gives necessary but not sufficient conditions for optimality
- In the zero-cost case, we could use geometric arguments to show sufficiency
- In general, we will need a more powerful tool. Our options:
  - Stronger versions of the PMP. Required conditions are **not satisfied** by SIRq.
  - Solve the Hamilton-Jacobi-Bellman PDE.

Hamilton-Jacobi-Bellman

$$u_t - \gamma y u_y = - \min_{0 \le \sigma \le \sigma_0} \left((u_y - u_x) \gamma x y \sigma(t) + c_2 \left(1 - \frac{\sigma(t)}{\sigma_0} \right)^2 \right)$$

$$u(x, y, T) = -x_\infty(x, y, \sigma_0)$$

u is the value of the objective function. The minimum is achieved with

$$\sigma(t) = \sigma_0 \min \left(1, \max \left(0, \left(1 - \frac{\sigma_0 \gamma}{2 c_2} x y (u_y - u_x) \right) \right) \right)$$

We solve the PDE backward from the final time.

Inserting the PDE solution in the formula above yields the optimal control.

Then, solving the ODE with this control yields the dynamics.

Hamilton-Jacobi-Bellman

- The domain of this PDE is the same space where solutions of the ODE system live.

- Information should flow only backward along the streamlines, so we need to use an “upwind” numerical method.

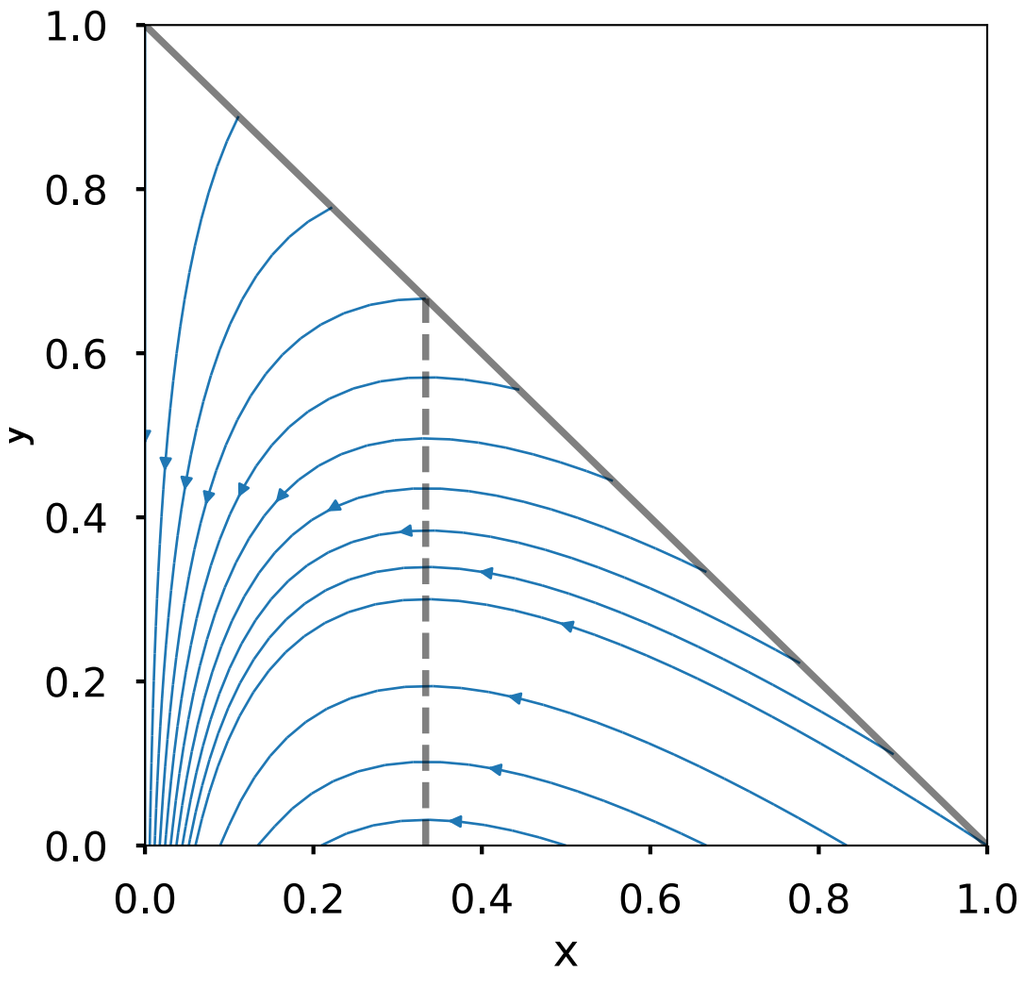

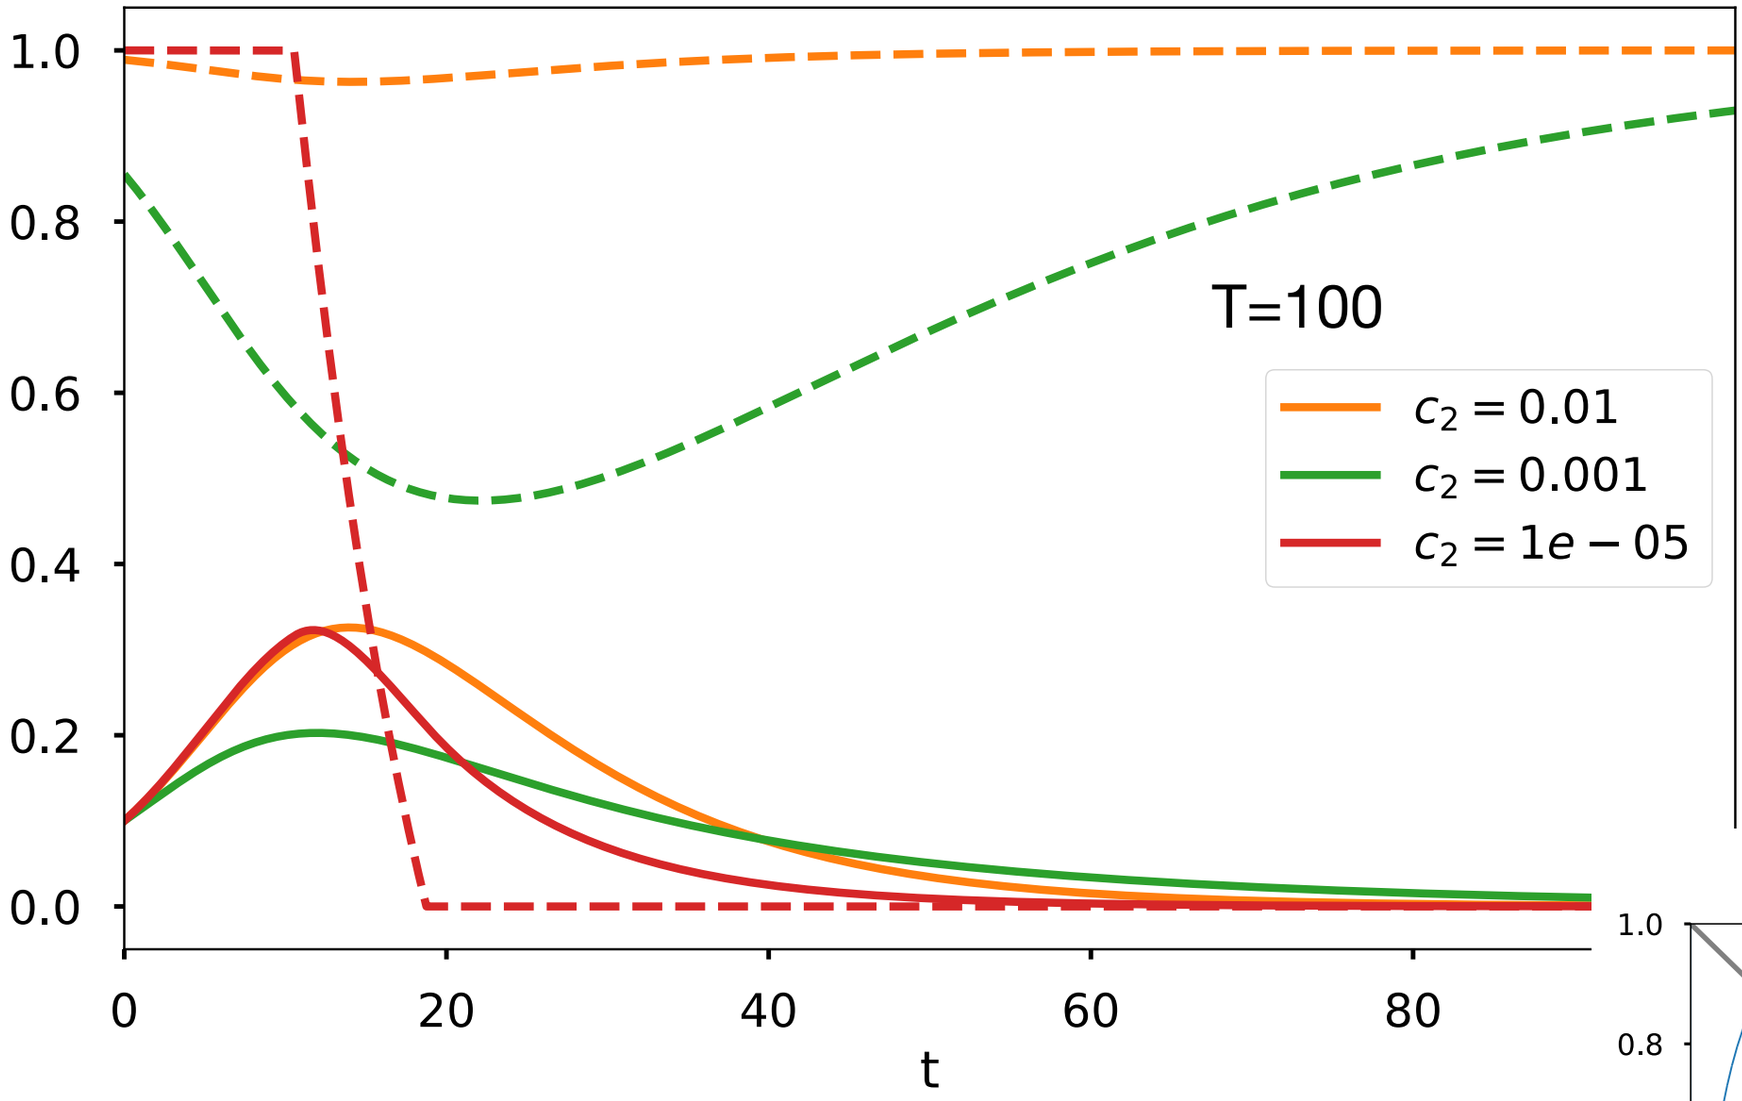

$$L = c_2 \left(1 - \frac{\sigma(t)}{\sigma_0}\right)^2$$

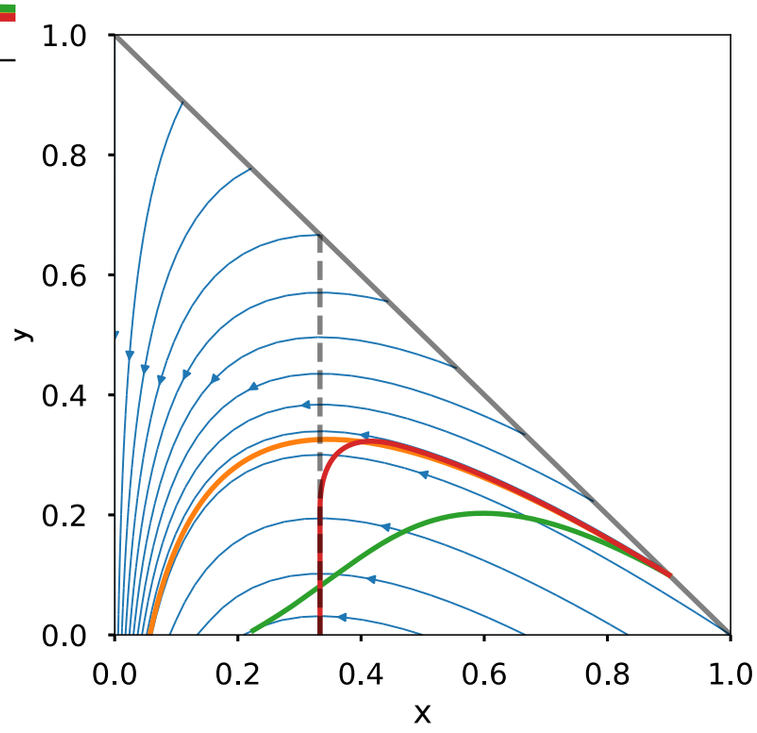

Solid lines: Infected fraction

Dashed lines: Control $\sigma(t)/\sigma_0$

Avoiding hospital overflow

$$L(x(t), y(t), \sigma(t)) = c_2 \left(1 - \frac{\sigma(t)}{\sigma_0}\right)^2 + c_3 g(y(t) - y_{\max}).$$

Economic cost of intervention

Penalty for hospital overflow

Hospital capacity

- A full model requires us to quantify:
  - The cost of intervention
  - The value of a life
  - The increased likelihood of death when hospital beds run out

$$g(v) = \frac{v}{1 + e^{-100v}}$$

Avoiding hospital overflow

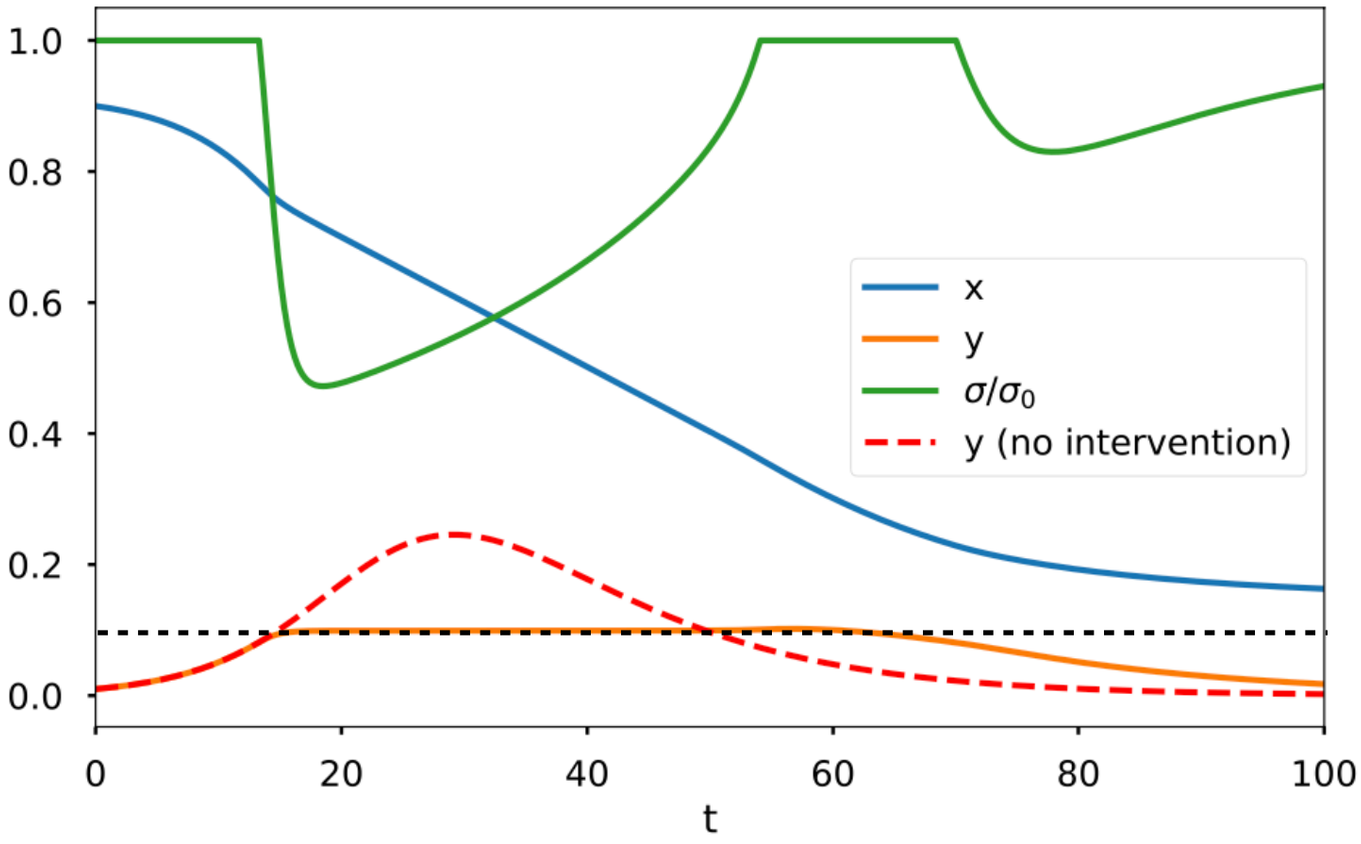

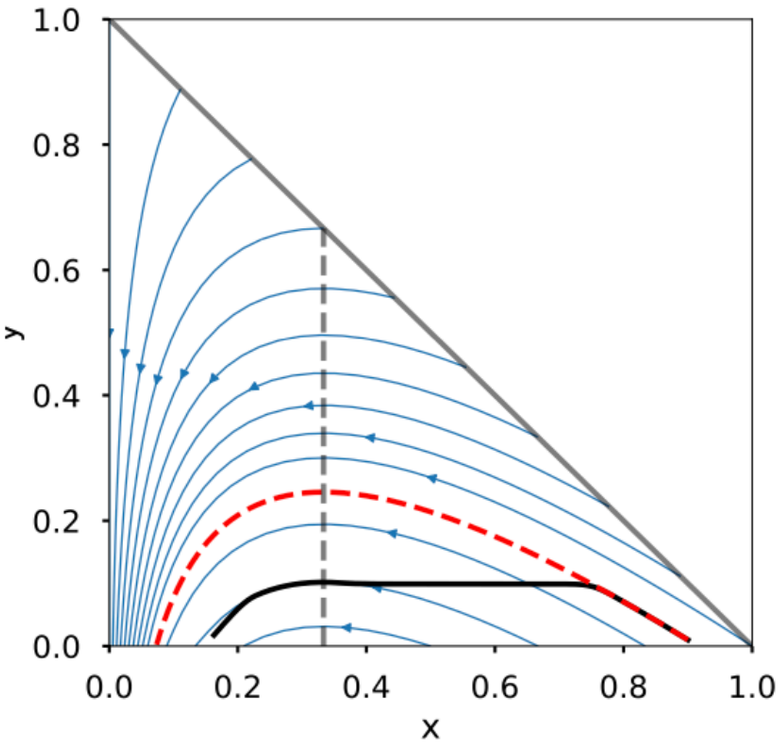

$$c_2 = 1/100$$

$$c_3 = 100$$

$$y_{\max} = 0.1$$

Avoiding hospital overflow

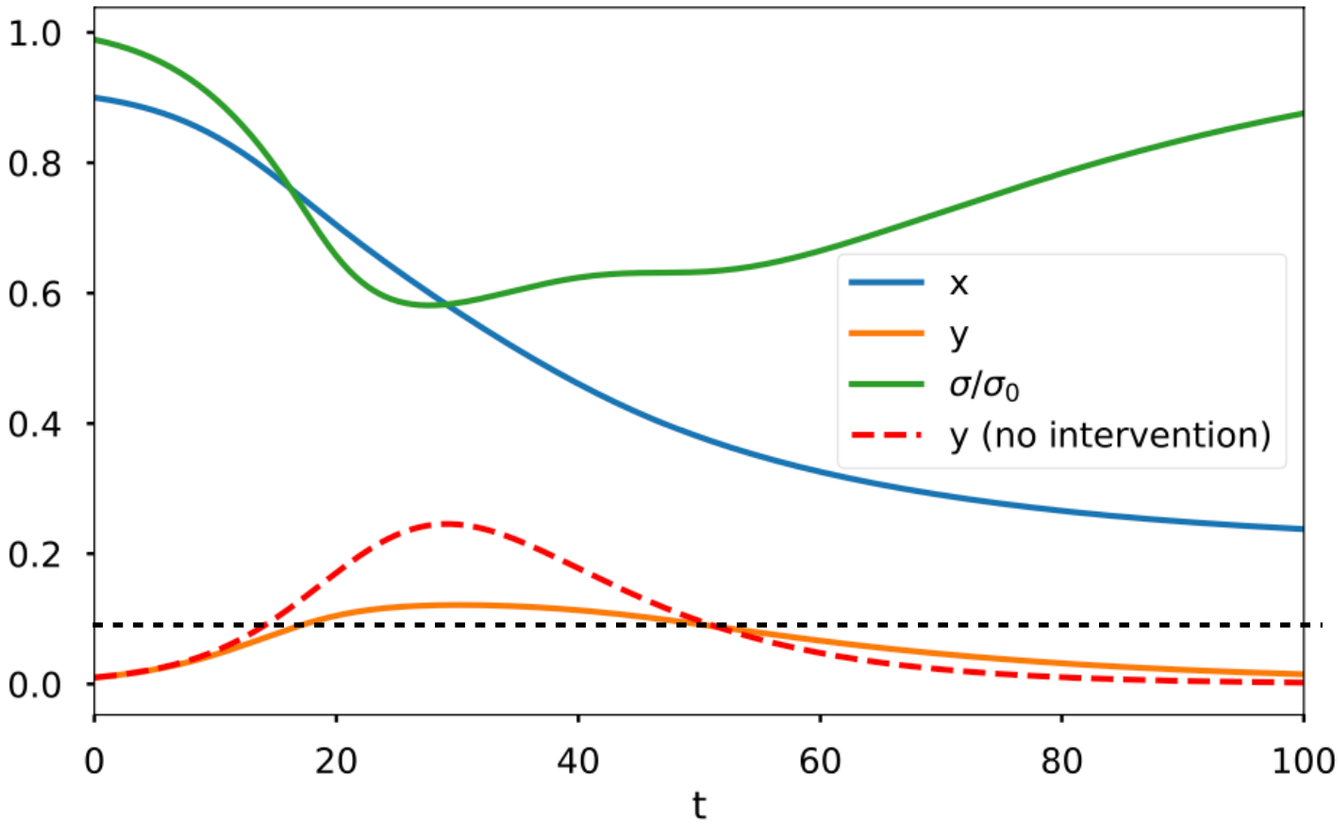

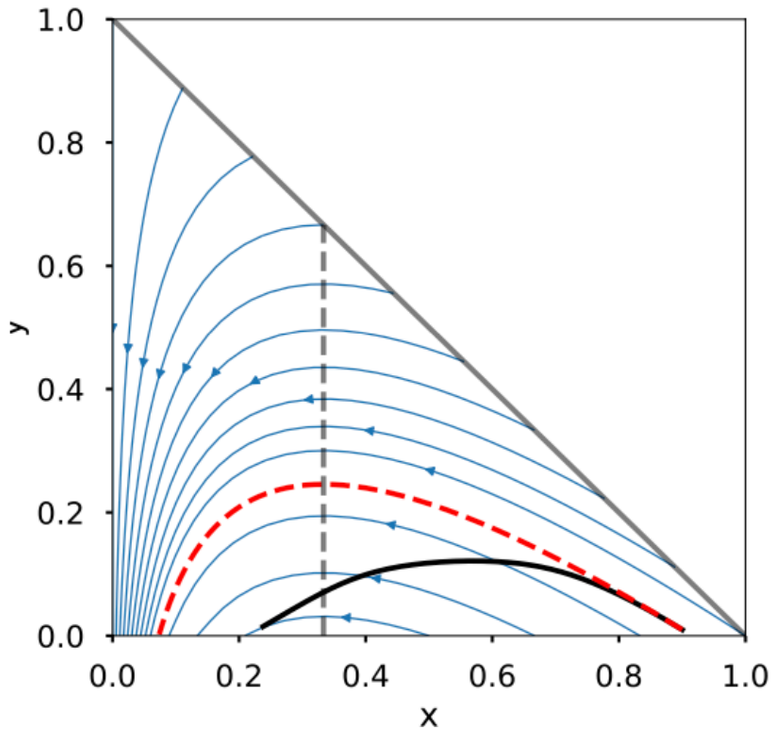

$$c_2 = 1/100$$

$$c_3 = 1$$

$$y_{\max} = 0.1$$

- Hospital capacity is exceeded
- Total infections significantly reduced

Covid-19 scenario A

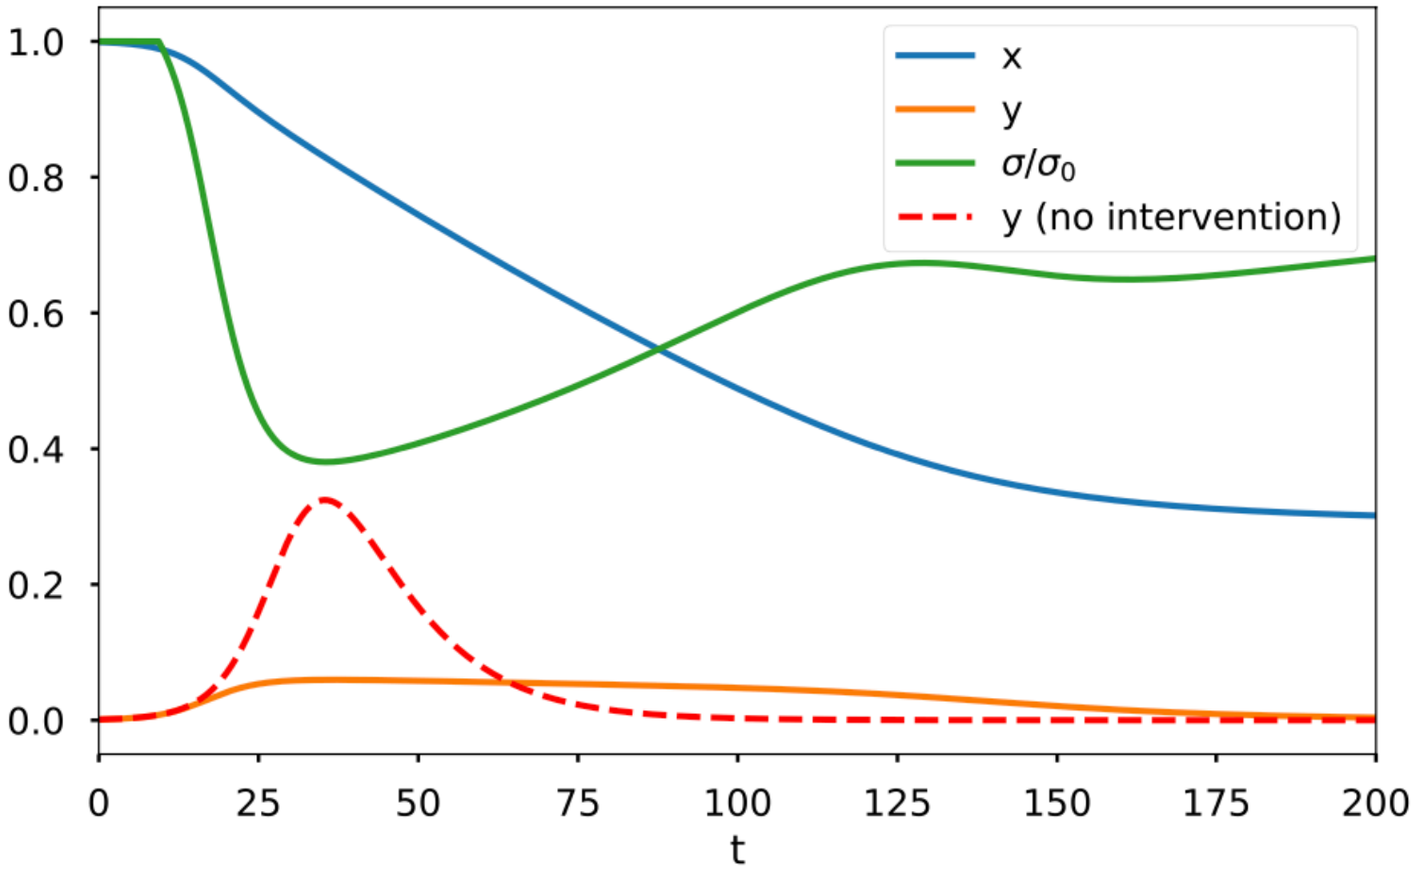

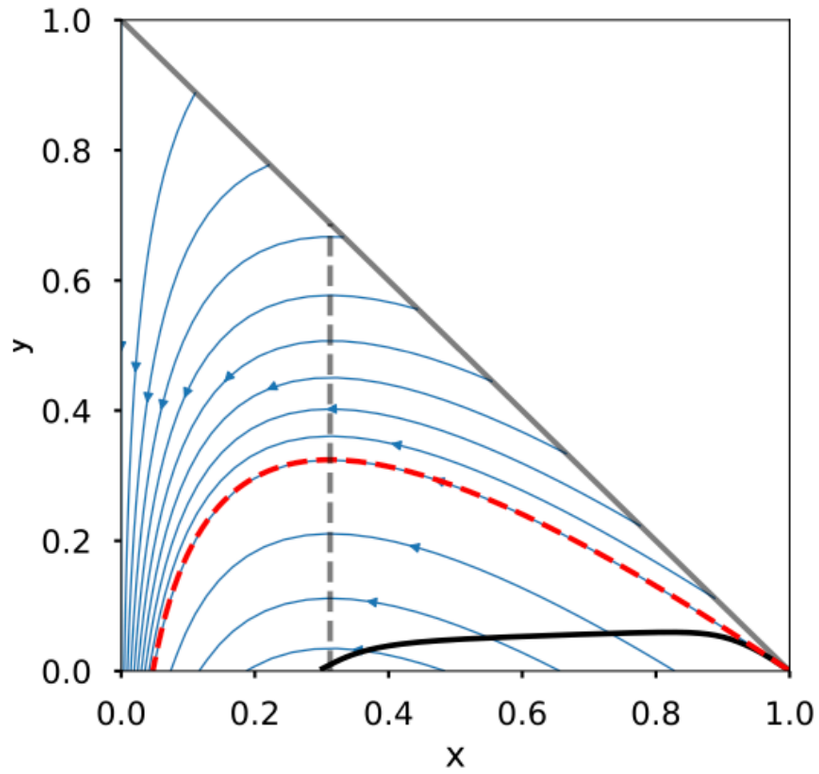

- Control focuses mainly on managing hospital occupancy
- Control lifted after 200 days
- Very little epidemiological overshoot

Covid-19 scenario B

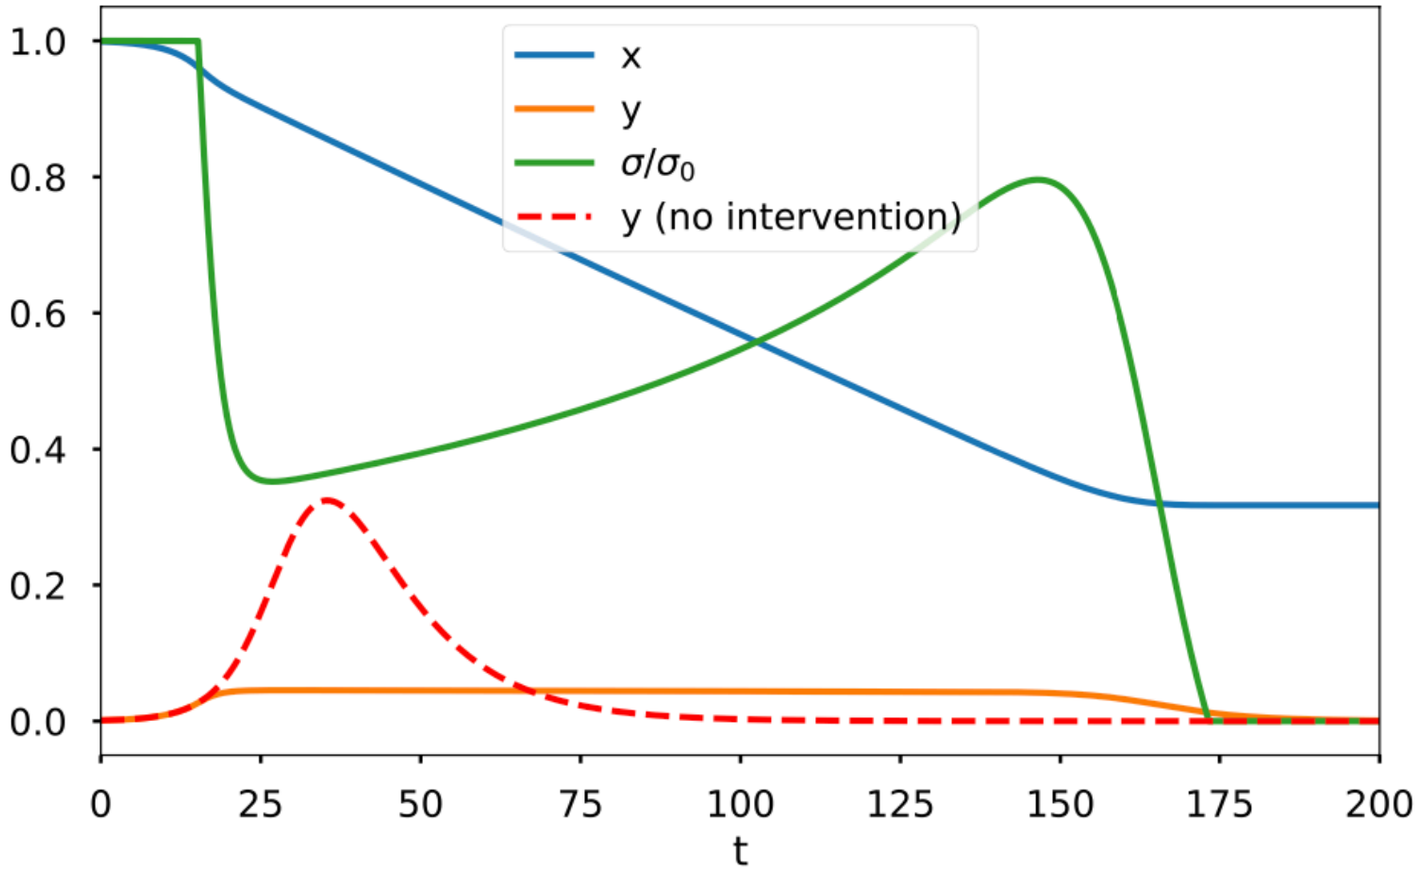

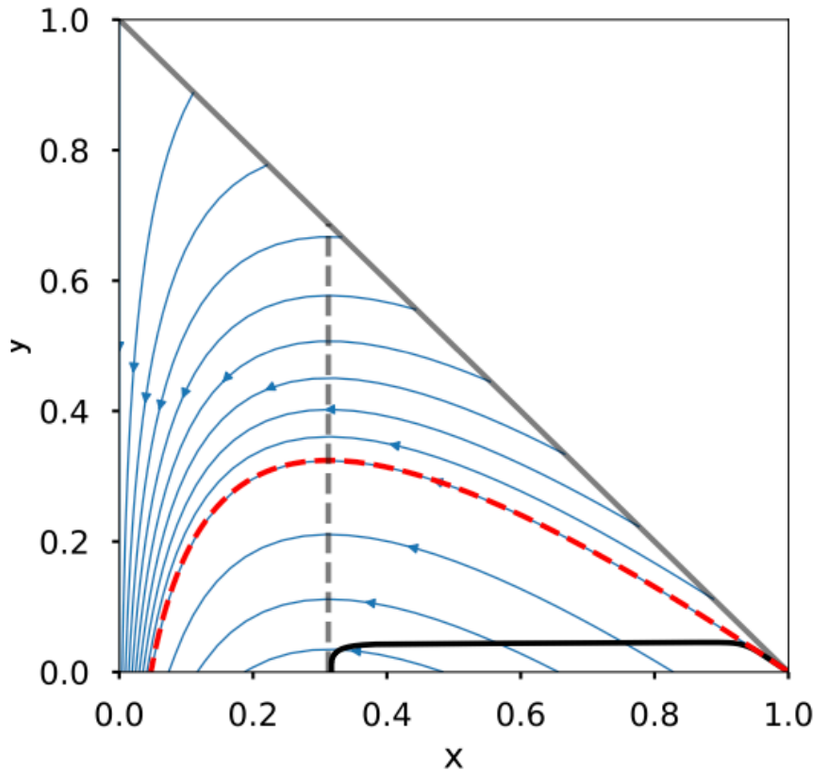

- Higher penalty for infections and lower cost of intervention
- Strong control at late time gives near-eradication
- Virtually no epidemiological overshoot

Covid-19 scenario C

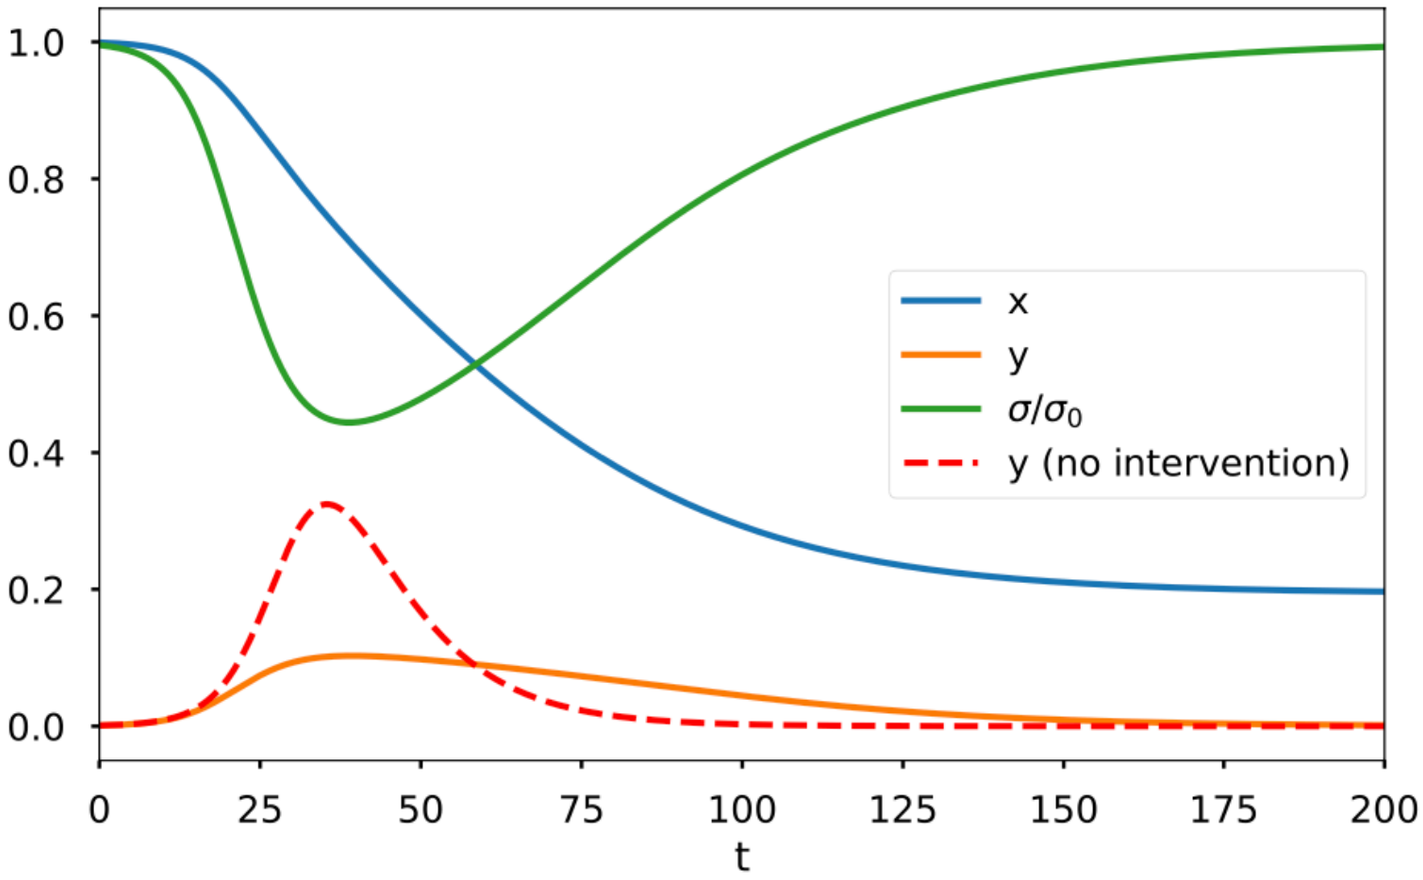

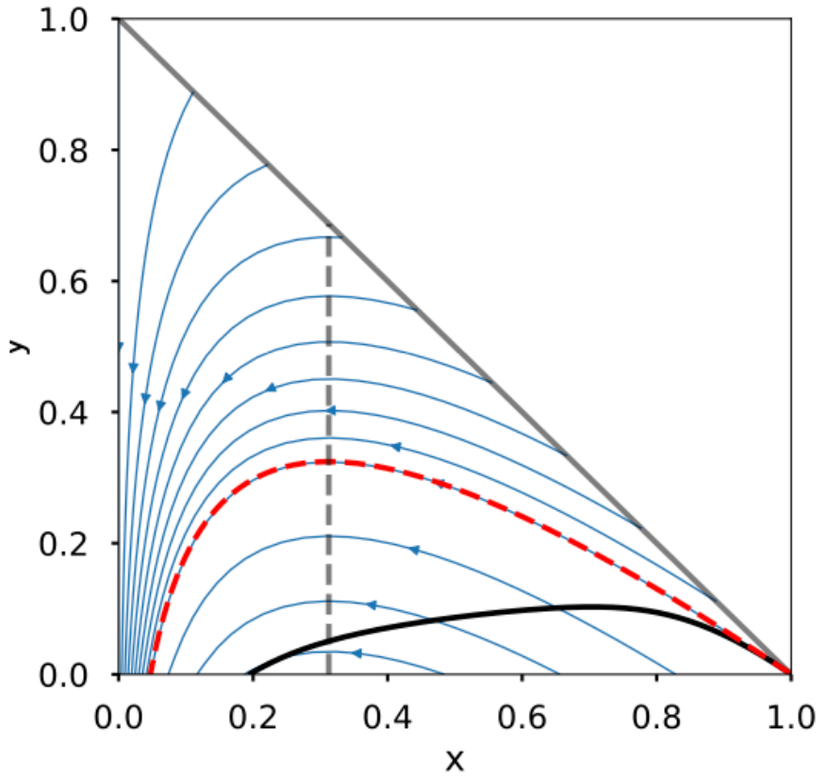

- Much higher cost of control
- Hospital capacity greatly exceeded
- Substantial epidemiological overshoot

Numerical solution

- BVP via Pontryagin
  - No-cost problem is extremely challenging; c_2 is a regularizing parameter
  - Use parameter continuation in c_2
- Hyperbolic PDE via HJB
  - Much more expensive in principle (PDE in 2D+time)
  - Discontinuous velocity field is difficult to capture accurately
  - Boundary conditions at “downwind” boundaries are problematic
  - Small errors in the solution can lead to substantial errors in the control

Conclusions and further directions

- All of this is based on an extremely simple model. Nevertheless, it may reveal useful conceptual or qualitative relations.

- Counterintuitively, waiting to intervene can be beneficial even if one ignores the cost of intervention.

- Extensions:
    - More complex compartmental models
    - Multiple demographic groups
    - Vaccines available in limited supply
    - etc.