Lecture 7: SIR Models

Radu Balan

Department of Mathematics, NWC
University of Maryland, College Park, MD

Version: February 16, 2024

Epidemiological Models

In this lecture we continue discussing epidemiological models. There are two main types of epidemic models:

- deterministic (or, compartimental) model
- stochastic (e.g., agent based) model

We focus on three deterministic models:

1. SI (Susceptible-Infected) Model
2. SIR (Susceptible-Infected-Removed) Model
3. SEIR (Susceptile-Exposed-Infected-Removed) Model

Today we discuss the SIR model.

The SIR Model with No vitals

Assume a system with three compartments: 'Susceptible' (S), 'Infected' (I) and 'Removed' or 'Recovered' (R). At time $t_0 = 0$ the system has a total of N individuals (initial total population). Most of them are susceptible $S(0)$, but some are infected, $I(0)$ and possibly some are in the recovered state, $R(0)$. Our intention is to model the time evoluation of these populations. We start with SI model:

$$\begin{cases} \frac{dS}{dt} &= -\beta S \frac{I}{N} \quad , \quad S(0) \\ \frac{dI}{dt} &= \beta S \frac{I}{N} \quad , \quad I(0) \end{cases}$$

where $\beta \geq 0$ is a parameter. We append a new term to model transition from $I \mapsto R$, assuming a constant rate of transition α:

$$\begin{cases} \frac{dS}{dt} &= -\beta S \frac{I}{N} \quad , \quad S(0) \\ \frac{dI}{dt} &= \beta S \frac{I}{N} - \alpha I \quad , \quad I(0) \\ \frac{dR}{dt} &= \alpha I \quad , \quad R(0) \end{cases}$$

The SIR Model

Deterministic simulations

Simulation of the SIR model:

$$\beta = 2\ ,\ \alpha = 1\ ,\ S(0) = 2000\ ,\ I(0) = 23\ ,\ R(0) = 0$$

Results were obtained with an Euler scheme with step size $h = 0.01$.

Note: The infected population $I(t)$ first increases and then decreases eventually to 0. The susceptible population decreases, but converges to some limiting value $S(\infty) > 0$. The removed population is monotone increasing and converges to some value $R(\infty) < N$. Some of the susceptible population who do not get infected are protected by the recovered population surrounding them. This is known as *herd immunity*.

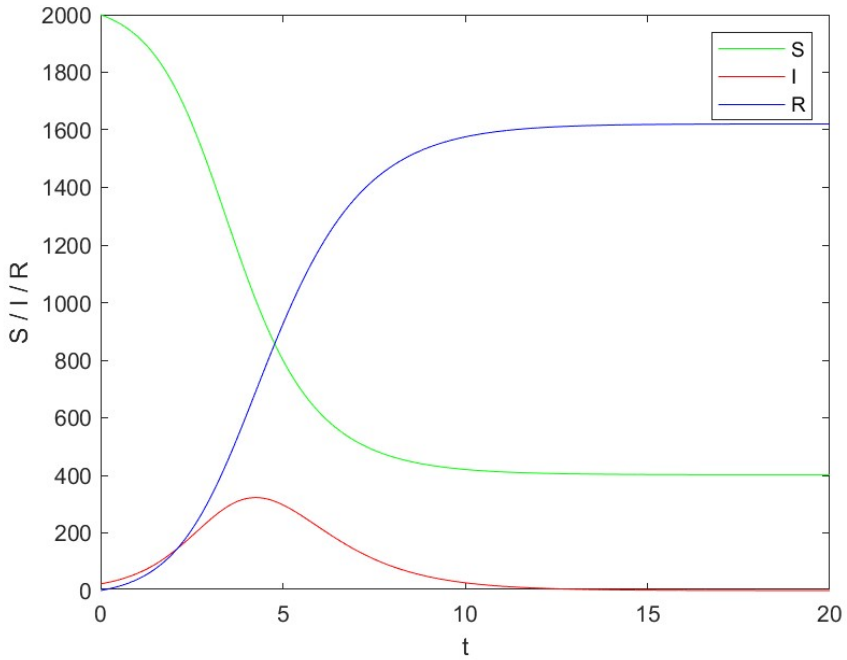

The SIR Model

Deterministic simulations (2)

Initial conditions: $S(0) = 2000,\ I(0) = 23,\ R(0) = 0.$

$\beta = 1 \quad , \quad \alpha = 1.$

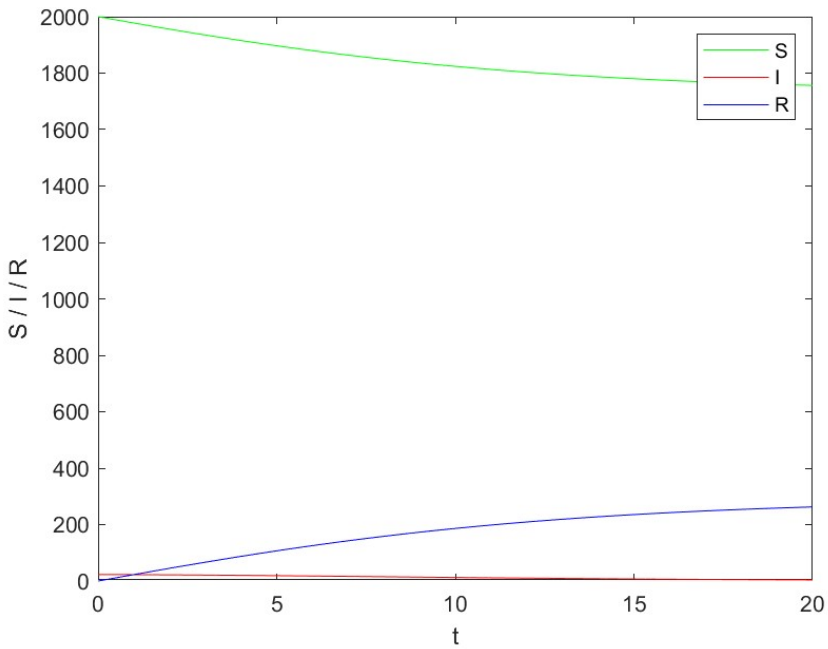

$\beta = 5 \quad , \quad \alpha = 1.$

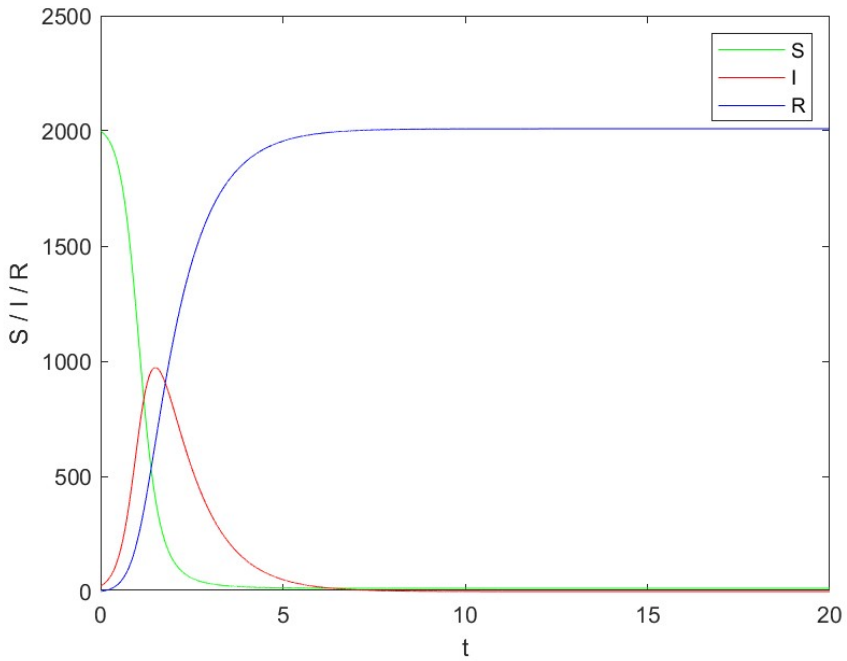

The SIR Model: The normalized form

For reasons of normalizations, we prefer to compute *fractions* of susceptible population, infected population, and of removed population:

$$s(t) = \frac{S(t)}{N} \quad , \quad i(t) = \frac{I(t)}{N} \quad , \quad r(t) = \frac{R(t)}{N}$$

In which case the model becomes:

$$\begin{cases} \frac{ds}{dt} &= -\beta si \quad , \quad s(0) = \frac{S(0)}{N} \\ \frac{di}{dt} &= \beta si - \alpha i \quad , \quad i(0) = \frac{I(0)}{N} \\ \frac{dr}{dt} &= \alpha i \quad , \quad r(0) = \frac{R(0)}{N} \end{cases} \qquad \text{(SIR Model)}$$

Note $s(t) + i(t) + r(t) = 1$ for all t (conservation of total population). The bad news: there is no closed form solution. The good news: some relationships can be expressed in closed form.

The SIR Model: Numerical solution vs. Agent Based Modeling

Similar to the SI model, one way of implementing an agent based simulation is to pretend the nonlinear term is linear in s $\tilde{\beta}(t) = \beta i(t)$. Thus the rate matrix is given by

$$\frac{d}{dt}\begin{bmatrix} s \\ i \\ r \end{bmatrix} = A \begin{bmatrix} s \\ i \\ r \end{bmatrix} \quad , \quad A = \begin{bmatrix} -\beta i & 0 & 0 \\ \beta i & -\alpha & 0 \\ 0 & \alpha & 0 \end{bmatrix}$$

and for a discretization step T_0, at time step $p > 0$, the transition matrix to transition from time $(p-1)T_0$ to pT_0 with $i_{p-1} = i((p-1)T_0)$ is given by

$$\Pi^{(p)} = \begin{bmatrix} e^{-\beta T_0 i_{p-1}} & 0 & 0 \\ \frac{\beta i_{p-1}}{\alpha - \beta i_{p-1}} e^{-\beta T_0 i_{p-1}} - \frac{\beta i_{p-1}}{\alpha - \beta i_{p-1}} e^{-\alpha T_0} & e^{-\alpha T_0} & 0 \\ 1 - \frac{\alpha}{\alpha - \beta i_{p-1}} e^{-\beta T_0 i_{p-1}} + \frac{\beta i_{p-1}}{\alpha - \beta i_{p-1}} e^{-\alpha T_0} & 1 - e^{-\alpha T_0} & 1 \end{bmatrix}, p = 1, 2, \ldots$$

The SIR Model: Numerical solution vs. Agent Based Modeling (2)

At every time step, the transition matrix can be approximated by the first order term in $e^{tA} \approx I + tA$:

$$\Pi_E^{(p)} = \begin{bmatrix} 1 - \beta T_0 i_{p-1} & 0 & 0 \\ \beta T_0 i_{p-1} & 1 - \alpha T_0 & 0 \\ 0 & \alpha T_0 & 1 \end{bmatrix}$$

This matrix has a direct interpretation, and allows for a more flexible implementation. for instance, for the transition $S \to I$ at time step p (i.e., from time $(p-1)T_0$ to time pT_0):

- If $State(Agent\ k\ at\ time\ (p-1)T_0) = Susceptible$ then:
  1. Draw a random variable u distributed uniformly over the set of all agents, $u \sim U\{1, 2, ..., M\}$.
  2. If $State(Agent\ u\ at\ time\ (p-1)T_0) = Infected$ then:
     1. Draw a random number $z \sim U[0, 1]$ uniformly distributed in $[0, 1]$.
     2. If $z \in [1 - \beta T_0, 1]$ then $State(Agent\ k\ at\ time\ pT_0) = Infected$.
     3. Otherwise $State(Agent\ k\ at\ time\ pT_0) = Susceptible$
  3. Otherwise $State(Agent\ k\ at\ time\ pT_0) = Susceptible$.

The SIR Model: Numerical Solution vs. Agent Based Modeling (3)

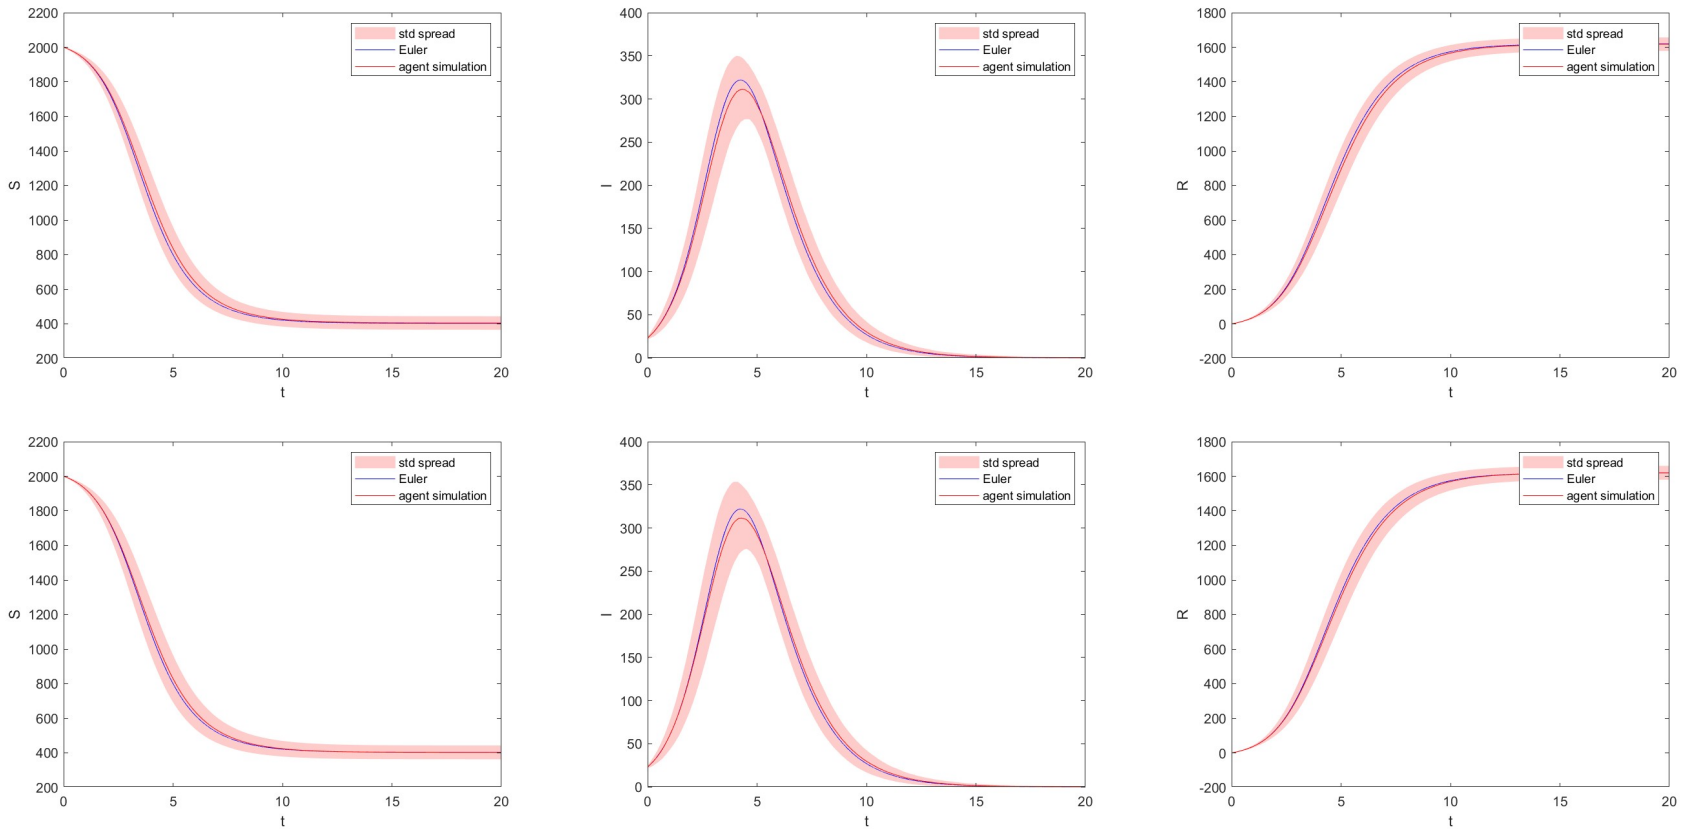

Figure: Results over 1000 simulations for $T_0 = 0.01$, with parameters $\beta = 2$ and $\alpha = 1$. The shaded area has semiwidth of one std of simulations. Top row utilizes the exponential matrix; the bottom row utilizes the linearized approximation for the transition matrix. The Euler scheme has a stepsize $h = 0.001$.

The SIR Model: Numerical Solution vs. Agent Based Modeling (4)

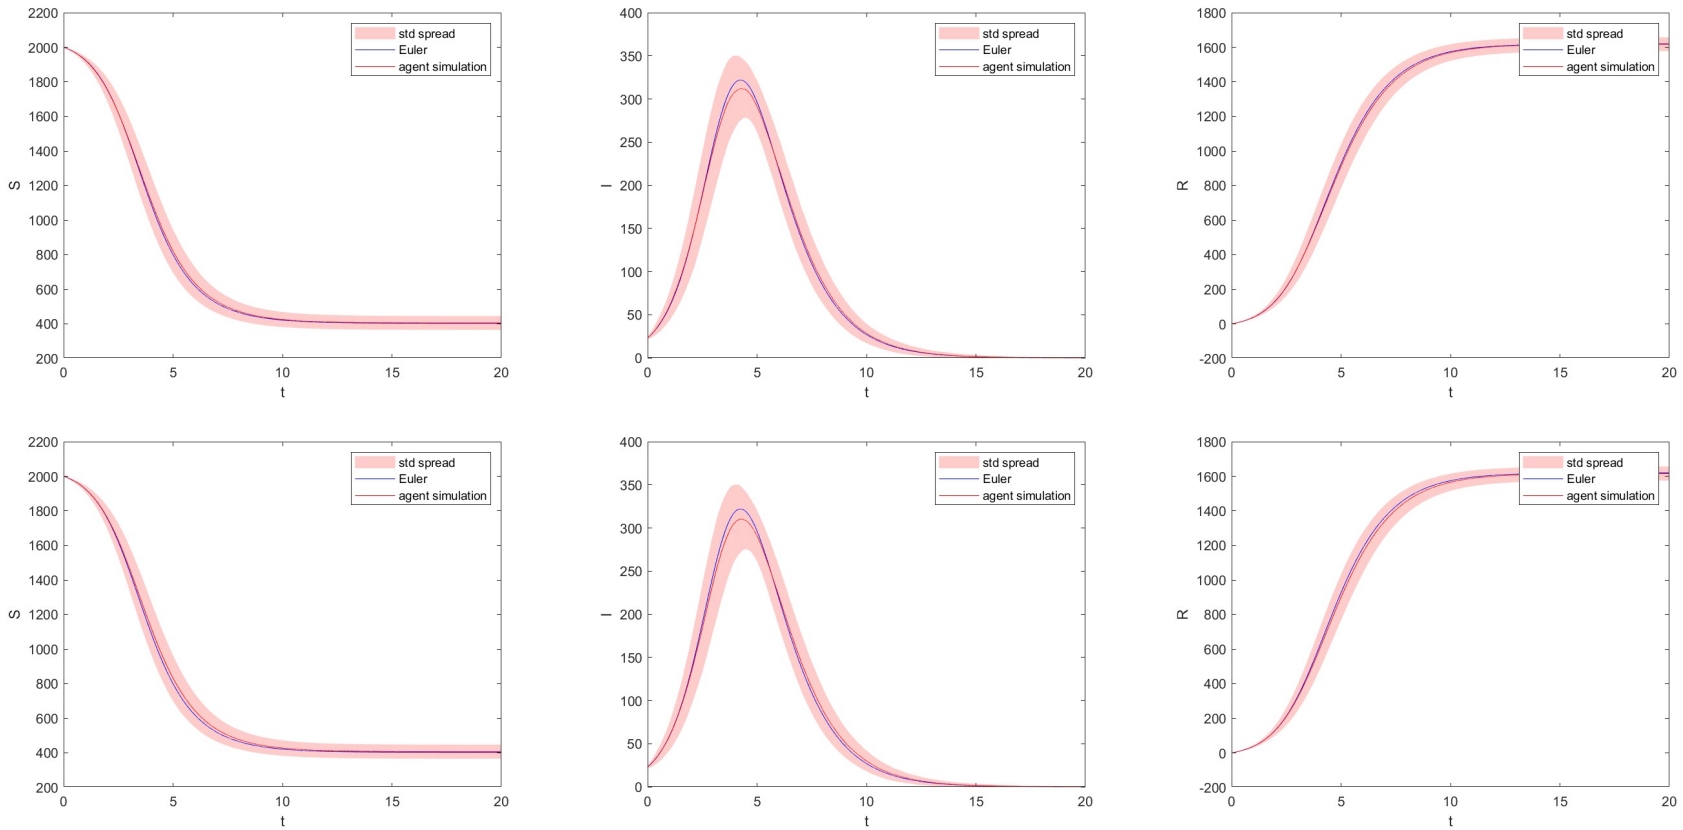

Figure: Results over 1000 simulations for $T_0 = 0.001$, with parameters $\beta = 2$ and $\alpha = 1$. The shaded area has semiwidth of one std of simulations. Top row utilizes the exponential matrix; the bottom row utilizes the linearized approximation for the transition matrix. The Euler scheme has a stepsize $h = 0.001$.

The SIR Model: Numerical Solution vs. Agent Based Modeling (5)

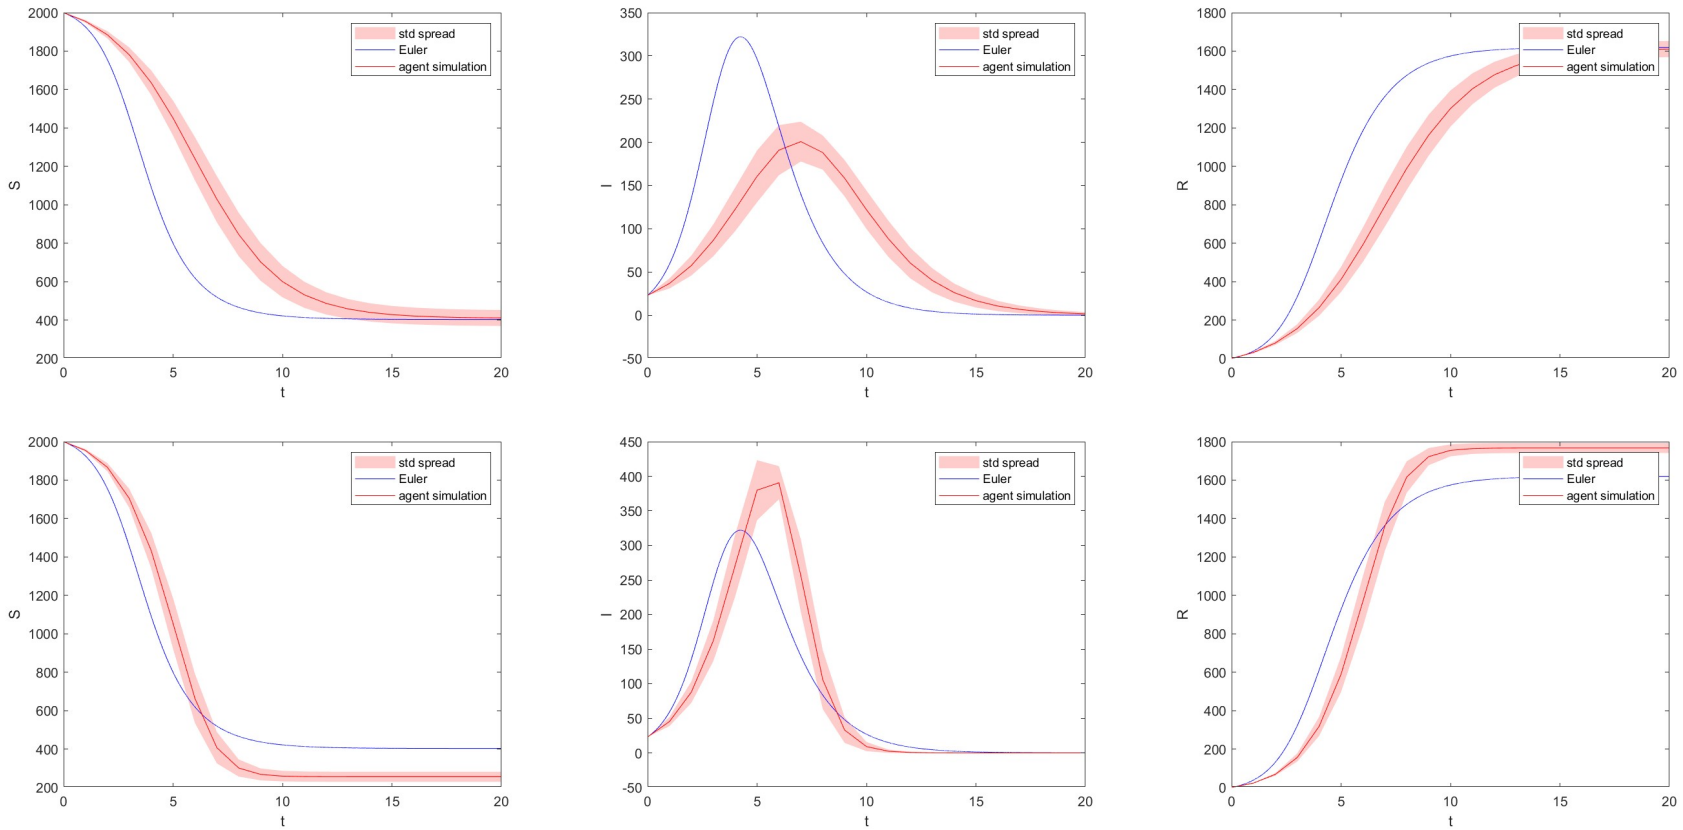

Figure: Results over 1000 simulations for $T_0 = 1.0$, with parameters $\beta = 2$ and $\alpha = 1$. The shaded area has semiwidth of one std of simulations. Top row utilizes the exponential matrix; the bottom row utilizes the linearized approximation for the transition matrix. The Euler scheme has a stepsize $h = 0.001$.

The SIR Model

Analytic expressions

For the normalized systems of equations, divide the equation for i by equation for s and use chain rule:

$$\frac{di}{ds} = -1 + \frac{\alpha}{\beta}\frac{1}{s}$$

The ratio $R_0 = \frac{\beta}{\alpha}$ is known as the *reproduction ratio*, or the *contact number*. Its meaning: β represents the number of close contacts per day per one infected individual; $\frac{1}{\alpha}$ is the average infectious period (or, the average number of days an infected person remains contagious). Hence R_0 represents the average number of close contacts per infected individual. Use separability of this Diff Eq. and integrate both sides:

$$i(t) - i(0) = s(0) - s(t) + \frac{1}{R_0}\left(log(s(t)) - log(s(0))\right)$$

Thus $i + s - \frac{1}{R_0} log(s)$ must stay constant over time.

The SIR Model

Analytic calibrations

We obtained that $i + s - \frac{1}{R_0} log(s)$ must stay constant over time.
Approximation: At time $t = 0$, $s(0) \approx 1$ (assuming little infections and recovered people) and $i(0) \approx 0$ (very few infected people compared to the total population).

$$i(t) + s(t) - \frac{1}{R_0} log(s(t)) = 1$$

What happens for $t \to \infty$?
One thing for sure: $i(\infty) = 0$. What happens with $s(\infty)$? We obtain the following equations:

$$1 = s(\infty) - \frac{1}{R_0} log(s(\infty))$$

$$R_0 = \frac{log(s(\infty))}{s(\infty) - 1}$$

The SIR Model

Herd Immunity

With the approximation $i(0) = 0$, $s(0) = 1$, the plot of $log(s(\infty))/(s(\infty)-1)$ as function of $s(\infty)$ is rendered in the left figure. For instance, if the contact number is $R_0 = 2$, then $s(\infty) \approx 0.2$. Thus 20% of population get protection from the 80% who have gotten infected and recovered.

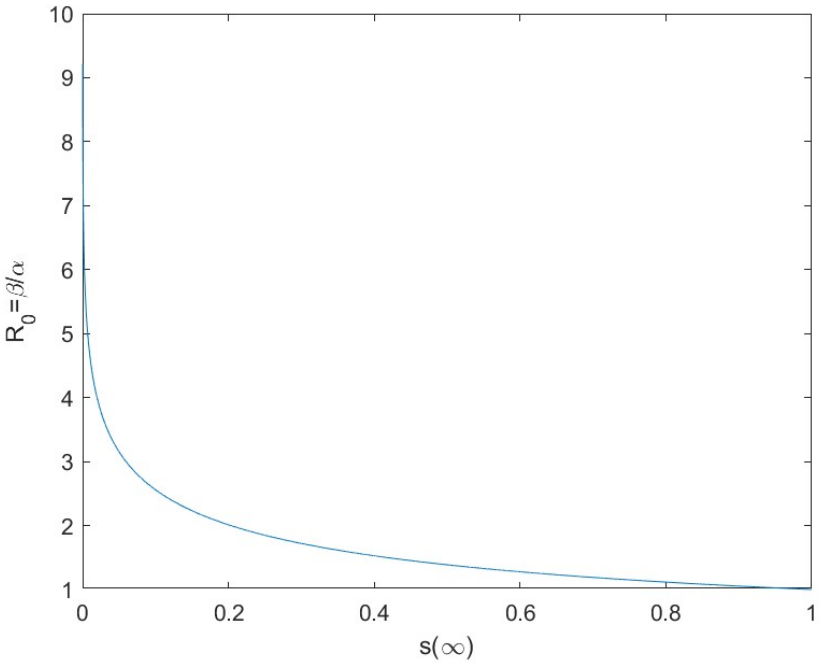

The SIR Model

Herd Immunity

With the approximation $i(0) = 0$, $s(0) = 1$, the plot of $log(s(\infty))/(s(\infty)-1)$ as function of $s(\infty)$ is rendered in the left figure. For instance, if the contact number is $R_0 = 2$, then $s(\infty) \approx 0.2$. Thus 20% of population get protection from the 80% who have gotten infected and recovered.

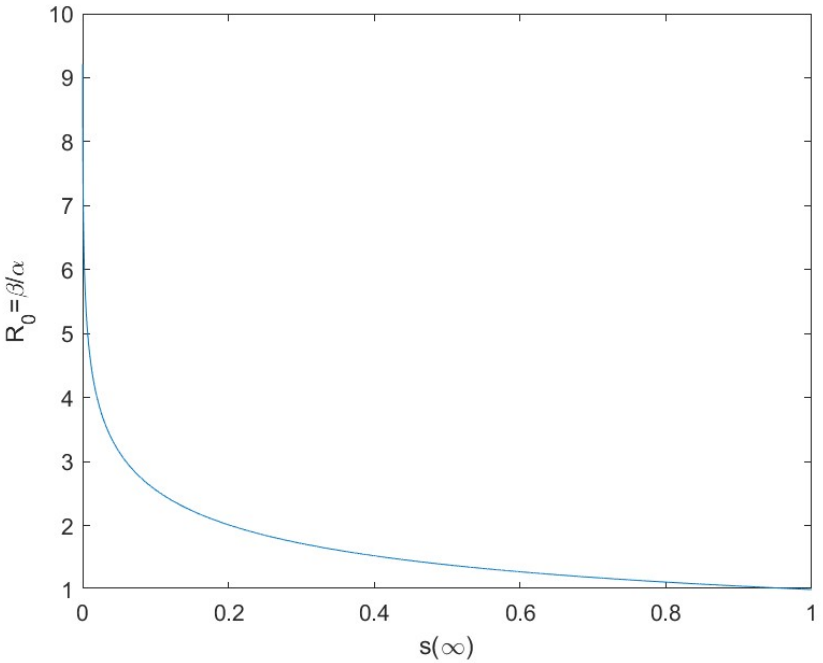

But what happens if $R_0 < 1$?

The SIR Model

Herd Immunity (2)

In reality, $i(0) > 0$ and $s(0) < 1$. Assume that still $r(0) = 0$, thus $s(0) + i(0) = 1$, but $log(s(0)) < 0$. We obtain:

$$i(t)+s(t)-\frac{1}{R_0}log(s(t)) = 1-\frac{log(s(0))}{R_0} \quad \Rightarrow \quad R_0 = \frac{log(s(\infty)) - log(s(0))}{s(\infty) - 1}$$

For $s(0) = 2000/2023$ we obtain the left plot. For $R_0 >> 1$, the previous approximation is still good.
For $R_0 < 1$, it follows that a significant number of susceptible individual do not get infected.

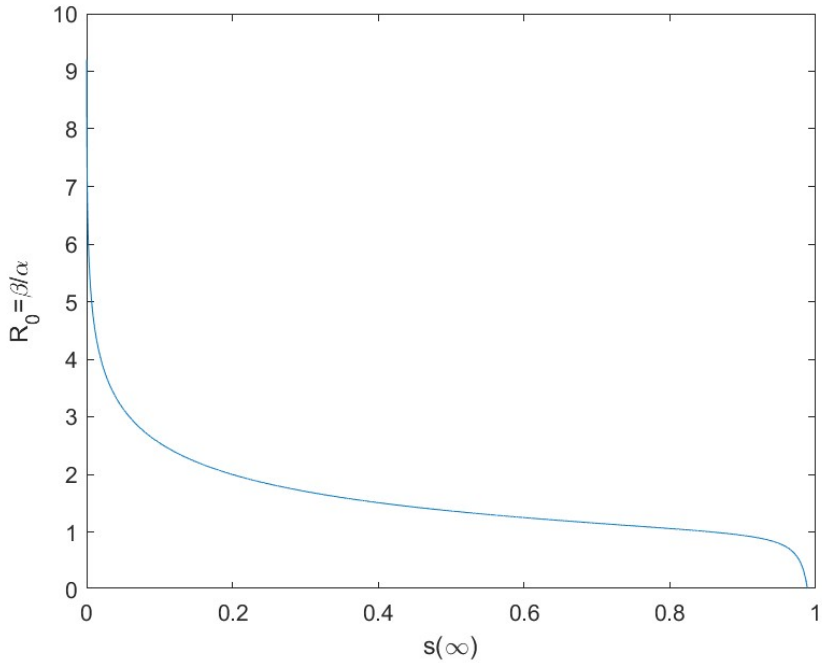

SIR Model with Two Outcomes

Since $I(t)$, or $i(t)$, is not monotone increasing sequence, the SIR model is appropriate for the time series of the *daily rates* of the number of infections. On the hand, the removed sequence $R(t)$, or $r(t)$, is monotone increasing. The "removed" compartment contains only two types of individuals: (1) individuals that recovered and gained full immunity $X(t)$, (so they will never get infected again), and (2) people who died, $Y(t)$. Thus $R(t) = X(t) + Y(t)$.
Assumption 1: $Y(t) = \gamma R(t)$ for all t. In other words, a fixed fraction γ of people who get infected eventually die, with the same infectious period as the individuals that eventually recovered and gained immunity:

$$\begin{cases} \frac{dS}{dt} &= -\beta S \frac{I}{N} \quad , \quad S(0) \\ \frac{dI}{dt} &= \beta S \frac{I}{N} - \alpha I \quad , \quad I(0) \\ \frac{dX}{dt} &= (1-\gamma)\alpha I \quad , \quad X(0) = (1-\gamma)R(0) \\ \frac{dY}{dt} &= \gamma \alpha I \quad , \quad Y(0) = \gamma R(0) \end{cases}$$

Note: For obtaining (S, I, X, Y) you need only to solve the SIR system and find $S(t), I(t), R(t))$, and then allocate, $X(t) = (1-\gamma)R(t)$ and $Y(t) = \gamma R(t)$.

Two-outcome SIR Model Analysis

The two-outcome SIR model has the form:

$$\begin{cases} \frac{dS}{dt} &= -\beta S\frac{I}{N} \quad , \quad S(0) \\ \frac{dI}{dt} &= \beta S\frac{I}{N} - \alpha I \quad , \quad I(0) \\ \frac{dX}{dt} &= (1-\gamma)\alpha I \quad , \quad X(0) = (1-\gamma)R(0) \\ \frac{dY}{dt} &= \gamma\alpha I \quad , \quad Y(0) = \gamma R(0) \end{cases}$$

Its normalized form in variables $s, i, r, x = \frac{X}{N}, y = \frac{Y}{N}$ is given by:

$$\begin{cases} \frac{ds}{dt} &= -\beta si \quad , \quad s(0) = \frac{S(0)}{N} \\ \frac{di}{dt} &= \beta si - \alpha i \quad , \quad i(0) = \frac{I(0)}{N} \\ \frac{dx}{dt} &= (1-\gamma)\alpha i \quad , \quad x(0) = (1-\gamma)\frac{R(0)}{N} \\ \frac{dy}{dt} &= \gamma\alpha i \quad , \quad y(0) = \gamma\frac{R(0)}{N} \end{cases}$$

Note the conservation laws: $N = S(t) + I(t) + X(t) + Y(t)$, $s(t) + i(t) + x(t) + y(t) = 1$. The deterministic system is initialzed by $(S(0), I(0), R(0))$ and its evolution is determined by the choice of three parameters: (α, β, γ). $X(t)$ and $Y(t)$ are computed from $R(t)$

SIR Model Analysis

The three parameters have the following meaning:

1. β represents the number of close contacts per day per one infected individual; differently said, it is the probability of disease transmission per contact(dimensionless) times the number of contacts per unit of time. Unit: day^{-1}
2. α is the removing rate of infectious individuals; its reciprocal is the infectious period. If no death, α represents the recovery rate from infections. Unit: day^{-1}
3. γ represents the probability of a fatal infection (death) once an individual gets infected. Unit: dimensionless.

$R_0 = \frac{\beta}{\alpha}$ is the *reproduction ratio* (or, the *contact number*) and represents the average number of infections caused by one infected individual.
Note: α and γ are parameters that characterize the infectious disease and cannot be controlled. Instead, β and ρ depend on human interactions, and therefore can be controlled by society/individuals (e.g., during a complete shut-down, $\beta \approx 0$).

SIR Model Calibration: LSE

For calibration and testing we are using two pieces of measured data: the *daily infection rates*, $\{I(0), \cdots, I(T_{max})\}$, and the time series of *cumulative deaths*, $\{Y(0), \cdots, Y(T_{max})\}$. Note that, if we know γ and N we can compute $R(0) = \frac{Y(0)}{\gamma}$ and $S(0) = N - I(0) - R(0)$. At the onset of an infectious disease it is the likely the case that $R(0) = Y(0) = 0$ and $I(0)$ is small and given by the first detected cases. Then $S(0) = N$.

Assumption 2: The time series of detected infections undercounts the actual number of infections. Specifically we assume $I(t) \approx \rho I_{sim}(t)$, where $\rho \leq 1$ is a parameter that represents the undercounting factor.

The least-squares estimator (LSE) finds parameters $\alpha, \beta, \gamma, \rho$ that minimize:

$$\begin{array}{c} \text{minimize} \\ \alpha, \beta, \gamma, \rho \geq 0 \\ \gamma, \rho \leq 1 \end{array} \quad I(\alpha, \beta, \gamma, \rho) := c_I \sum_{t=0}^{T_{max}} (I(t) - \rho I_{sim}(t))^2 + c_Y \sum_{t=0}^{T_{max}} (Y(t) - \gamma R_{sim}(t))^2.$$

Weights $c_I, c_Y \geq 0$ are chosen by user depending on how accurate are the two measured time series. If infections go unreported, set $c_I = 0$, $c_Y = 1$. Alternatively, choose $c_I = c_Y = 1$.

SIR Model Calibration: LSE (2)

The procedure works like this: $(S_{sim}(t), I_{sim}(t), R_{sim}(t))$ are simulated with a numerical solver for the SIR model with parameters (α, β) initialized at $(S(0), I(0), R(0))$. Parameters γ and ρ are then obtained by solving two independent optimization problems:

$$\hat{\rho} = \underset{0 \leq \rho \leq 1}{\text{argmin}} \sum_{t=0}^{T_{max}} (I(t) - \rho I_{sim}(t))^2 \quad , \quad \hat{\gamma} = \underset{0 \leq \gamma \leq 1}{\text{argmin}} \sum_{t=0}^{T_{max}} (Y(t) - \gamma R_{sim}(t))^2$$

Then compute the objective function $J(\alpha, \beta) = I(\alpha, \beta, \hat{\gamma}, \hat{\rho})$ and minimize over the set of pairs (α, β) used in simulations.

SIR Model with Vitals

A simple modification of the SIR vanilla model is to consider vital signals, such as births and deaths at separate processes. In normalized form this becomes:

$$\left\{\begin{array}{rcl} \frac{ds}{dt} & = & \frac{\Lambda}{N} - \beta si - \mu s \quad , \quad s(0) = \frac{S_0}{N} \\ \frac{di}{dt} & = & \beta si - \alpha i - \mu i \quad , \quad i(0) = \frac{I_0}{N} \\ \frac{dr}{dt} & = & \alpha i - \mu r \quad , \quad r(0) = \frac{R_0}{N} \end{array}\right. \qquad \text{(SIR Model)}$$

where $\Lambda \geq 0$ is the constant source of births (=number of births/day) and $\mu \geq 0$ is the natural death rate (i.e., in the absence of this virus). Its reciprocal $1/\mu$ represents the average life expectancy.