

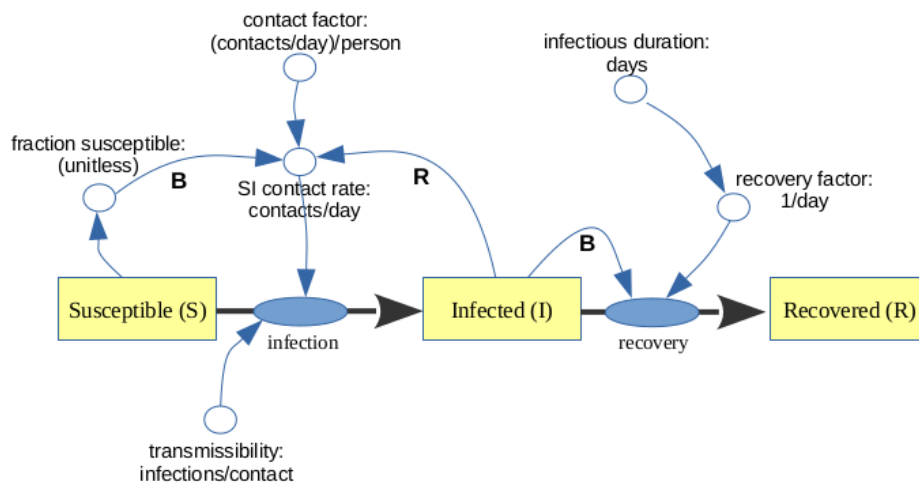
The SIR epidemiology model

Epidemiology is the study of disease outbreaks, distribution, and control. One of the simplest models of disease propagation in a “well-mixed” society is the classic **SIR model**, whose letters stand for “Susceptible,” “Infected,” and “Recovered.” In this basic model, those are the three stocks which contain human beings in various states. It is interesting to see what we can learn about the degree to which a disease will permeate a society based on some simple yet sensible assumptions.

The **Susceptible** stock denotes persons who are vulnerable to getting a contagious disease but who have not yet been exposed to it. (Non-susceptible persons, if they exist, don’t enter into the model’s dynamics at all and are therefore ignored.*) **Infected** individuals are those who are currently suffering from the disease *and are still contagious*. Once a person is no longer contagious, they are effectively in the **Recovered** state even though they may still experience symptoms.

One basic assumption is that we have a **closed population**, meaning that our stocks transfer quantities from one to another but never to or from *clouds*. This means that people never enter or leave the population but only (possibly) transfer between categories. Further, the transfers are all “one way”: a person can move from the susceptible state to the infected state, but never back; and from the infected state to the recovered state, but never back. This implies, among other things, that the disease is like the chicken pox: once you have it, you’re effectively inoculated against ever having it again.

A version of the entire model is reflected in the following stock-and-flow diagram:



*If present in sufficient numbers, however, they could possibly impact the `fraction_susceptible`.

The theory behind the model

Like the Lotka-Volterra model, the SIR model uses the notion of “encounters” which are proportional to the *product* of two populations. Just as both a bat and a mouse were required to have a killing event, so an Infected individual must come into contact with a Susceptible one to pass on a disease. The rate of infection is therefore based on several factors:

1. The `contact_factor` expresses how many other persons an individual comes into physical contact with per unit time. It is a property of a society, not a disease.
2. A disease’s `transmissibility` reflects how easily communicable the disease is between two people who have come in contact. Some diseases are more contagious than others.
3. In addition to the above two constant factors, there are also the changing fractions of the society that are in the S and I states. We’ll model the `frac_susc` (“fraction susceptible”) of the population, so that for a given infected person, we can estimate how many *susceptible* people (rather than just people-in-general) he/she comes into contact with each day.

That third item is the trickiest, so let’s unpack it. Suppose Filbert receives the bad news that he has the disease. He is therefore in the **Infected** stock. The question we want to ask is: how many *Susceptible* people per day will the poor Filberts come in contact with (and thus possibly infect) each day?

The stock-and-flow diagram calls this the “**SI contact rate**” since it’s the rate at which S’s and I’s will come in contact. (This is the only case in which the disease can be transmitted.) It’s in units of $\frac{\text{contacts}}{\text{day}}$. We compute:

$$\text{SI_contact_rate} = \text{frac_susc} \times \text{contact_factor} \times I_{i-1}$$

A quick units check confirms we’re not insane:

$$\frac{\text{contacts}}{\text{day}} = \text{unitless} \times \frac{\frac{\text{contacts}}{\text{day}}}{\text{person}} \times \text{people} = \frac{\text{contacts}}{\text{day}} \quad \checkmark$$

So for example, if every Filbert comes into physical contact with 4 other people each day, about $\frac{3}{4}$ ths of whom are Susceptible, and there are 50 such Filberts, then there will be $\frac{3}{4} \times 4 \times 50 = 150$ contacts between S/I pairs every day. That’s 150 opportunities for things to go bad!

On the recovery side of things, all that really matters is (1) how many Infected individuals there are, and (2) how long, on average, each one remains infected. If a person on average is contagious for 5 days, that means that $\frac{1}{5}$ th of the infected population recovers every day. This rate controls how quickly people are removed from the dangerous Infected pool where their presence threatens more susceptible individuals.

Implementation

Here's a realization of the SIR model in Python:

```
mean_infectious_duration = 2          # days
transmissibility = .25                # infections/contact
contact_factor = 5.1                 # (contacts/day)/person
recovery_factor = 1/mean_infectious_duration # 1/day

start_x = 0                          # days
end_x = 365 / 3                      # days
delta_x = 1/24                       # days
x = np.arange(start_x, end_x, delta_x) # days

S = np.empty(len(x))                # individuals
I = np.empty(len(x))                # individuals
R = np.empty(len(x))                # individuals
S[0] = 2000
I[0] = 1
R[0] = 0

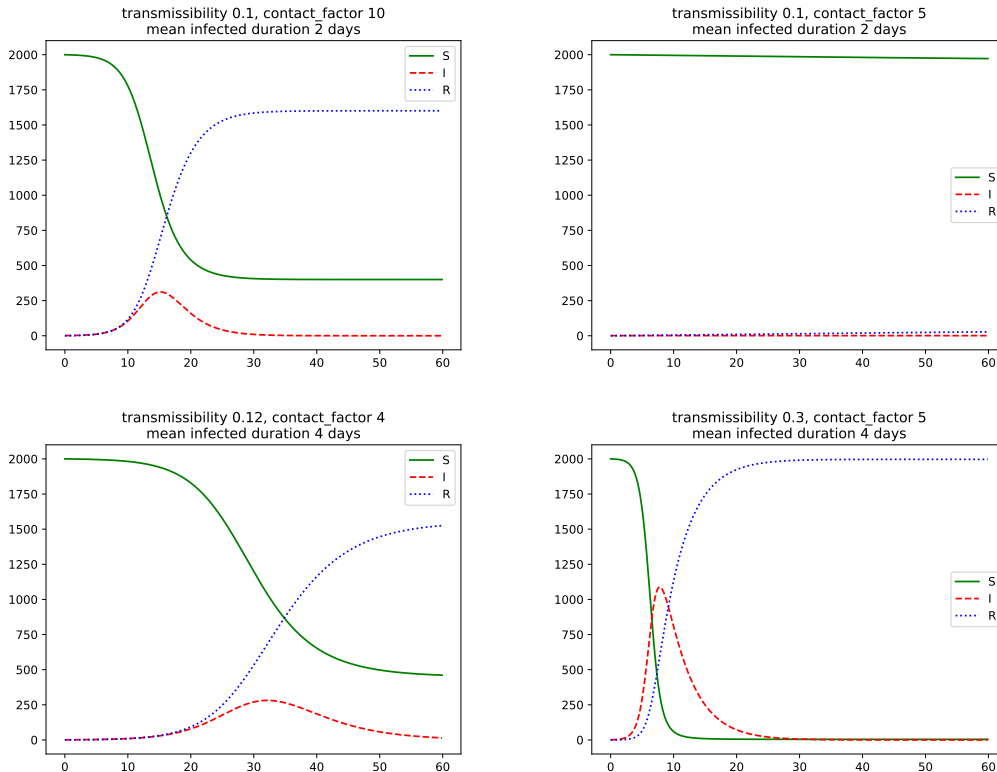
for i in range(1,len(x)):

    # Flows.
    frac_susc = S[i-1]/(S[i-1] + I[i-1] + R[i-1])
    SI_contact_rate = frac_susc * contact_factor * I[i-1] # contacts/day
    infection = SI_contact_rate * transmissibility        # infections/day
    recovery = I[i-1] * recovery_factor                   # recoveries/day

    # Primes.
    S_prime = -infection
    I_prime = infection - recovery
    R_prime = recovery

    # Stocks.
    S[i] = S[i-1] + S_prime * delta_x
    I[i] = I[i-1] + I_prime * delta_x
    R[i] = R[i-1] + R_prime * delta_x
```

Here are some sample outputs for various values of the parameters:



As expected, the populations level off at an equilibrium in which they are all either recovered from the disease, or never had it. The relative levels of these stocks is an interesting issue.

The “basic reproduction number” R_0

It turns out that an interesting quantity can be computed from the various parameters in this model. Suppose we multiply the **transmissibility** (in $\frac{\text{infections}}{\text{contact}}$), the **contact_factor** (in $\frac{\text{contacts/day}}{\text{person}}$), and the **mean_infectious_duration** (in days). We then get a quantity in units of $\frac{\text{infections}}{\text{person}}$ which we call “ R_0 ” (pronounced “R-nought”):

$$R_0 = \text{transmissibility} \times \text{contact_factor} \times \text{mean_infectious_duration}$$

This number, which combines features of both the community and the disease, can be thought of as “the number of further infections to be expected for each person who already has the disease.” The actual number of further infections of course depends on how much of the population has already been infected: early in an outbreak, it is higher per infected person than later. So perhaps a better way to think of R_0 is: “the expected number of secondary cases produced by a typical infected individual early in an epidemic.”

(Incidentally, R_0 can be thought of as a unitless quantity, since it's sort of in units of “new people who'll be infected per current person infected,” or “people per person.”)

Now it turns out this number has a critical threshold on which all else turns. If R_0 is less than 1, then although a few people will get the disease, it is doomed to die out before it gets much momentum, because *each infected person infects on average less than one other person*. This means that a sustained chain reaction isn't possible: the disease lacks enough staying power to increase its victims. The limiting factor is the inability of the disease itself. But when R_0 is greater than 1, the disease is bound to move through the entire population, causing an epidemic. This doesn't mean that *every* single person gets the disease, of course. But it does mean that a significant fraction will. In this case, the limiting factor is that so many people have already had the disease that there isn't a high enough percentage of susceptible victims left to keep it going.

In a simulation, this sort of threshold is called a “**tipping point**” because the behavior of the model changes qualitatively (not just quantitatively) when parameters are altered within a certain sensitive, critical range.