Transmission of infections on contact networks I: An introduction

Winfried Just
Department of Mathematics, Ohio University

Daniel Figueiredo, Olena Kostylenko, Ana Paião
Department of Mathematics, University of Aveiro

July 19, 2018
2^{nd} Portuguese Meeting in Biomathematics
University of Aveiro

Olá!

WJ: Olá!

Audience and Ana: Olá!

WJ: I'm Vinny.

Ana: Nice to meet you, Vinny. I'm Ana.

WJ: Nice to meet meet you, Ana. I'm a mathematician.

Ana: Oh . . . Ok. What do mathematicians do for a living?

WJ: We play.

Ana: You mean, like, with toys??

WJ: Yes. Our toys are called mathematical models.

Ana: And you get paid for this???

WJ: Yes. Like children's toys, our models can teach us and others how the real world works. . . .

Ana: Are you serious????

WJ: . . . We are also toymakers. Our models can be used by others to understand the real world.

Olena wants to play with the toys

Ana: Hmm . . . well. I need to go now. Nice talking with you!

WJ: Same here.

(Ana exits)

WJ: (Shrug)

Olena: Can I play with your toys?

WJ: Absolutely! Why don't you join me up here?

Olena: Hi! I'm Olena.

WJ: Hi! Nice to meet you, Olena.

Olena: Nice to meet you, Vinny! So where are your toys?

WJ: First I need to tell you what my models are all about. I study mathematical epidemiology.

Olena: Is this like mathematical medicine?

Epidemiology

Medicine studies the causes, symptoms, and progression of infectious diseases.

Its goal is to **cure** diseases.

Epidemiology studies how infections spread.

Its goal is to find ways to **prevent** or at least **limit** the spread of infections.

The latter can be achieved by appropriate **control measures:**

- **vaccination**,
- **quarantine**,
- **behavior modifications** (for human diseases).

Mathematical models

Ana: What's math got to do with that?

WJ: Welcome back, Ana! That's exactly where mathematical models enter the picture!

In order to prevent infections, we need to understand how they spread.

But real life is complicated, messy, rich in meanings.

If we were to pay attention to everything at once, we would be in a mental fog.

If we want to understand what really goes on, we need to focus on the essentials and ignore that there is more.

- **Mathematical models** are toy worlds that ignore most of the messy details, cut right through the fog, and reduce a situation to its driving forces.

Daniel: Would this be like in that old song by Tina Turner?

What's love got to do with it? By Tina Turner

You must understand how the touch
Of your hand makes my pulse react
That it's only the thrill of boy meeting girl
Opposites attract

It's physical, only logical
You must try to ignore
That it means more than that

What's love got to do, got to do with it
What's love but a second hand emotion
What's love got to do, got to do with it
Who needs a heart when a heart can be broken

It may seem to you that I'm
Acting confused when you're close to me
If I tend to look dazed I've read
It someplace I've got 'cause to be

There's a name for it
There's a phrase that fits
But whatever the reason you do it for me

What's love got to do, got to do with it . . .
I've been taking on a new direction
But I have to say
I've been thinking about my own protection
It scares me to feel this way

What's love got to do, got to do with it . . .

Hi Daniel!

WJ: Would you like to join us?

Daniel: I'm Daniel. Nice to meet you, Vinny, Olena, and Ana!

All: Nice to meet you, Daniel!

WJ: Yes, I remember that song. It isn't relevant here though.

Daniel: Why not?

WJ: I have to give my presentation ...

Ana: Can you stop lecturing for a moment and just listen to what this woman is singing?

The song's true meaning explained to the professor ...

You must understand how the touch of your hand . . .
I've been thinking about my own protection . . .
Who needs a heart when a heart can be broken . . .

Olena: She is afraid that her heart will be broken and wants to take protective measures. So she needs to understand what drives the situation.

Acting confused when you're close to me
If I tend to look dazed I've read . . .

Daniel: She is in a mental fog, and cuts right through it:

Opposites attract. It's physical, only logical.

Ana: Can you relate to that?

You must try to ignore that it means more than that . . .

WJ: Yes! That's exactly how mathematical models work!!

So what's math got to do with it?

WJ: Mathematical epidemiology studies models were infections can spread in an abstract, greatly simplified setting.

Olena: Mathematical epidemiology would then be, like, spreading a "second-hand infection" in a model?

WJ: So to speak, yes.

Olena: And nobody gets hurt by it?

WJ: Absolutely nobody. Most importantly, without causing any actual suffering, we can explore many different options for control measures, like, for example, different policies for vaccinating people.

Olena: To see which one works best?

WJ: Right.

Olena: And then recommend that the best option be implemented?

WJ: That's exactly the point. You have put it in a nutshell.

Olena: Wow!!

So what's math got to do with it?

Ana: But control measures may be costly. For example, not enough vaccine may be available to protect the whole population.

WJ: Good point. This may be a problem.

Ana: So "best" would mean the strategy that will be expected to give the largest reduction in suffering among all options that are economically affordable?

WJ: I couldn't have said it any better.

Daniel: But you need to make sure your models are telling you the correct thing, or else real people might suffer!!!

WJ: Excellent point. So we will need to make our models as realistic as possible. Mathematical models are based on ignoring a lot of messy real-world details, but we need to be careful about which details we cannot ignore.

Summary: Mathematical models in epidemiology

- **Mathematical models** are toy worlds that ignore most of the messy details, cut right through the fog, and reduce a situation to its driving forces.
- In particular, **mathematical epidemiology** studies models were infections can spread in an abstract, greatly simplified setting.
- We can study how an infection spreads in a mathematical model.
- This will allow us to compare the effects of various control measures.
- The control measure that gives the largest reduction in suffering among all economically affordable ones should be implemented.
- Before basing an actual policy recommendation on it, we need to verify that the model's assumptions and predictions are realistic.

The art of mathematical modeling

Daniel: And how do you know which aspects you can ignore, and which ones you can't?

WJ: There is no algorithm for that. Making good choices about what to incorporate is the art of modeling.

Daniel: But how can one base decisions that may impact human suffering on an art?

Ana: I'm also wondering about this. Bad decisions may be costly.

Olena: Perhaps one can learn this art?

WJ: Like any art, you learn it by practicing it. We will get a little practice here.

Olena: Yes, please show us how to do it!

WJ: But first I need to do some more lecturing.

Epidemic models in a nutshell: Some terminology

Infectious diseases that are triggered when **pathogens** such as viruses or bacteria enter a **host** (human, animal, plant).

Mathematical epidemiology studies how the infection **spreads between hosts** of a given **population.**

Here we will focus on diseases whose transmission requires **direct contact** (of a certain type) between hosts.

Epidemic models in a nutshell: Compartments

All known mathematical models of disease transmission are based on the notion of **compartments.**

Each host is assumed to be, at any given time, in one of a few **states.**

For example, the states could be:

- State S: **Susceptible** (to infection).
- State I: **Infectious** (has the disease and can infect others).
- State R: **Removed** (recovered with immunity or died from the disease. Will not infect others).

The **S**-, **I**-, and **R**-compartments are the **sets** of hosts that are susceptible, infectious, and removed, respectively.

The spread of the infection is modeled as **movement** (of sorts) of hosts from one compartment into another.

SIR-models

Models of **type** *SIR* may be suitable for **immunizing infections** like measles or chicken pox where recovery confers permanent immunity.

Schematically, such a model can be represented as:

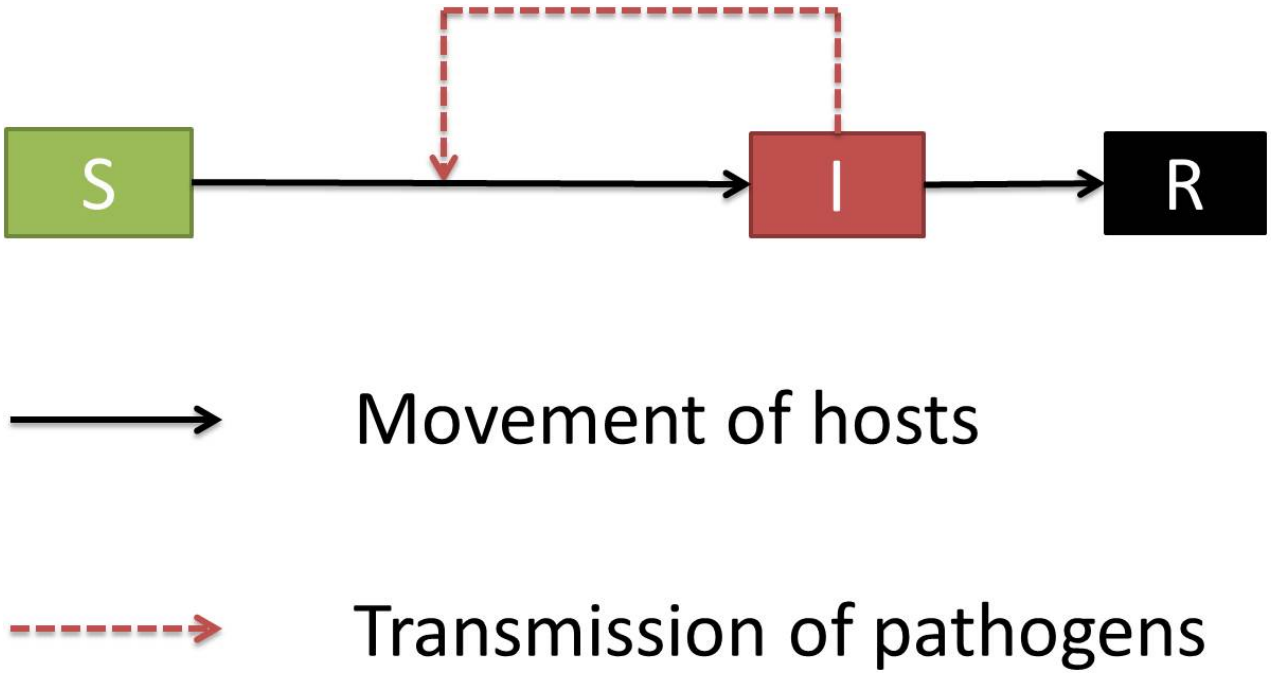

How would an outbreak get started?

A single infected host is introduced into a population of susceptibles.

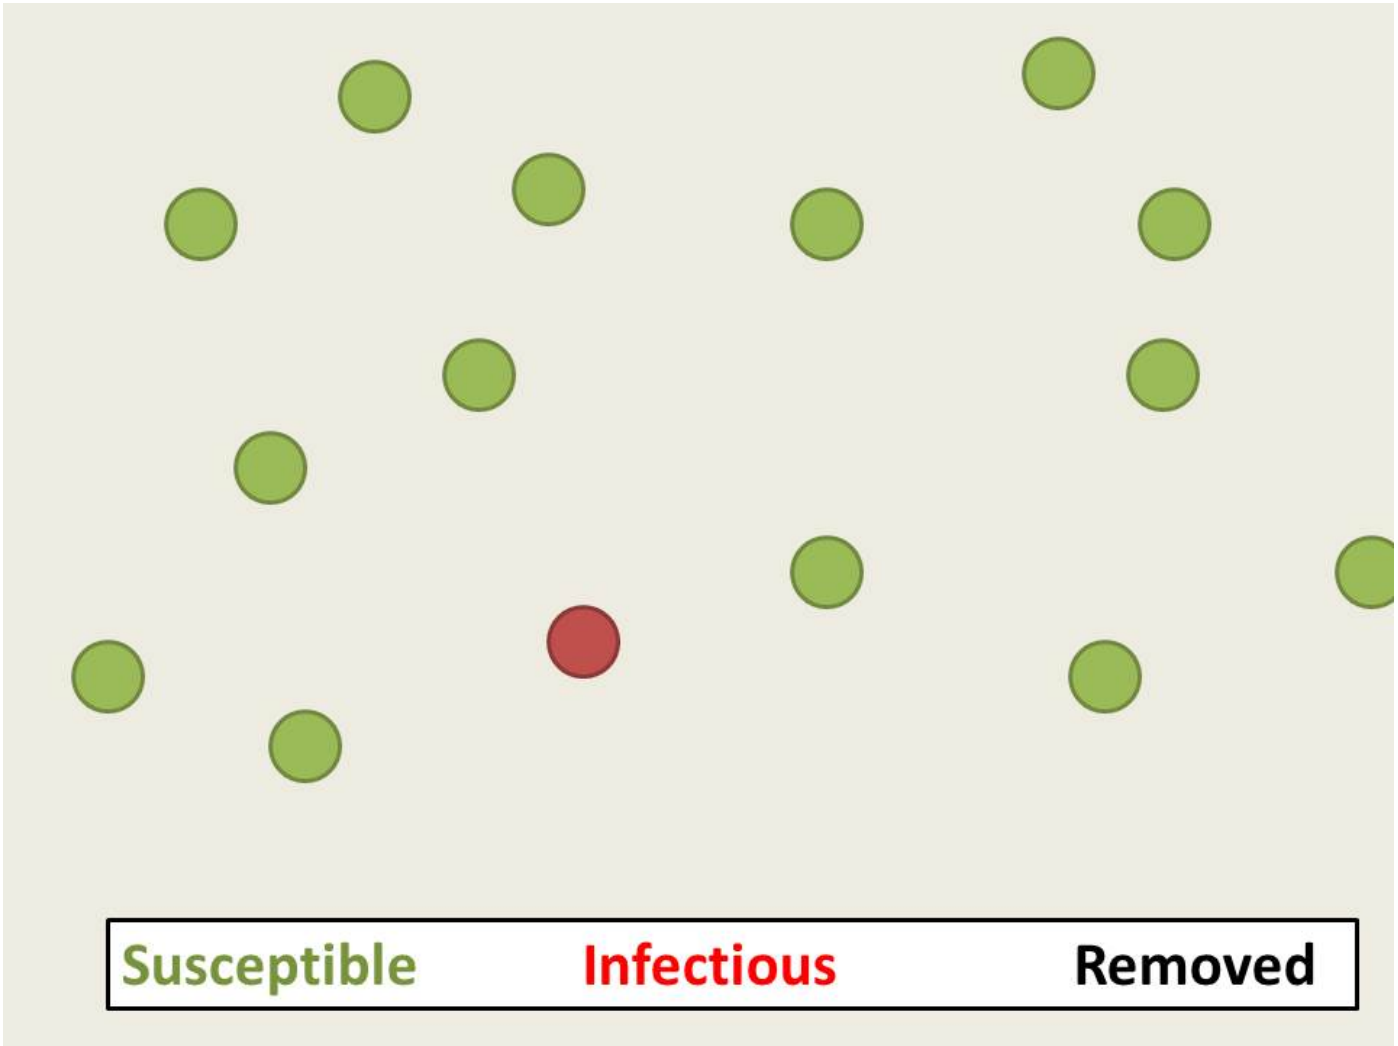

A hypothetical outbreak

A new infection occurs.

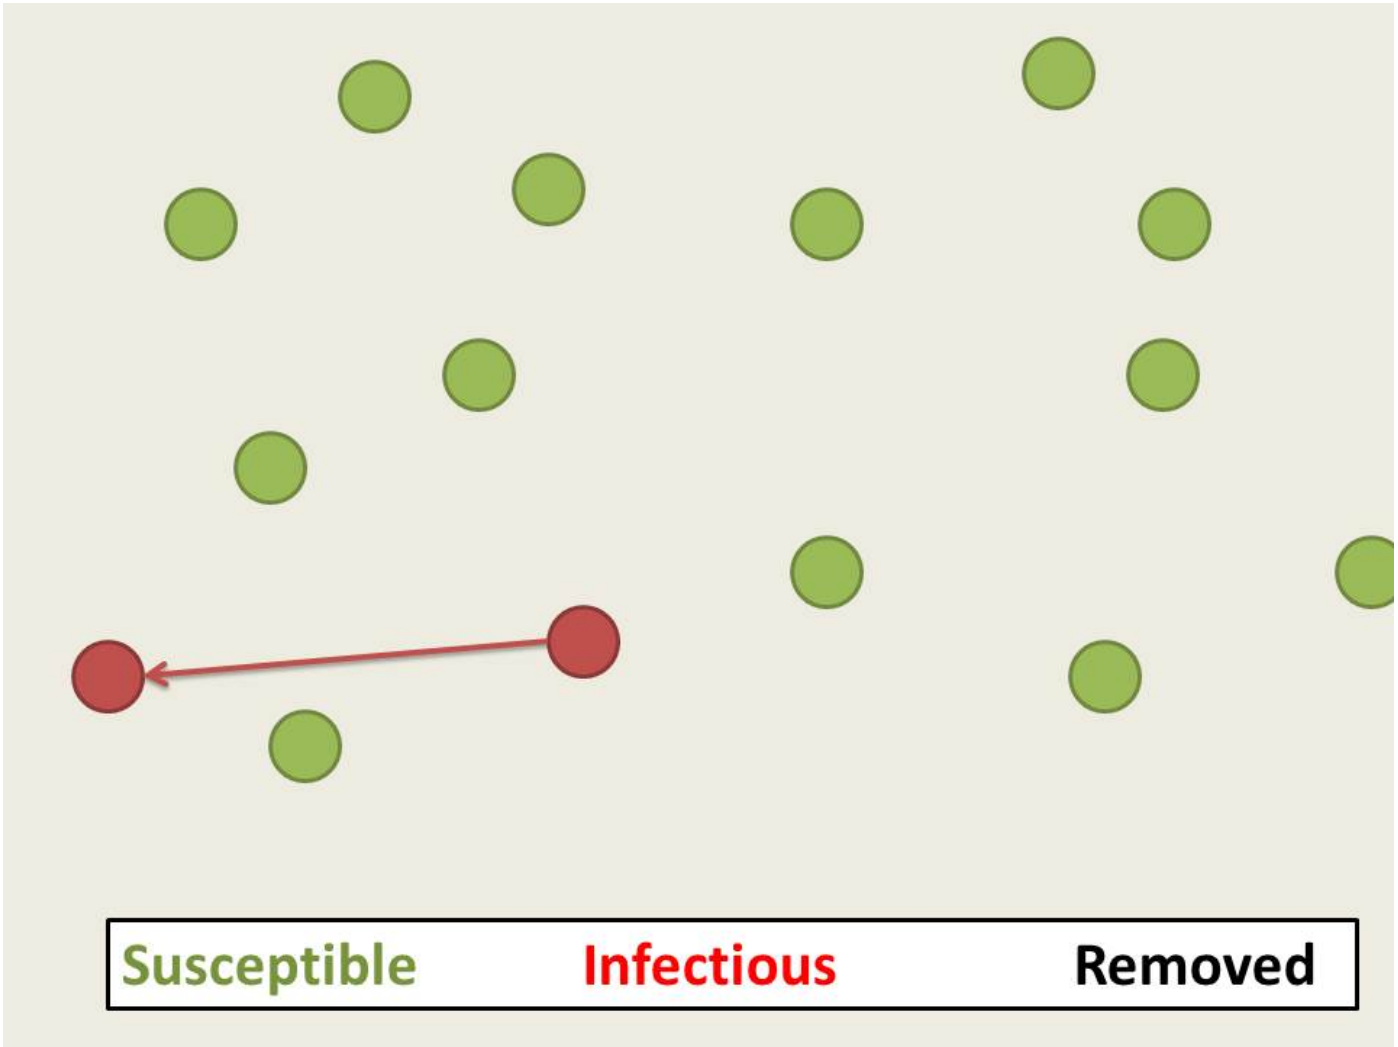

A hypothetical outbreak

A new infection occurs.

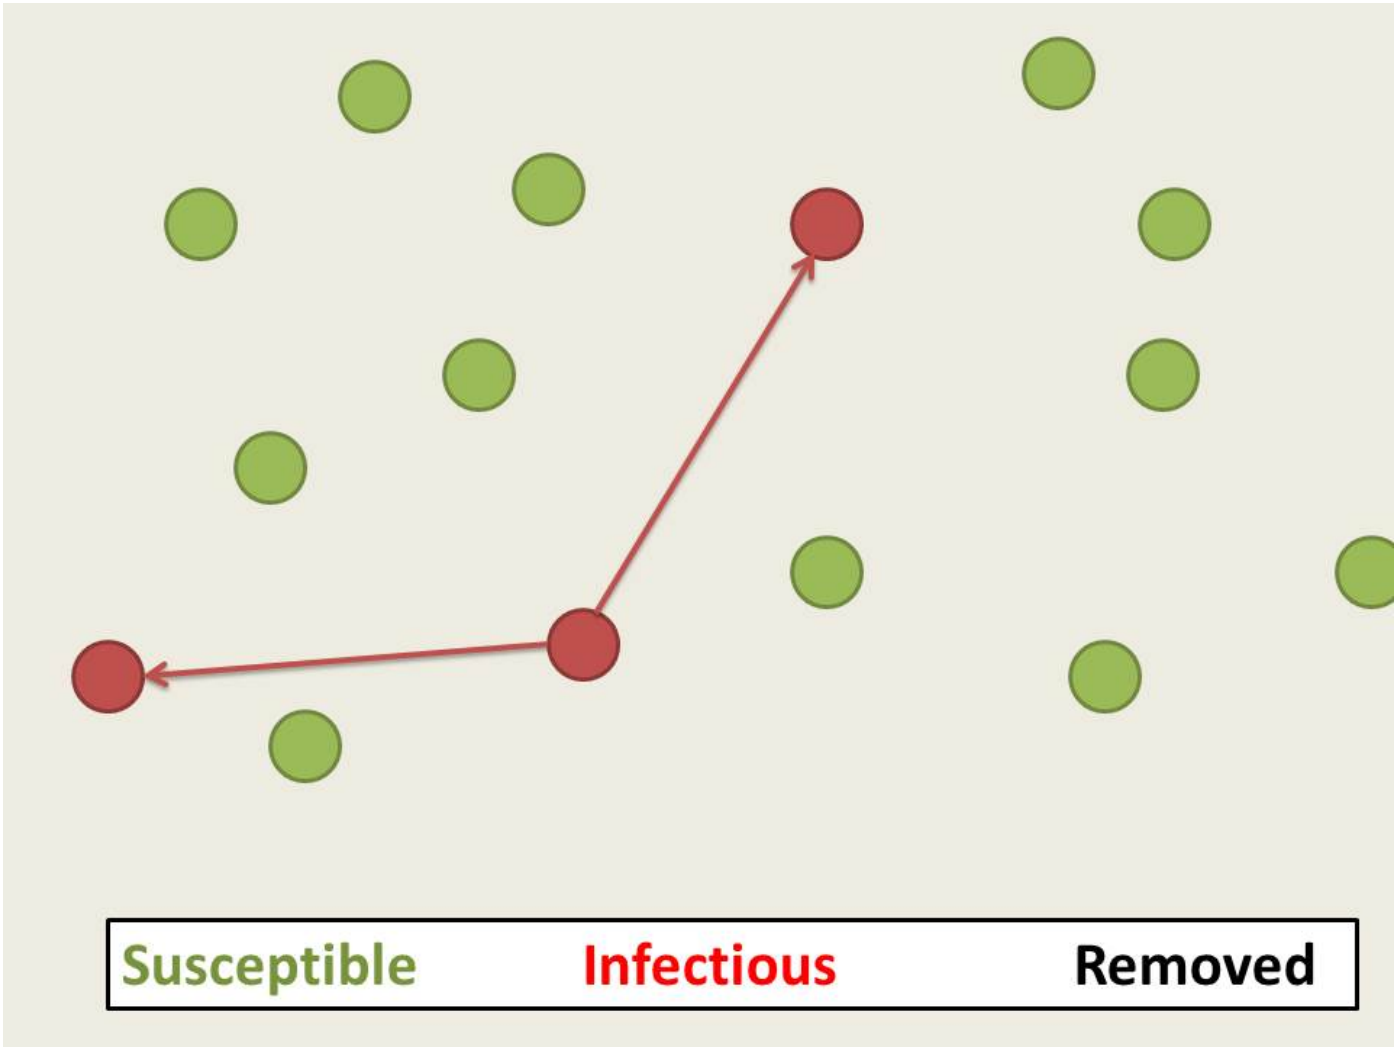

A hypothetical outbreak

A new infection occurs.

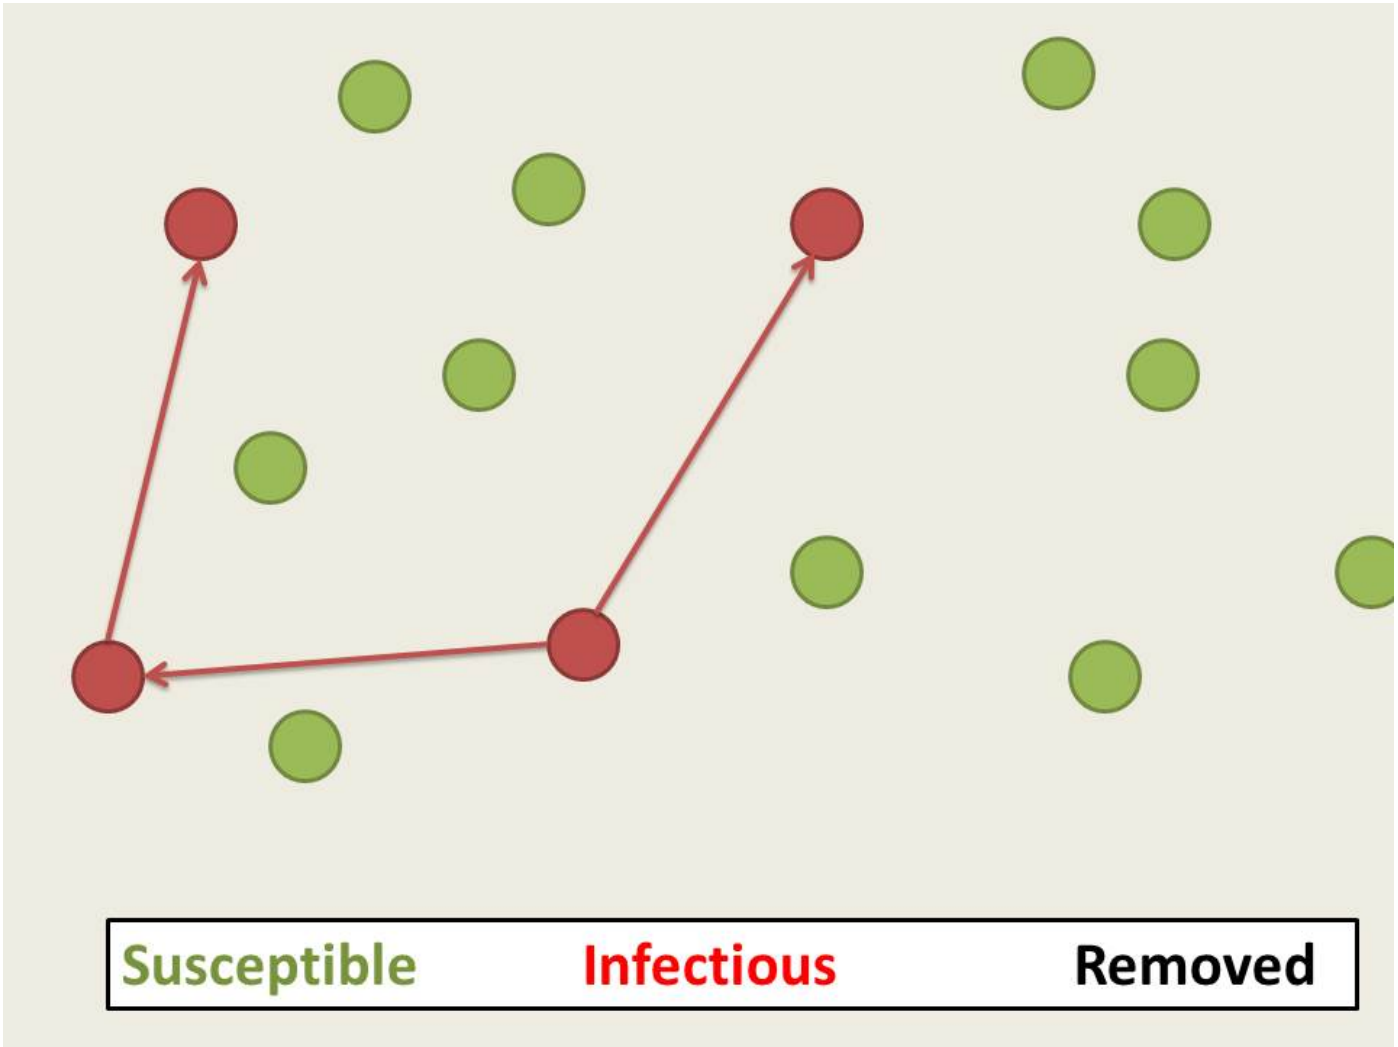

A hypothetical outbreak

An infectious host is removed.

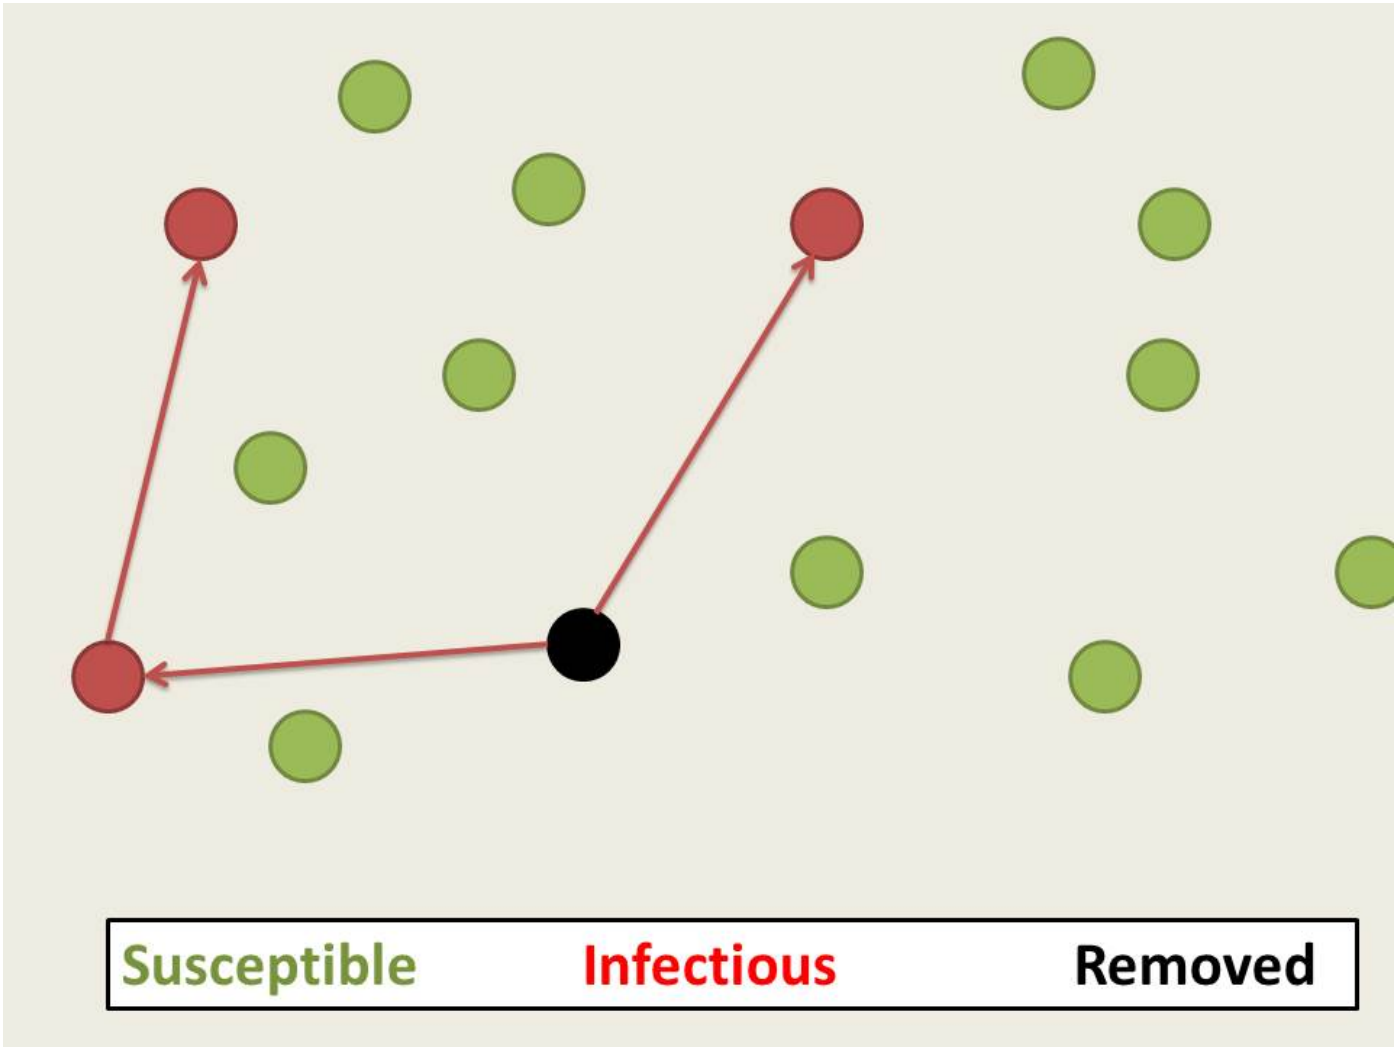

A hypothetical outbreak

An infectious host is removed.

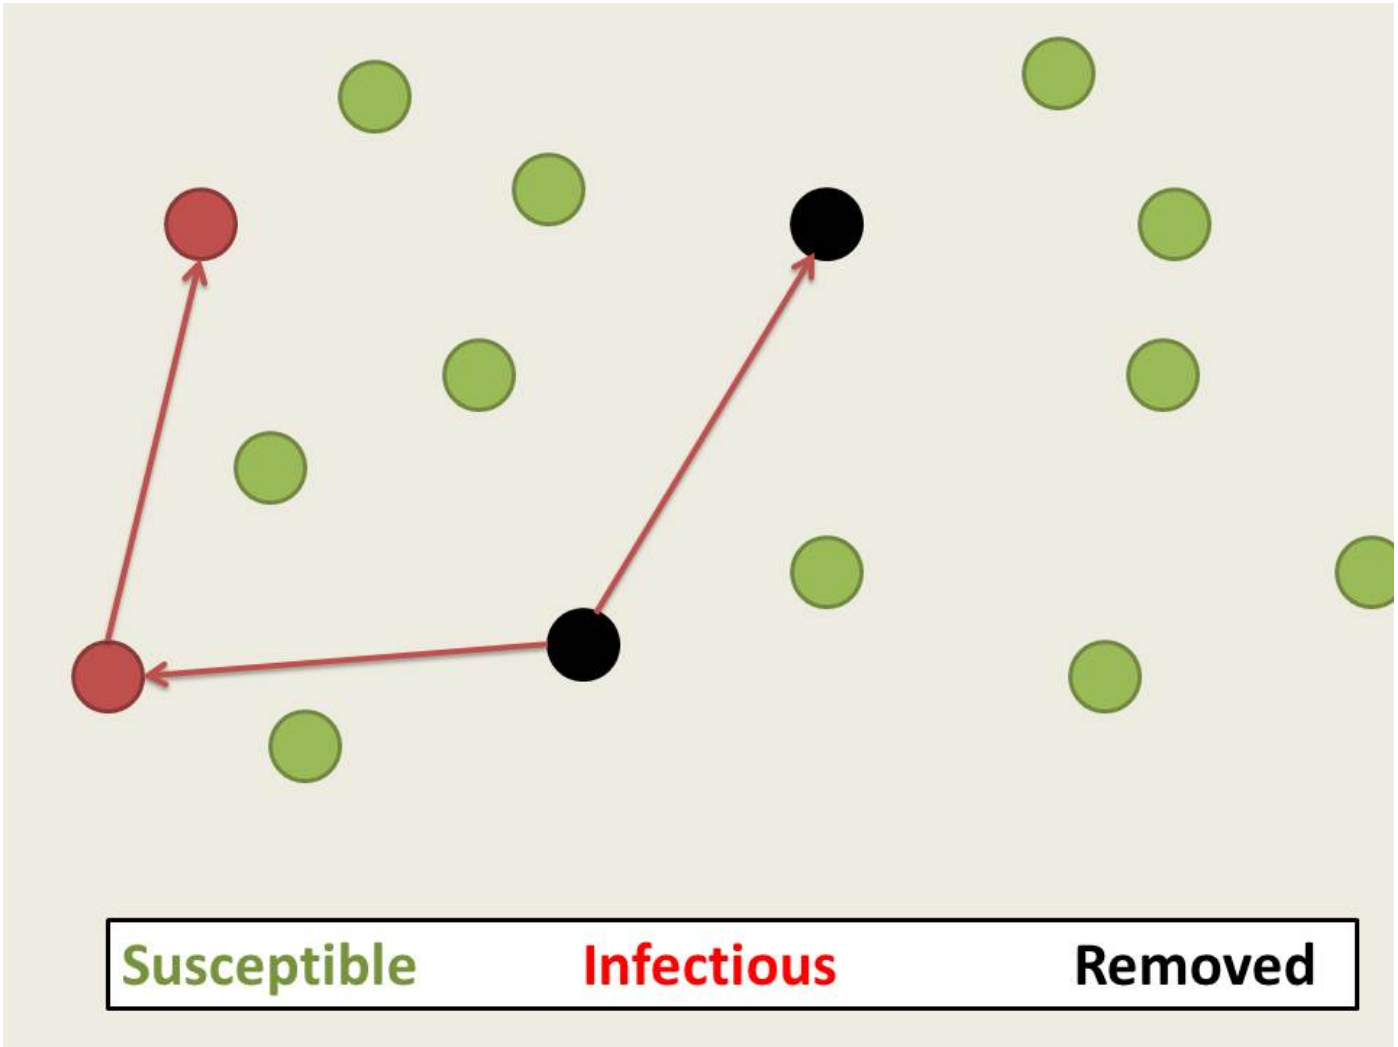

A hypothetical outbreak

A new infection occurs.

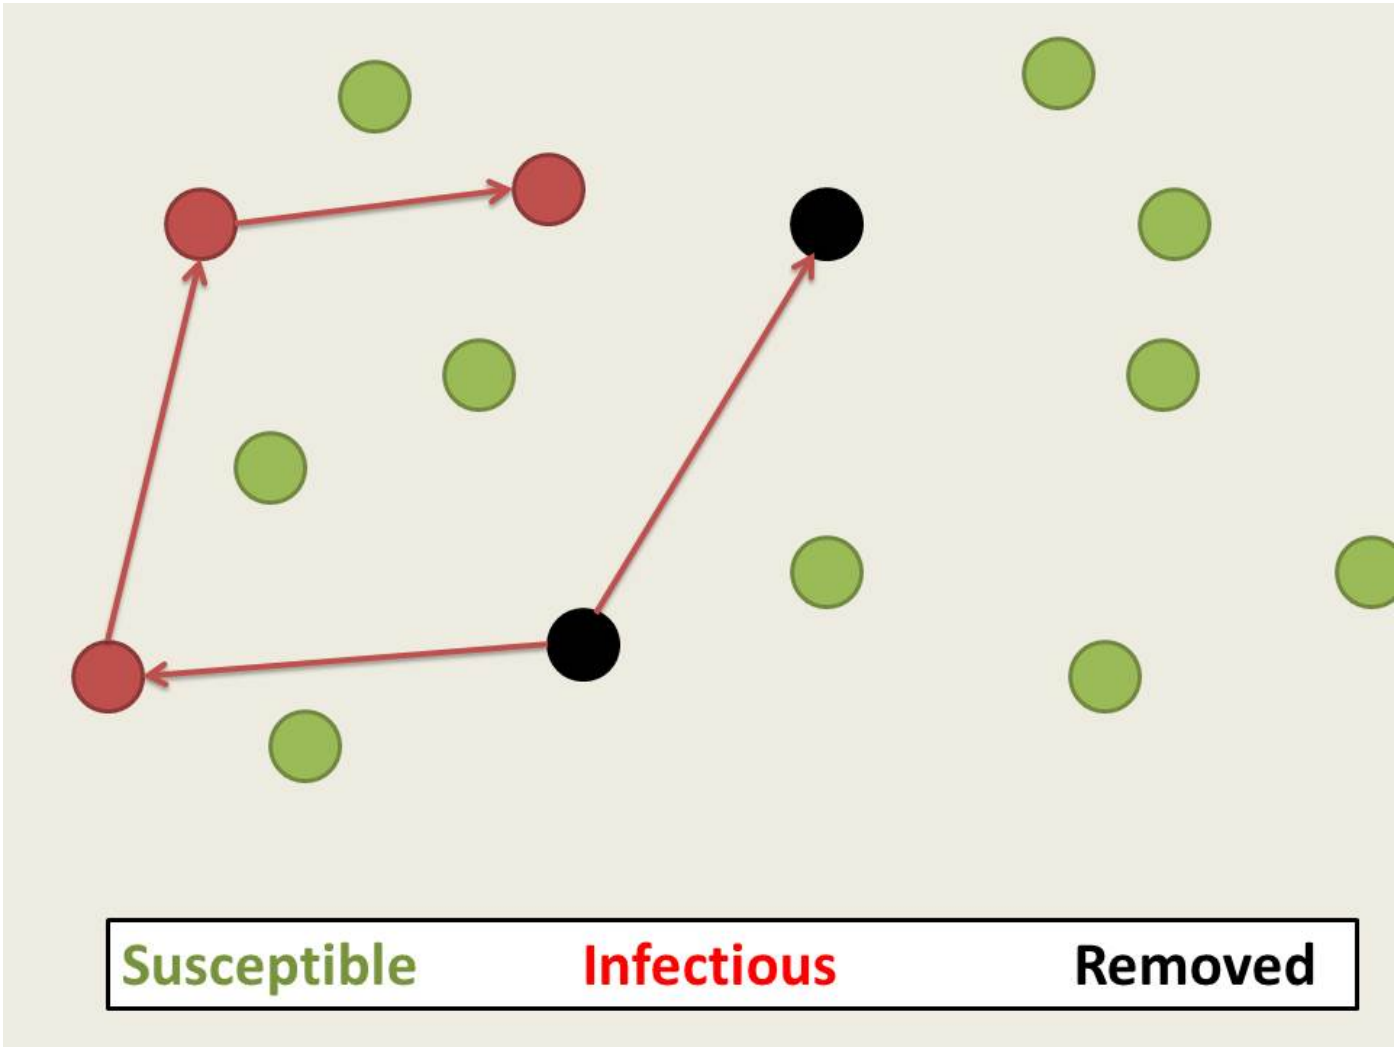

A hypothetical outbreak

An infectious host is removed.

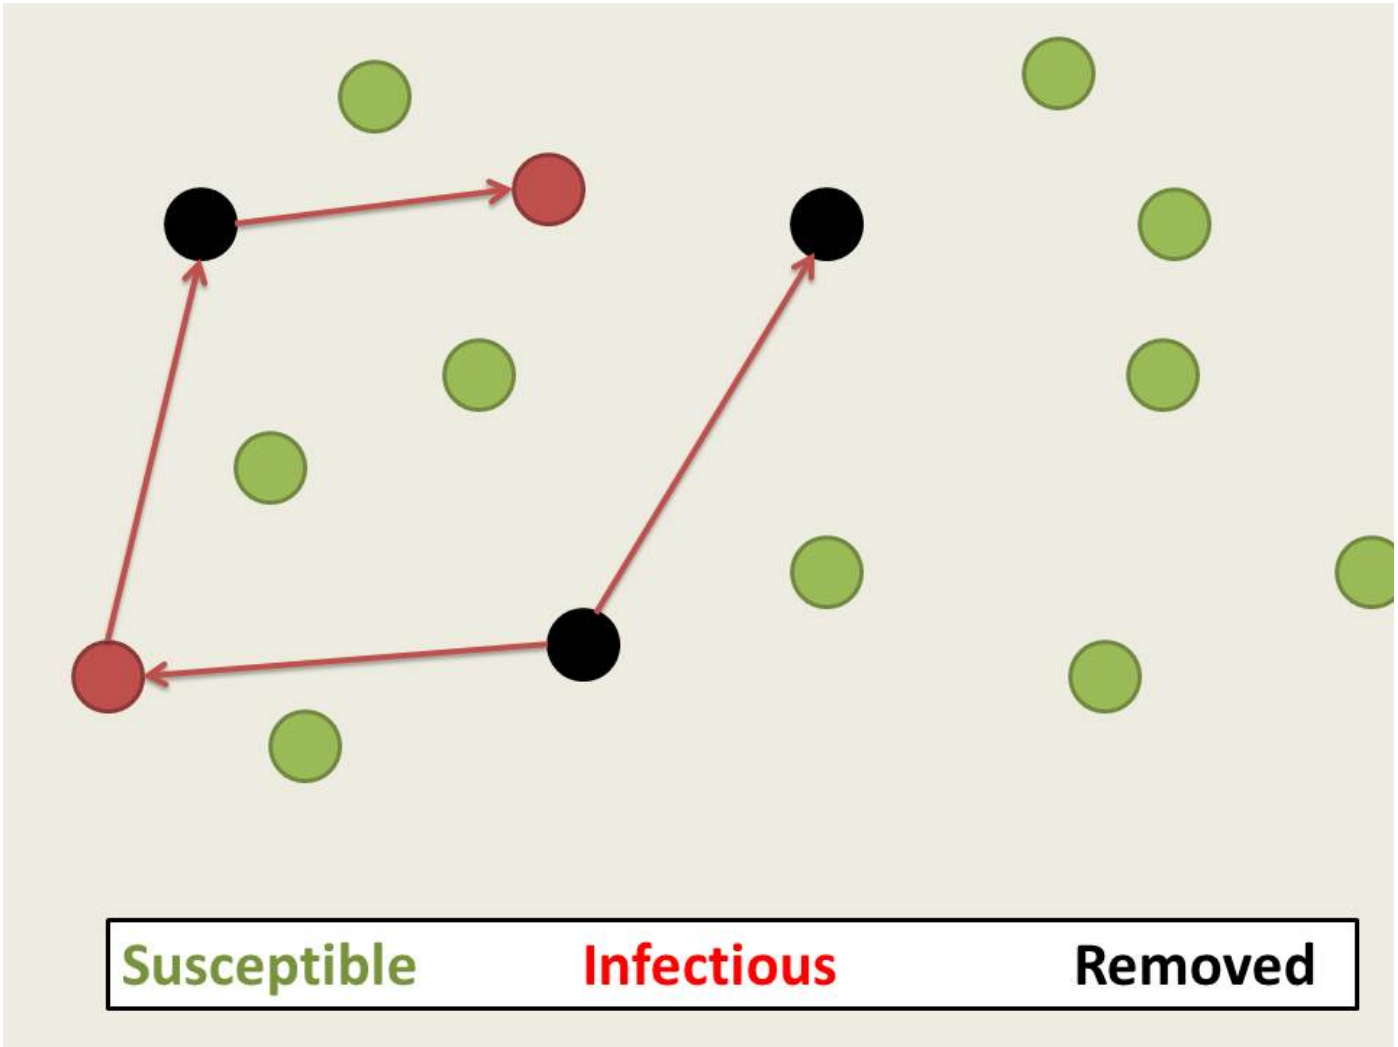

A hypothetical outbreak

An infectious host is removed.

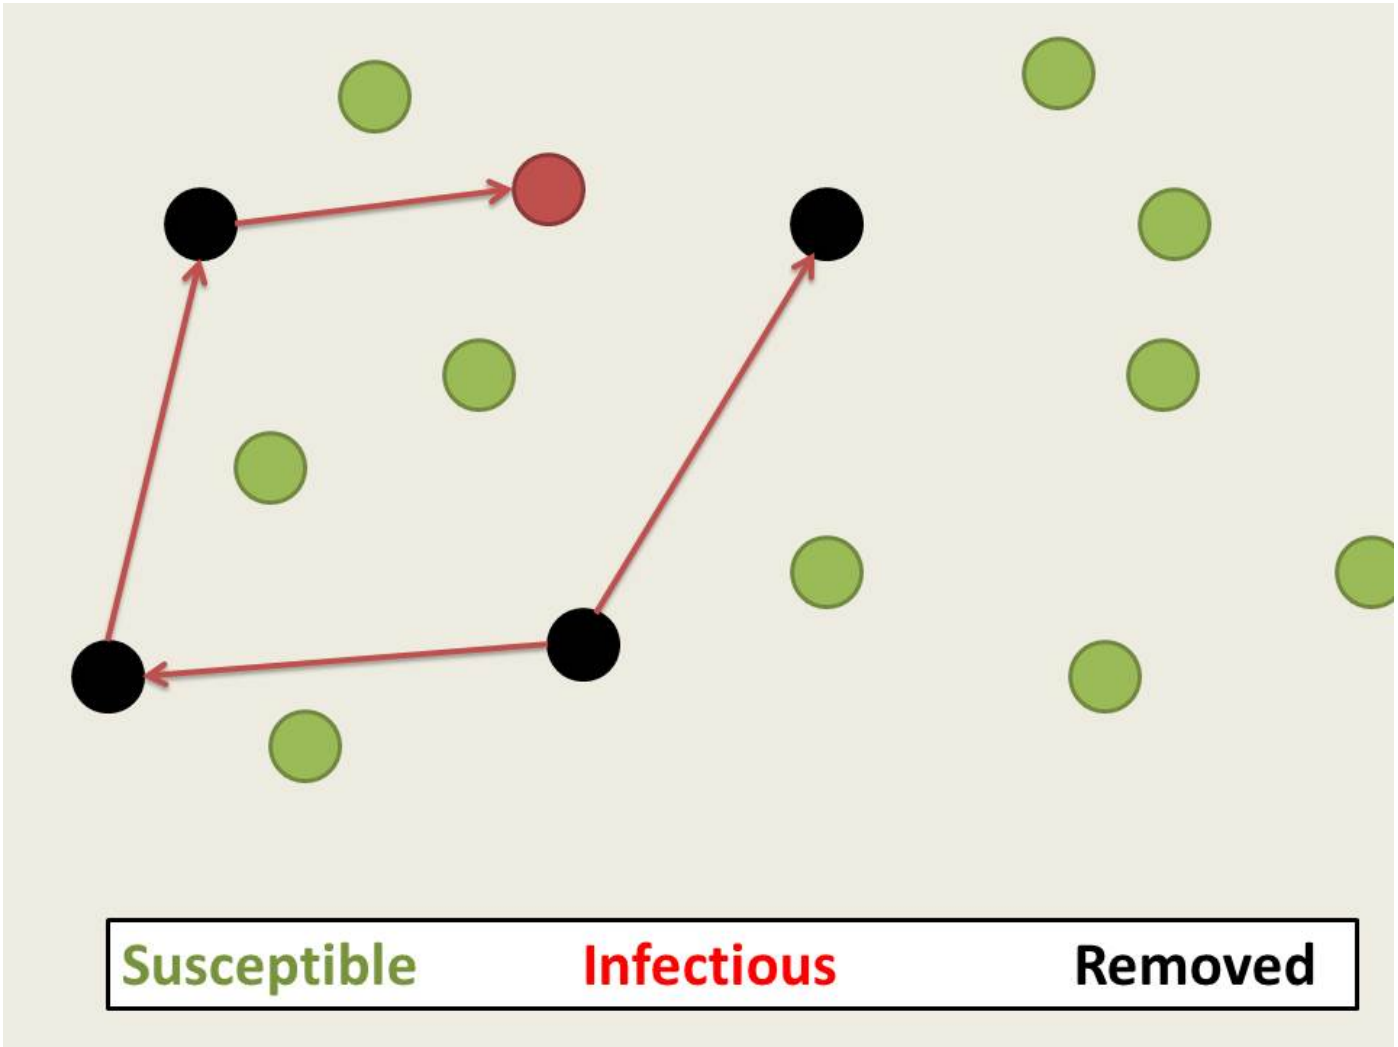

A hypothetical outbreak

A new infection occurs.

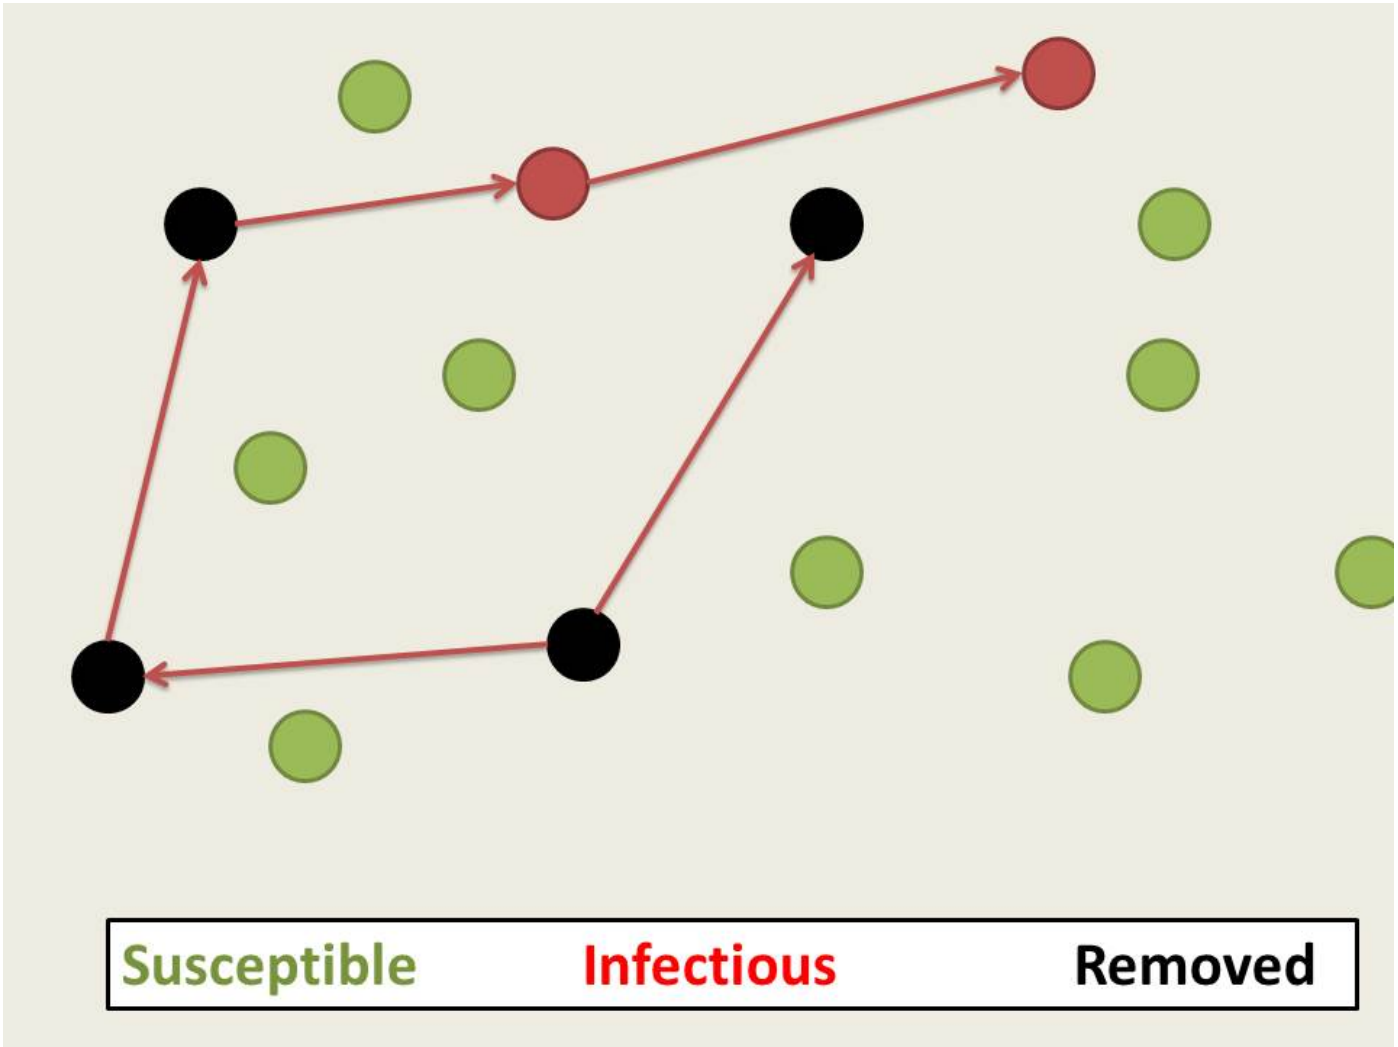

A hypothetical outbreak

An infectious host is removed.

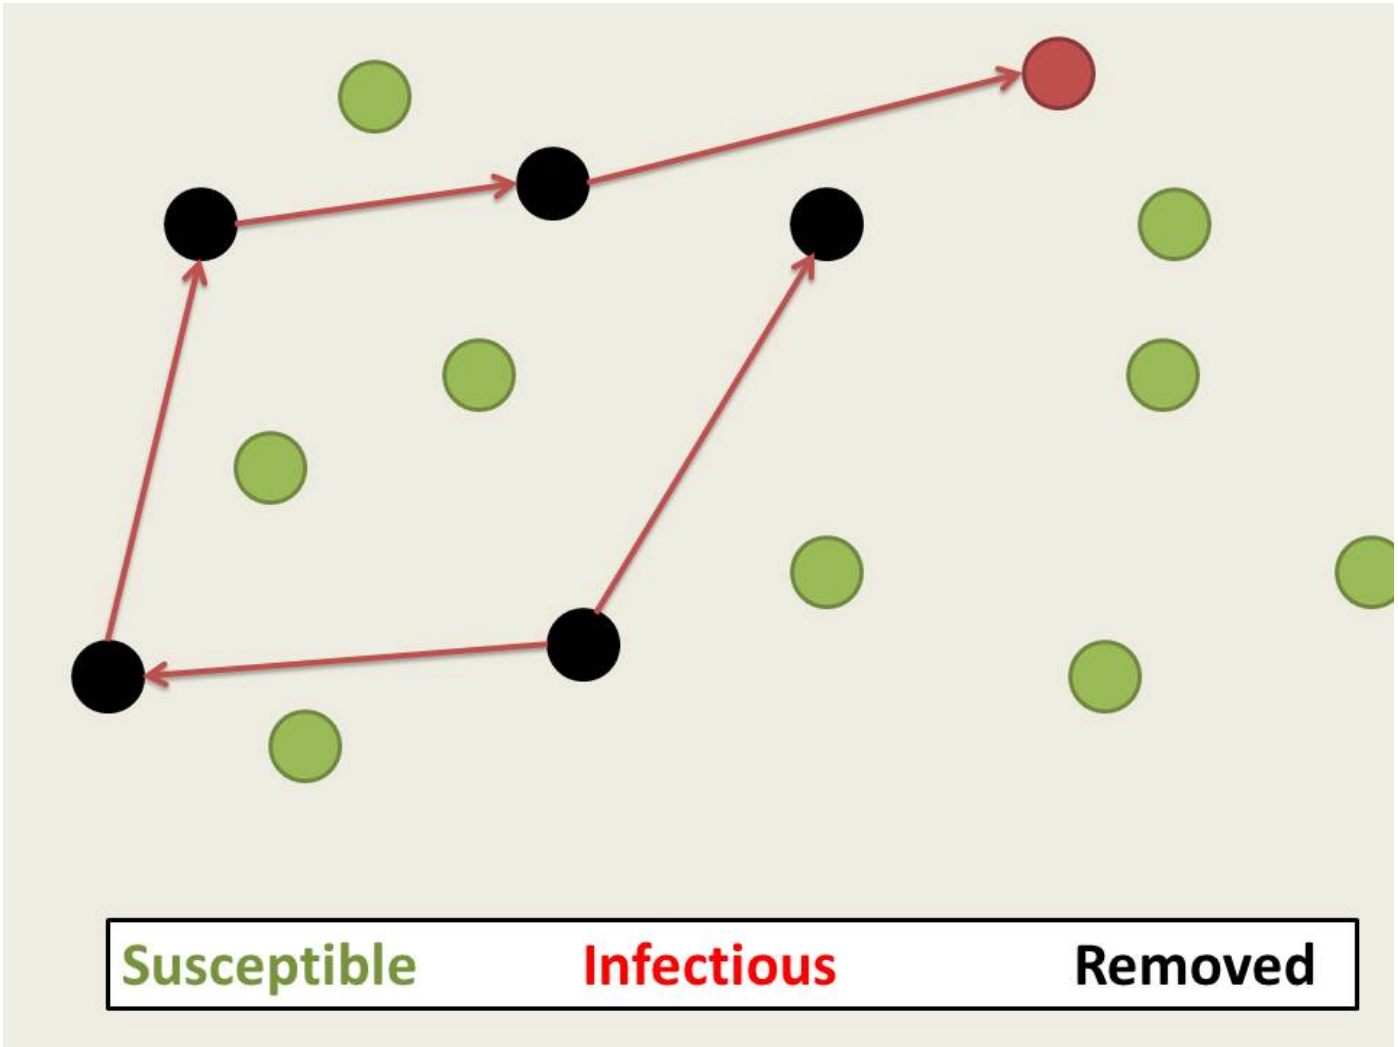

A hypothetical outbreak

An infectious host is removed. The infection has died out.

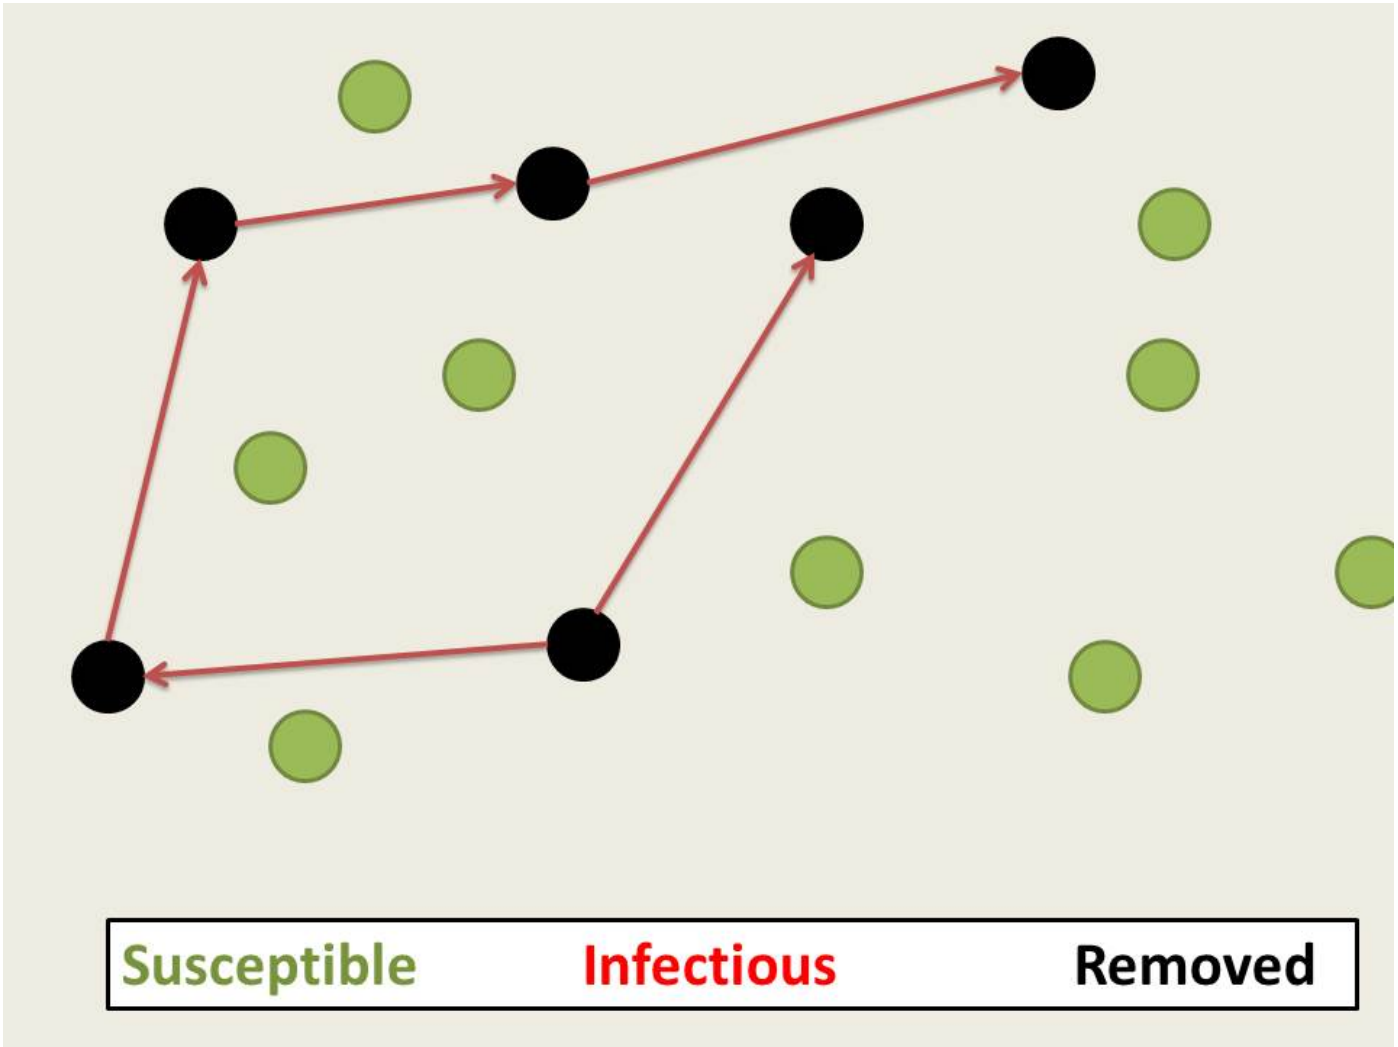

Some questions that we are trying to answer

- If one **index case** is introduced into an entirely susceptible population, will a **major outbreak** result? That is, will a significant fraction of the population eventually get infected?
- If a major outbreak does occur, what proportion of hosts will experience infection? This proportion is called the **final size** (of the outbreak).
- Which **control measures** are most effective in either preventing a major outbreak or reducing its final size?

Was this a major outbreak?

The infection has died out.

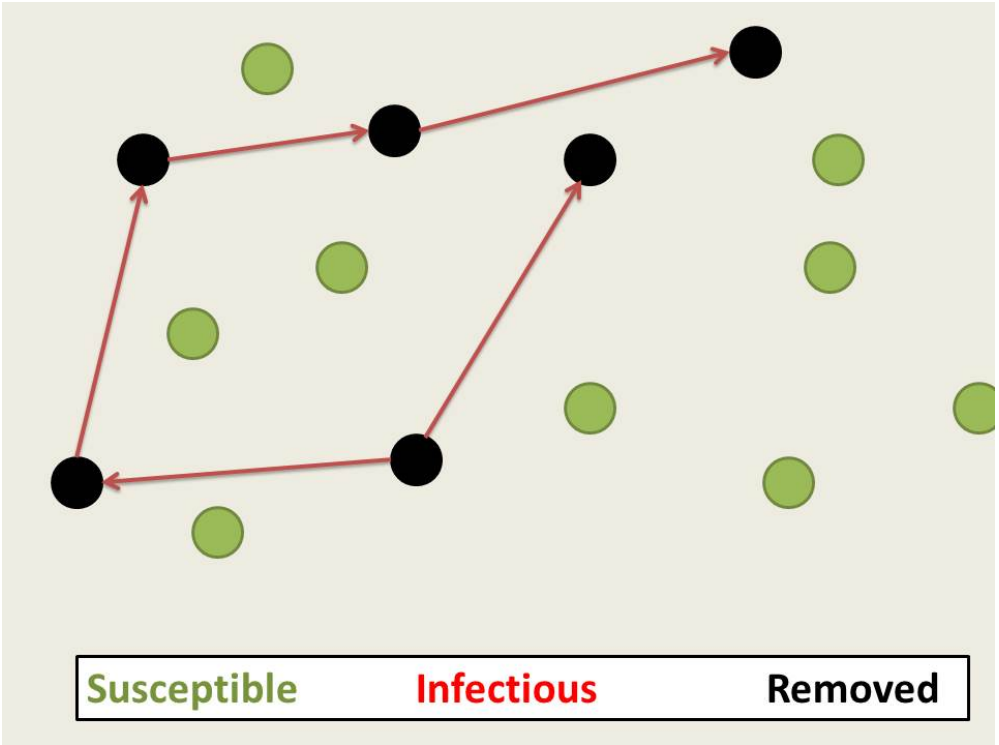

What is the final size of the outbreak that we just modeled?

Was this a major outbreak?

Was this a major outbreak?

Daniel: 6 hosts experienced infection.

Olena: Out of a population of 15. So the final size is $\frac{6}{15} = 0.4$. That is a significant fraction!

Ana: 40% of the population experienced infection. Definitely a major outbreak.

Daniel: But how do you know that the picture shows the entire population? Note that one circle is cut off.

WJ: Right, we don't know. What if the picture just shows some of the student population of the University of Aveiro?

Olena: Then the final size would be $\frac{6}{12300} \approx 0.0005$. Much less, but if the disease were very nasty, I would still be worried about it.

Ana: On a scale of the entire Portugal though, the final size would be approximately one in half a million. Definitely a minor outbreak.

Major *vs.* minor outbreaks?

Olena: So what at what fraction do you draw the dividing line between major and minor outbreaks?

Ana: It might depend on the severity of the disease and the overall impact. But one in half a million would definitely be minor, I think.

Daniel: It doesn't make sense to draw the dividing line at a fixed percentage.

WJ: You are right. In mathematics we draw it differently.

Daniel: How would you do that?

Major *vs.* minor outbreaks?

We make certain assumptions on how the disease spreads. One such possible assumption, that it would be immunizing and would conform to an *SIR*-model, we have already discussed.

We then try to estimate certain **parameters** of our model from data and to derive predictions about the final size $F(N)$ if a disease with such parameters spreads in a population that consists of N hosts.

When the resulting models show that $\lim_{N\to\infty} F(N) = 0$, then we say that the models **predict minor outbreaks**.

When the resulting models show that $\lim_{N\to\infty} F(N) = F > 0$, then we say that the models **predict major outbreaks**.

Ana: Cool. So how do control measures enter the picture?

So how do control measures enter the picture?

WJ: To study the effect of control measures, we consider a model without control, that will give us $\lim_{N\to\infty} F(N) = F > 0$.

Then we study models with, say two different control measures C_1 and C_2, for which we may get $\lim_{N\to\infty} F_1(N) = F_1$ and $\lim_{N\to\infty} F_2(N) = F_2$, respectively.

Olena: And if $F_1 = 0$ but $F_2 > 0$, then you would recommend control measure C_2, so that no major outbreaks will occur!

Daniel: But what if both $F_1, F_2 > 0$?

WJ: Whenever $F_1 < F_2, F$, then we would consider control measure C_1 **more effective.**

It may not always be possible to find a control measure that will prevent all major outbreaks.

What if a control measure is too expensive?

Ana: But what if the more effective control measure is too expensive to implement?

Daniel: How can you think about saving money when people's life may be at stake?

Ana: Not all infectious diseases are life-threatening. And the money you save on achieving a small but very costly reduction of final size in a mild disease could perhaps be put to better use in preventing more cases of a more dangerous disease.

WJ: Good point, Ana!
Even if very effective control measures exist in theory, it may not always be practically feasible to implement them.

Olena: This discussion is all getting very abstract. Can you show us a concrete example of a model and of its parameters?

An ODE-based *SIR*-model

We need:

- **Variables:** s, i, r They will represent the proportions of hosts in the **S**-, **I**, and **R**-compartments, respectively.
- **A state space:** Each of the variables can take values in $[0, 1]$ and they must add up to 1.
- **Some parameters:** α, β. Positive constants whose values need to be empirically determined.
- **A rule of change:** This governs the change of the variables over time.

$$\frac{ds}{dt} = -\beta\, i\, s,$$
$$\frac{di}{dt} = \beta\, i\, s - \alpha\, i,$$
$$\frac{dr}{dt} = \alpha\, i.$$

The basic reproductive ratio R_0

$$\frac{ds}{dt} = -\beta\, i\, s,$$
$$\frac{di}{dt} = \beta\, i\, s - \alpha\, i,$$
$$\frac{dr}{dt} = \alpha\, i.$$

To make predictions about the final size, we really need only one parameter here, the **basic reproductive ratio** $R_0 = \frac{\beta}{\alpha}$.

When $R_0 \leq 1$, then $\frac{di}{dt}$ will always be negative whenever $i > 0$, so that the model predicts that all outbreaks will be minor.

When $R_0 > 1$, the model predicts major outbreaks.

For example, when $R_0 = 1.5$, then the model predicts a final size of 0.5828, so that 58.28% of the entire population will experience infection.

Does this make sense?

Daniel: This doesn't make any sense to me.

WJ: Good.

Daniel: Are you making fun of me?

WJ: No. Tell me why it doesn't make sense to you.

Daniel: Because, as you said, if all parameters are fixed, then the predicted percentage is also fixed. If it happens to be 58.28%, then the model would supposedly predict that out of a population of 10,000 hosts, always exactly 5,828 will experience infection.
This is nonsense!!!

WJ: And why, exactly, would this be nonsense?

Daniel: Stop making fun of me!!
Everybody in the audience will tell you why!

Stochastic processes

- Transmission of pathogens between two hosts is a random event.
- The spread of infections is a **stochastic process,** which means that random stuff happens at random times.
- The predictions of ODE and other deterministic models of disease transmission are meaningful only if interpreted as expected values for large populations.
- R_0 would in this context be interpreted as the mean number of secondary infections caused by introduction of a single infectious host into an entirely susceptible population.
- Due to random chance, even for very large values of R_0, some outbreaks may be minor. The final size predicted by the ODE model should be interpreted as the mean final size of major outbreaks.

Agent-based models

Stochastic process models of disease transmission are more realistic than ODE models.

Instead of studying a stochastic process model as a formal mathematical construct, we can use a version of it that is embodied in computer code and run simulations.

Here the hosts are represented as so-called **agents** in the code. Therefore these models are called **agent-based models.**

In agent-based models of type *SIR* we need:

- A population of N agents that represent hosts.
- A time line $t \in \mathbb{N}$ for discrete-time models or $t \in \mathbf{R}$ for continuous-time models.
- States. At any given time, an agent can be either in state S (susceptible), I (infectious), or R (removed).
- For discrete-time models:
  - A transmission probability b_{ij} that agent i, if currently infectious, will transmit a critical number of pathogens to agent j during the current time step.
  - A removal probability a_i that agent i, if currently infectious, will transition into state R by the next time step.
- For continuous-time models:
  - A transmission rate β_{ij} at which events of infections of agent j by agent i occur, if agent i is currently infectious and agent j is currently susceptible.
  - A removal rate α_i at which agent i, if currently infectious, will transition into state R.

A nice toy, but ...

(After some simulations of agent-based *SIR* models with the simulation tool IONTW were shown)

Olena: I like this toy!

But what you were showing us doesn't agree with your last slide.

WJ: So how is this slide different from what I have shown in the simulations?

Olena: Let's look at the discrete-time version. There you have separate parameters a_i for each agent and separate parameters b_{ij} for each pair of agents, but your interface allows the user to input only one parameter a and one parameter b.

Homogeneity of hosts

WJ: All models use simplifying assumptions. In the software IONTW, we assume **homogeneity of hosts,** so that there are single parameters a (for discrete time) or α (for continuous time) with $a_i = a$ or with $\alpha_i = \alpha$, respectively.

Ana: But wouldn't this assumption imply that all hosts recover equally quickly from the disease? This doesn't seem realistic. For example, older people might typically take longer to recover from an infection than younger ones.

WJ: You are absolutely right. Like most models in this area, ours assume that these differences in individual duration of infectiousness differenced will average out.

Olena: But do these differences average out?

WJ: Good question! Dr. Sofia Rodrigues, whom you might have heard of, asked me a similar one in a related context. Your question translates into lots of interesting open research problems.

The uniform mixing assumption

Daniel: But you also have only a single parameters b (for discrete time) or β (for continuous time) with $b_{ij} = b$ or with $\beta_{ij} = \beta$, respectively.

WJ: This is called the "uniform mixing assumption," which means that every host is equally likely to make contact with every other host. In the simulations I showed you I did make this assumption.

Daniel: This cannot be true!

WJ: It will be approximately true for populations of animals that move around a lot, encounter each other rarely, and have no social or territorial structure.

Daniel: But not for human populations! Think about your chances of catching a disease from your next-door neighbor or from Ronaldo!

Contact networks

WJ: I see your point. So how would you model the social and territorial structure of human contacts?

Daniel: We make contact mostly with the people in our contact network.

WJ: Yes ...

Daniel: So you need to assume that transmission between two hosts can occur only when these hosts are connected in the contact network.

WJ: This is exactly what network-based models of disease transmission do. We will still need only one disease transmission parameter b or β. Then we set $b_{ij} = b$ or $\beta_{ij} = \beta$ when i and j are connected in the contact network, and $b_{ij} = 0 = \beta_{ij}$.

The assumptions of homogeneity of hosts and of uniform mixing

Under the assumption of homogeneity of hosts there are single constants a, α such that $a_i = a$ and $\alpha_i = \alpha$ for all i.

Under the uniform mixing assumption there are single constants b, β such that $b_{ij} = b$ and $\beta_{ij} = \beta$ for all i, j with $i \neq j$.

The uniform mixing assumption will be approximately true for populations of animals that move around a lot, encounter each other rarely, and have no social or territorial structure.

In network-based models of disease transmission we also have single parameters b, β, but assume that

- $b_{ij} = b, \beta_{ij} = \beta$ when i and j are connected in the contact network,
- $b_{ij} = 0, \beta_{ij} = 0$ otherwise.

Contact networks as graphs

Daniel: OK. But how do you model the contact network?

WJ: We can model contact networks as a **graph.**

The **nodes** of the graph represent hosts, the **edges** connect any two hosts who have reasonable frequent contact of the kind that may lead to the transmission of the particular infection.

Giving a precise meaning to "reasonably frequent" here is a decision that the modeler has to make.

The **degree** $deg(i)$ **of a node** i is the number of edges that connect i with other nodes.

Olena: Can we think of two people connected by an edge as friends? And of the degree of a node as the number of friends of that person?

WJ: This will give you a good intuition about these concepts.

Contact networks as graphs

Olena: Do Facebook friends count?

Daniel: No way!

Ana: Unless you study the spread of computer viruses.

WJ: Or of gossip, for that matter. In general, the relevant contact network will be different for different kinds of infections.

The uniform mixing assumption can be modeled in this framework by assuming that the network is the **complete graph** where all possible edges between the N nodes are included.

Example: The Sisters of the Round Table

Consider the monastic order of the Sisters of the Round Table. The sisters spend most of their lives in their individual cells, where they devote themselves to prayer and meditation. The only time they have contact with each other is during meals that they take seated in a fixed order around a giant round table. Within this community, diseases can be transmitted only during mealtime.

The probability of transmission will decrease with the distance at the table. It may depend on the particular nature of the disease how far the infectious agents can travel. When constructing a network model, we need to make a decision on how we want to set the cutoff.

In IONTW, the parameter **d** in network-type option **Nearest-neighbor 1** represents our assumption of how far the pathogenes can travel.

Let us look at some graphs in IONTW

(Some graphs are created in the software, including one of a rectangular grid.

Olena: Beautiful!!

Daniel: But real human contact networks are more messy. Let's say, more . . . random.

WJ: If we could construct a graph that reflects the somewhat random way in which humans make new contacts, might this give us a realistic picture of some human contact networks?

Daniel: Good idea. We could add new nodes one by one and attach them randomly with an edge to a previous node.

Olena: But new contacts are not made entirely randomly!! People who are already very popular attract more new friends.

Daniel: "Randomly" doesn't mean "with the uniform distribution." We could attach them preferentially to hosts who have already many friends.

The preferential attachment model in IONTW

WJ: Or, as in mathematical terms, to "nodes with high degree."

What you have described is the famous "preferential attachment model" of network growth.

It is implemented in our software IONTW.

(An IONTW movie of how network grows is shown)

Olena: Cool!!!

WJ: Could this be a halfway realistic contact network?

Olena: Yes!!

Daniel: Well, let's say it might give us an important insight.

Random graphs

In reality, contact networks are more messy and less structured, even the Sisters of the Round Table will exchange a few kind words with their next-cell neighbors on their way to the table.

It is also difficult to collect enough reliable data on a contact network on an actual population. We can often estimate some network parameters like the mean degree of a node in the network, or, more generally, the **degree distribution.**

We can then assume that the actual network will be somewhat typical for networks that are drawn from a certain distribution and study the expected dynamics on such (instances of) **random networks.**

Important classes of random networks

- **Erdős-Rényi** random graphs are the most basic example. Here each possible edge is included with probability p, independently for all pairs of nodes i, j. Here p is the model parameter.
- By adding the edges of Erdős-Rényi random graphs to the edges of nearest-neighbor graphs, we obtain **small-world models**. Such networks might more realistically represent the contact network of the Sisters of the Round Table who also will engage in conversation on their way to and from the table.
- **Scale-free** random graphs have degree distributions that roughly follow a power law. Such networks can be created by the **preferential attachment model** of Barabasi and Albert.

I want to get this toy!

Ana: Where can I buy your IONTW toy?

Daniel: There are things money can't buy . . . .

You can download it for free from

`https://people.ohio.edu/just/IONTW/`

Or use the free web-based version at

`https://qubeshub.org/iontw`

Olena: And the best things in life are free!!