Dynamics of Disease Transmission Networks

Winfried Just
Department of Mathematics
Ohio University

Presented at the workshop
Teaching Discrete and Algebraic Mathematical Biology to Undergraduates
MBI, Columbus, OH, July 31, 2013

What processes are we modeling?

We are interested in diseases that are triggered when infectious agents such as viruses or bacteria (called microparasites) enter the organism of a host (human, animal, plant).

We are not interested here in the actual changes that the disease causes in the organism of the host, or how the infectious agents multiply within the host. We only care about how the disease spreads between hosts of a given population.

In this lecture we will focus on diseases whose transmission requires direct contact (of a certain type) between hosts, as opposed to diseases that require a third type of organisms, called vectors for transmission between hosts (mosquitoes in the case of malaria), or diseases where the infectious agents are taken up from the shared environment of the hosts (drinking water in the case of cholera).

Which questions are we trying to answer?

- If some disease agents are introduced into a population of hosts that have not previously been exposed to the disease, will an epidemic result? That is, should we expect that a significant fraction of hosts in the population will eventually get infected?
- If an epidemic does result, what proportion of hosts will be infected? The proportion of hosts that will be infected at some (not necessarily the same) time during the epidemic is called the final size (of the epidemic).
- What control measures are most effective in either preventing an epidemic or reducing the final size as much as possible?
- Possible control measures include vaccination, quarantine, culling (for animal and plant diseases), or behavior modifications (for human diseases).

When is a mathematical model good (enough)?

Our goal is to construct mathematical models that give us correct answers to the questions on the previous slide. For that, we need the model to:

- Give us answers in the first place. Thus the model needs to be simple enough to be tractable either by mathematical analysis or computer simulations.
- Be sufficiently realistic. It needs to take into account sufficiently many biological details that influence the dynamics so as to make reasonably correct predictions.
- Be based on data that we actually can collect.

This may be too much to ask for. In practice, we may not know whether a given model is sufficiently realistic. But it is sometimes possible to study the mathematical problem of how more detailed, or finer-grained models relate to simplified coarser-grained models.

What happens to a host during a disease?

Since we only aim at modeling the dynamics between hosts, we only make the following assumptions about the disease that gets transmitted to host number i at time T_E^i, which stands for the time of exposure:

- All (potential) hosts start out being susceptible to the disease (at times $t < T_E^i$).
- At all times $T_E^i \leq t < T_I^i$ the host will not (yet) be able to infect others.
- At all times t with $T_I^i \leq t < T_R^i \leq \infty$ the host will be infectious, that is, will transmit the disease with positive probability during contacts with susceptible hosts.
- At all times $t \geq T_R^i$ the host will neither be infectious nor susceptible.

We are not assuming that T_I^i marks the onset of symptoms of the disease. Neither does T_R^i always mark their cessation.

What happens at time T_R^i? What kind of diseases do not satisfy the above assumptions?

The building blocks of disease models: compartments

Let's summarize:

- T_E^i is time of exposure of host number i .
- T_I^i is time of onset of infectiousness of host number i .
- T_R^i is time of removal of host number i, which means the host number i either dies from the disease or acquires permanent immunity at time T_R^i.

This suggests a partition of host population at time t into up to four compartments:
S comprises all susceptible hosts (for which $T_E^i > t$),
E comprises all exposed hosts with $T_E^i \leq t < T_I^i$),
I comprises all infectious hosts with $T_I^i \leq t < T_R^i$, and
R comprises all removed hosts with $T_R^i \leq t$.

SEIR-models

If all four compartments are considered, we get an *SEIR*-model.

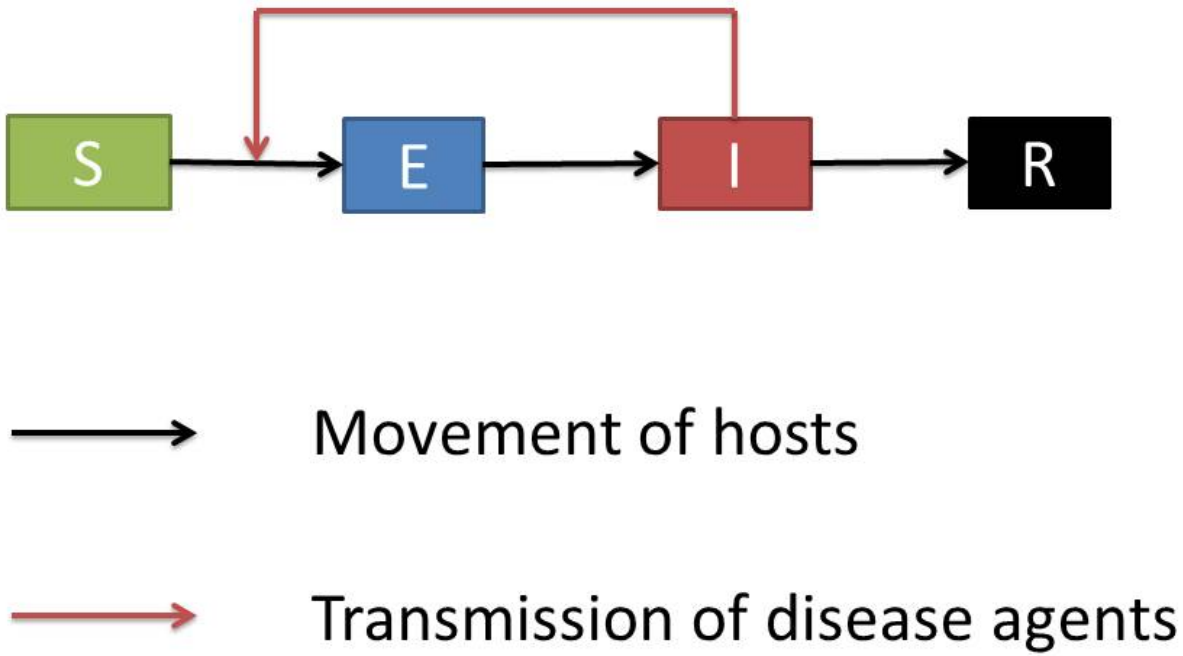

SIR-models

The simplifying assumption $T_E^i = T_I^i$ eliminates the E-compartment and we get an *SIR*-model.

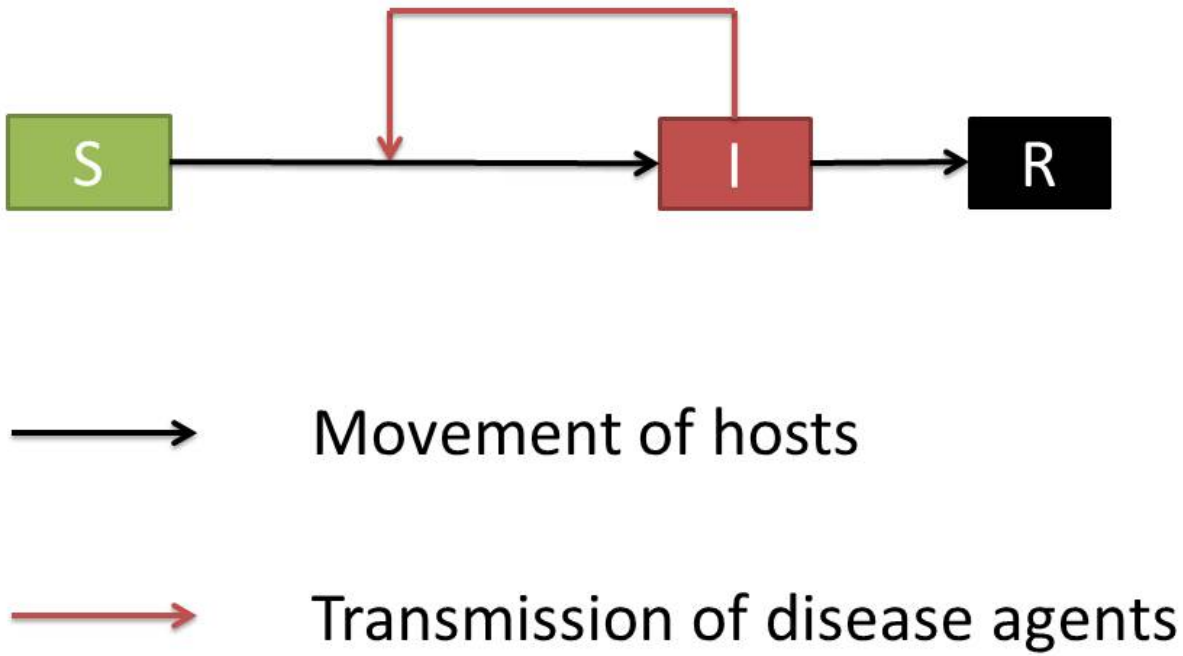

SI-models

If in addition $T_R^i = \infty$, then the R-compartment becomes redundant and we get an SI-model.

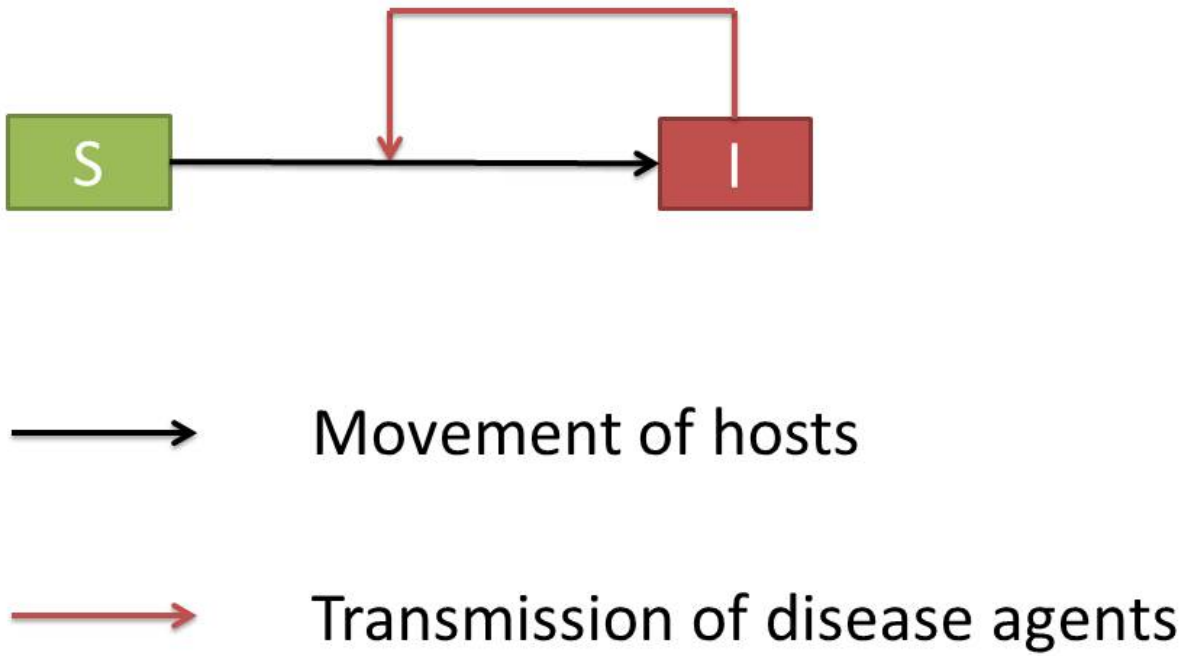

SIS-models

If we assume instead that at time T_R^i hosts recover and become susceptible to reinfection instead of acquiring immunity, then we get an *SIS*-model.

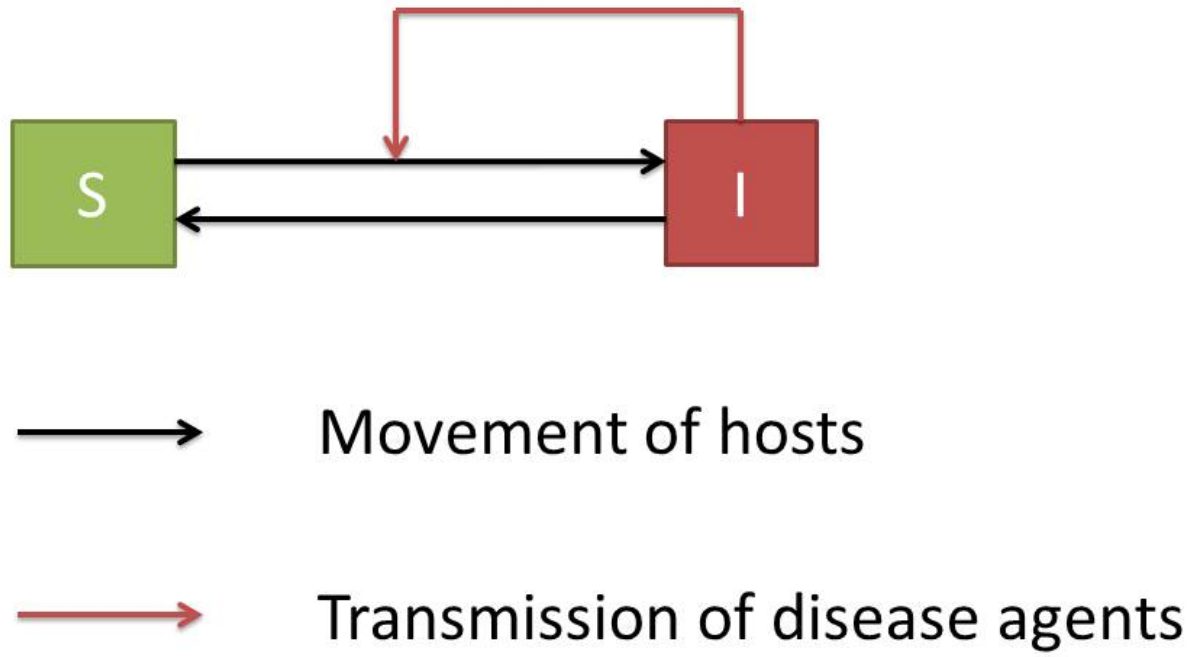

Four basic types of compartment models

S comprises all susceptible hosts (for which $T_E^i > t$),
E comprises all exposed hosts with $T_E^i \leq t < T_I^i$),
I comprises all infectious hosts with $T_I^i \leq t < T_R^i$, and
R comprises all removed hosts with $T_R^i \leq t$.
The times T_E^i, T_I^i, T_R^i are specific to each host. Membership in the compartments changes over time, and we can think of hosts moving (but not in space!) from S to E to I to R.

- The above is called an *SEIR*-model.
- The simplifying assumption $T_E^i = T_I^i$ eliminates the E-compartment and gives an *SIR*-model.
- If in addition $T_R^i = \infty$, then the R-compartment becomes redundant and we get an *SI*-model.
- If we assume instead that $T_E^i = T_I^i$ and at time T_R^i hosts simply recover and become susceptible to reinfection instead of acquiring immunity, then we get an *SIS*-model.

Can you think of other types of meaningful compartment models?

Are compartments good enough for our modeling?

Let us see whether we can translate our guiding questions into compartmentalese.
Assume an SIR-model. Consider a population of N individuals, all initially in S, and assume K of them become infected from outside sources at time T_E, where $K \ll N$. Now let the disease run its course, and consider
$F(K,N) = \lim_{t\to\infty} \frac{N - \#S(t)}{N}$.
This limit is the final size, *i.e.* fraction of hosts that eventually experience infection.

There is danger of an epidemic unless for fixed K we have $\lim_{N\to\infty} F(K,N) = 0$, which would mean that the disease will affect only a negligible fraction of a large population.

How would vaccination at time $T_v < T_E$ of a fraction r of hosts translate into compartmentalese?
As moving rN hosts into R at time T_v.
Compartmentalese seems to be a convenient language for us.

Compartment-based ODE models

- In the remainder of this talk we will ignore the E-compartment and assume that the onset of infectiousness coincides with the time of exposure, that is, $T_I^i = T_E^i$.
- We will moreover assume that the population size N is fixed, which ignores demographics, that is, births, deaths from unrelated causes, immigration and emigration.
- In the resulting ODE models the state of the population is represented by three variables S, I, and R. These variables could either represent the proportions of hosts in the respective compartments, or their numbers.
- In order to make the use of derivatives somewhat respectable, one can think of population size being expressed in units of a thousand or a million individuals so that at least some fractional values of the variables make sense.

Basic ODE versions of the *SI*-, *SIR*- and *SIS*-models

***SI*-model:**

$$\frac{dS}{dt} = -\beta SI$$
$$\frac{dI}{dt} = \beta SI$$

***SIS*-model:**

$$\frac{dS}{dt} = -\beta IS + \alpha I$$
$$\frac{dI}{dt} = \beta IS - \alpha I$$

***SIR*-model:**

$$\frac{dS}{dt} = -\beta IS$$
$$\frac{dI}{dt} = \beta IS - \alpha I$$
$$\frac{dR}{dt} = \alpha I$$

The rate of infection β may or may not depend on N; the removal rate α is always independent of N.

What do the ODE models predict?

These ODE models are easy to study analytically.

- If $S(0), I(0) > 0$, then the SI-model predicts that $\frac{dI}{dt} > 0$ at all times and $\lim_{t\to\infty} S(t) = 0$, so the whole population will eventually be infected.
- Since $S = N - I$, the SIS-model simplifies to the logistic growth model
  $\frac{dI}{dt} = \beta I(N - I) - \alpha I = \beta I(N - \frac{\alpha}{\beta} - I)$,
  which predicts both a disease-free equilibrium $I^* = 0$ and an endemic equilibrium $I^{**} = N - \frac{\alpha}{\beta}$.
- The SIR-model allows both for predicting whether or not an epidemic will occur and for predicting its final size if it does.

Would the final size be 1 if β is sufficiently large relative to α?

But wait a minute ...

In our definition of final size we assumed that K out of N individuals became initially infected and considered $F(K,N) = \lim_{t\to\infty} \frac{N - \#S(t)}{N}$.

This limit exists all right in each repetition of the "experiment," but disease transmission is inherently a stochastic process and the outcome will differ between repeated runs of the "experiment." To make the above definition of final size meaningful for a given compartmental model we need to treat both sides as expected values.

Even in this interpretation though, $F(K,N)$ will not in general be determined by K and N alone. Which K individuals are initially infected? Some socially withdrawn loners or highly gregarious ones with a lot of social interactions?

Advantages and disadvantages of compartment models

While compartment-based models, and ODE models in particular, often make fairly realistic predictions, they have certain advantages and certain disadvantages.

- ODE models are relatively easy to study.
- They involve few variables and require estimation of very few parameters.
- ODE models ignore the stochastic nature of disease transmission.
- Compartment models ignore heterogeneities between individual hosts.
- Compartment models are based on the often unrealistic assumption of uniform mixing between individual hosts.

We will present a different type of models that can alleviate these three disadvantages to some extent. But first let us illustrate the nature of these problems with some examples.

An example

Consider the spread of flu in a dorm with initially $K = 1$ student infected. If we ignore heterogeneities, then we can estimate the probability that an epidemic will result if the infected student is "average." Look at two scenarios:

Scenario 1: The infected student caught it in a bar.

Does the probability estimate based on the assumption of an "average" student appear to apply in this scenario, or does it appear to be too high or too low?

Scenario 2: The infected student caught it from the janitor.

Does the probability estimate based on the assumption of an "average" student appear to apply in this scenario, or does it appear to be too high or too low?

How would an epidemic get started anyway? Example 1

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

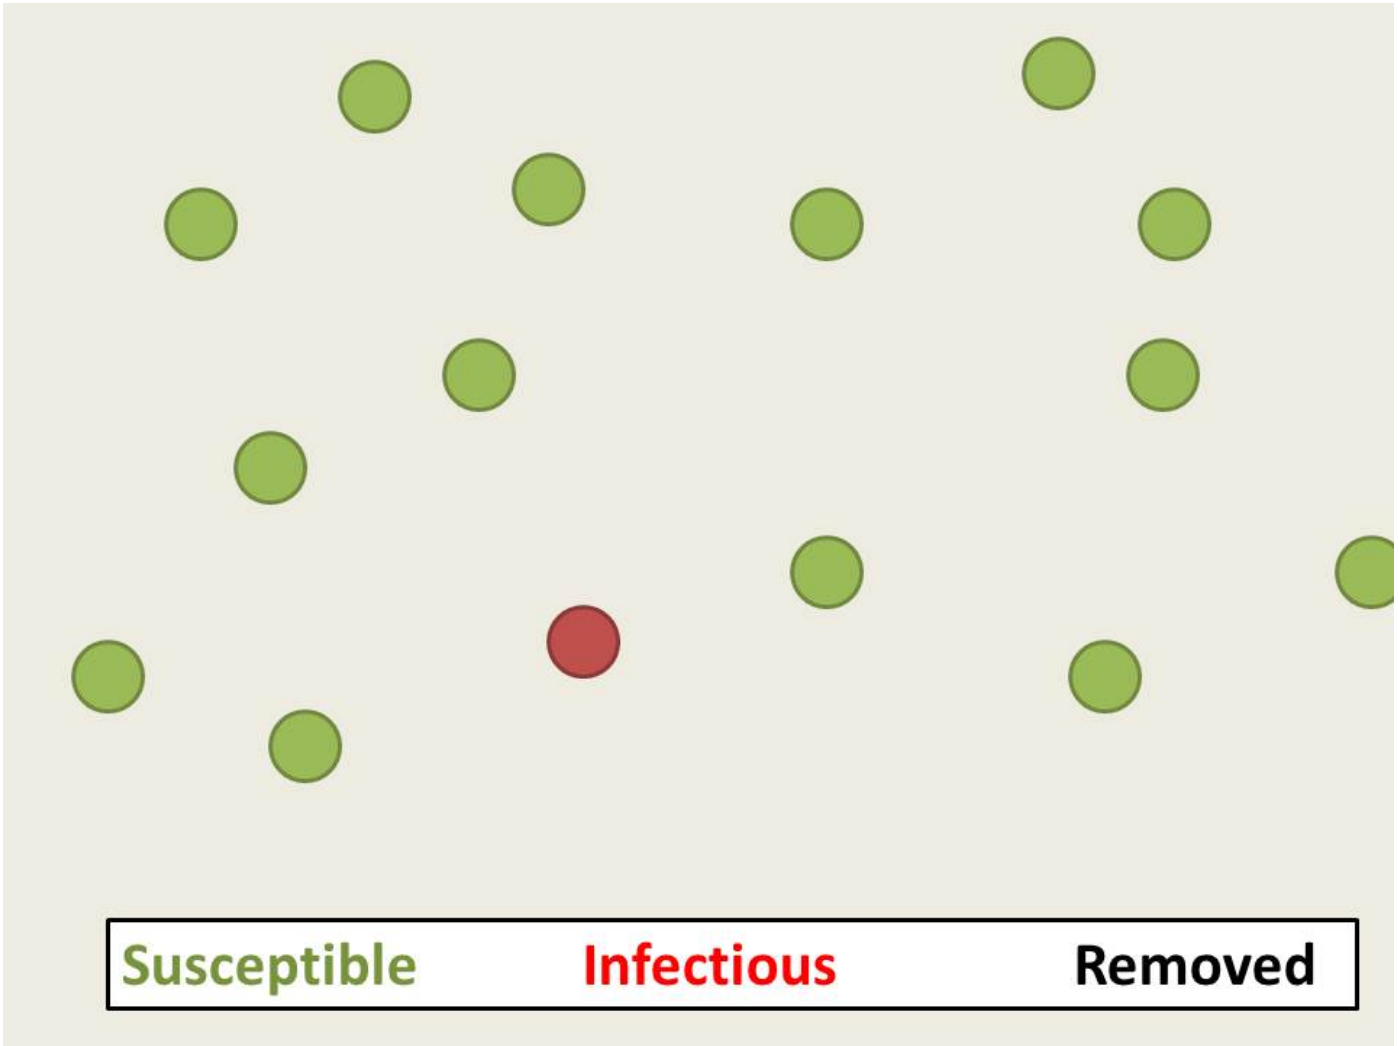

How would an epidemic get started anyway? Example 1

A new infection occurs.

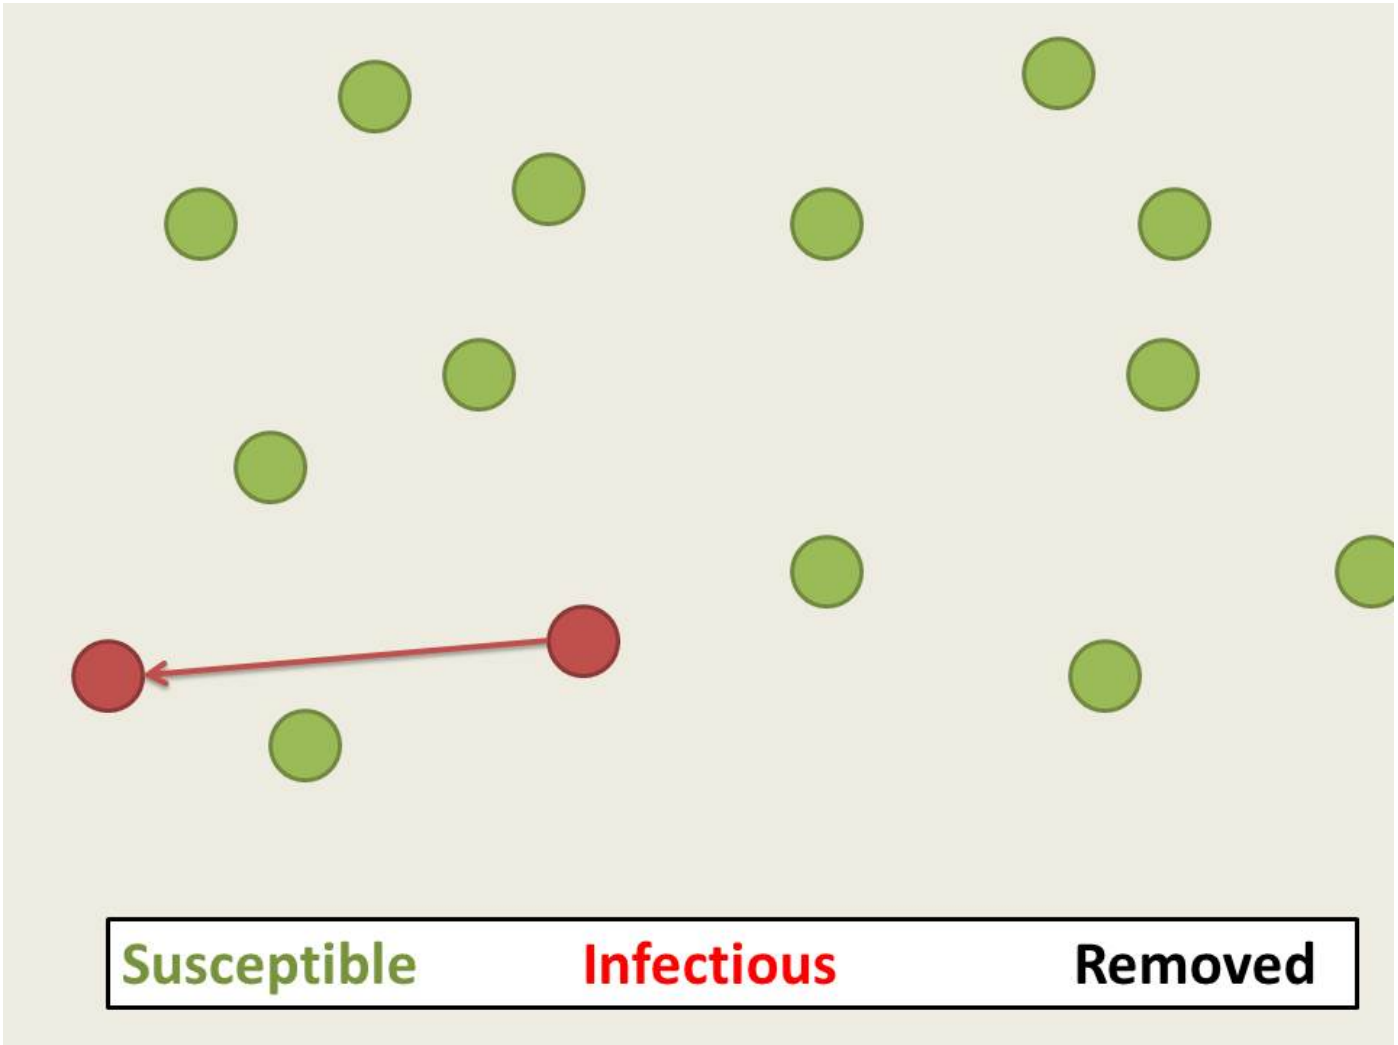

How would an epidemic get started anyway? Example 1

A new infection occurs.

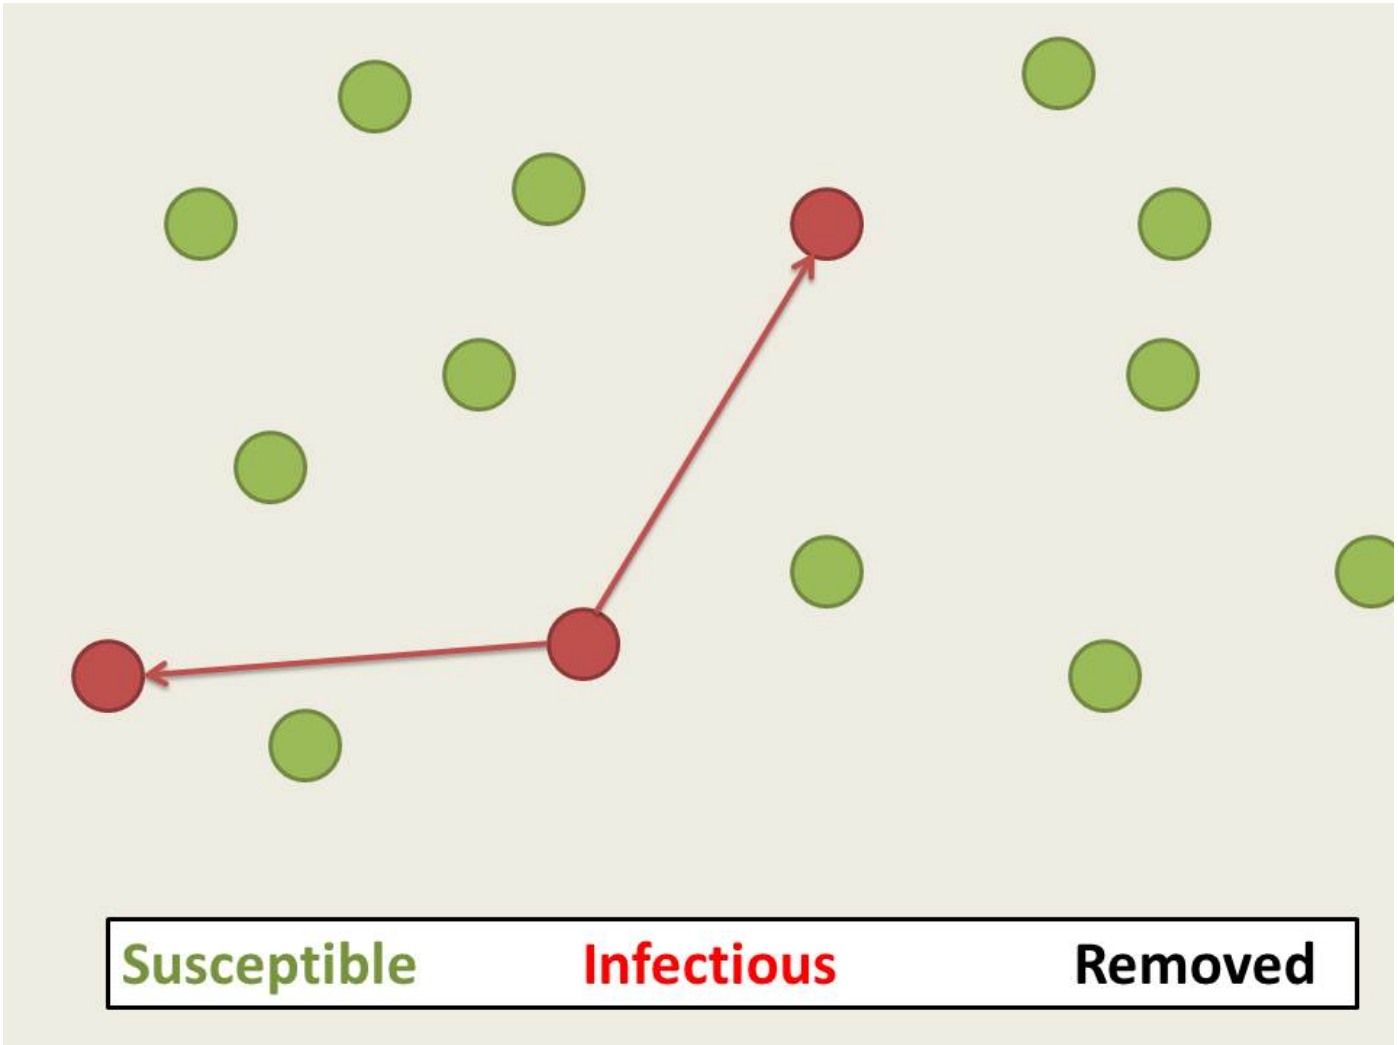

How would an epidemic get started anyway? Example 1

A new infection occurs.

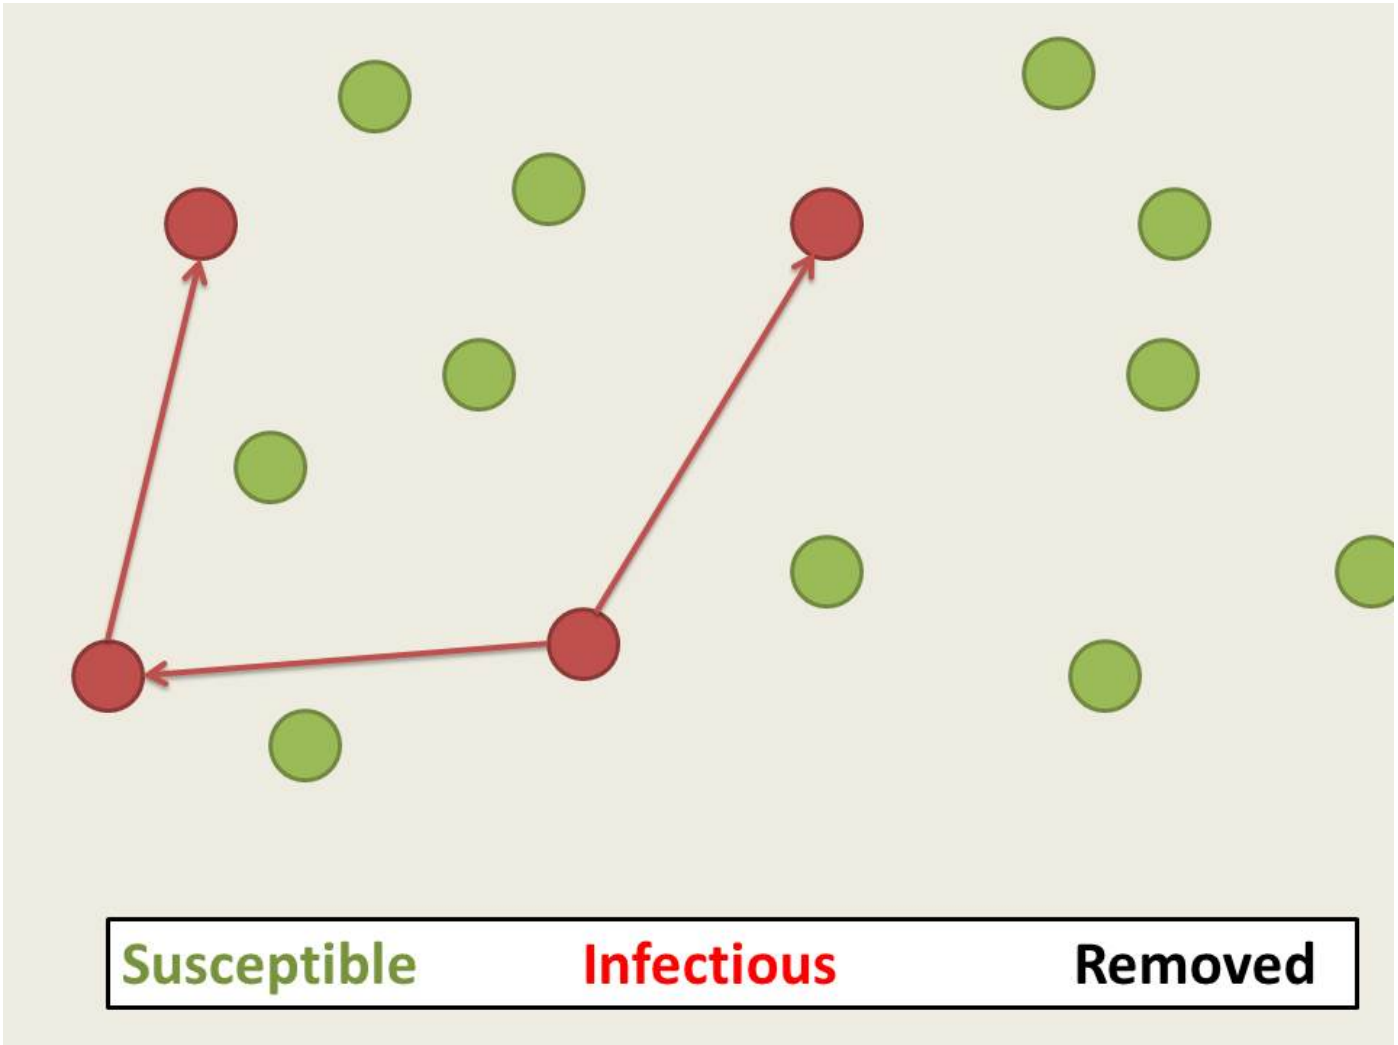

How would an epidemic get started anyway? Example 1

An infectious host is removed.

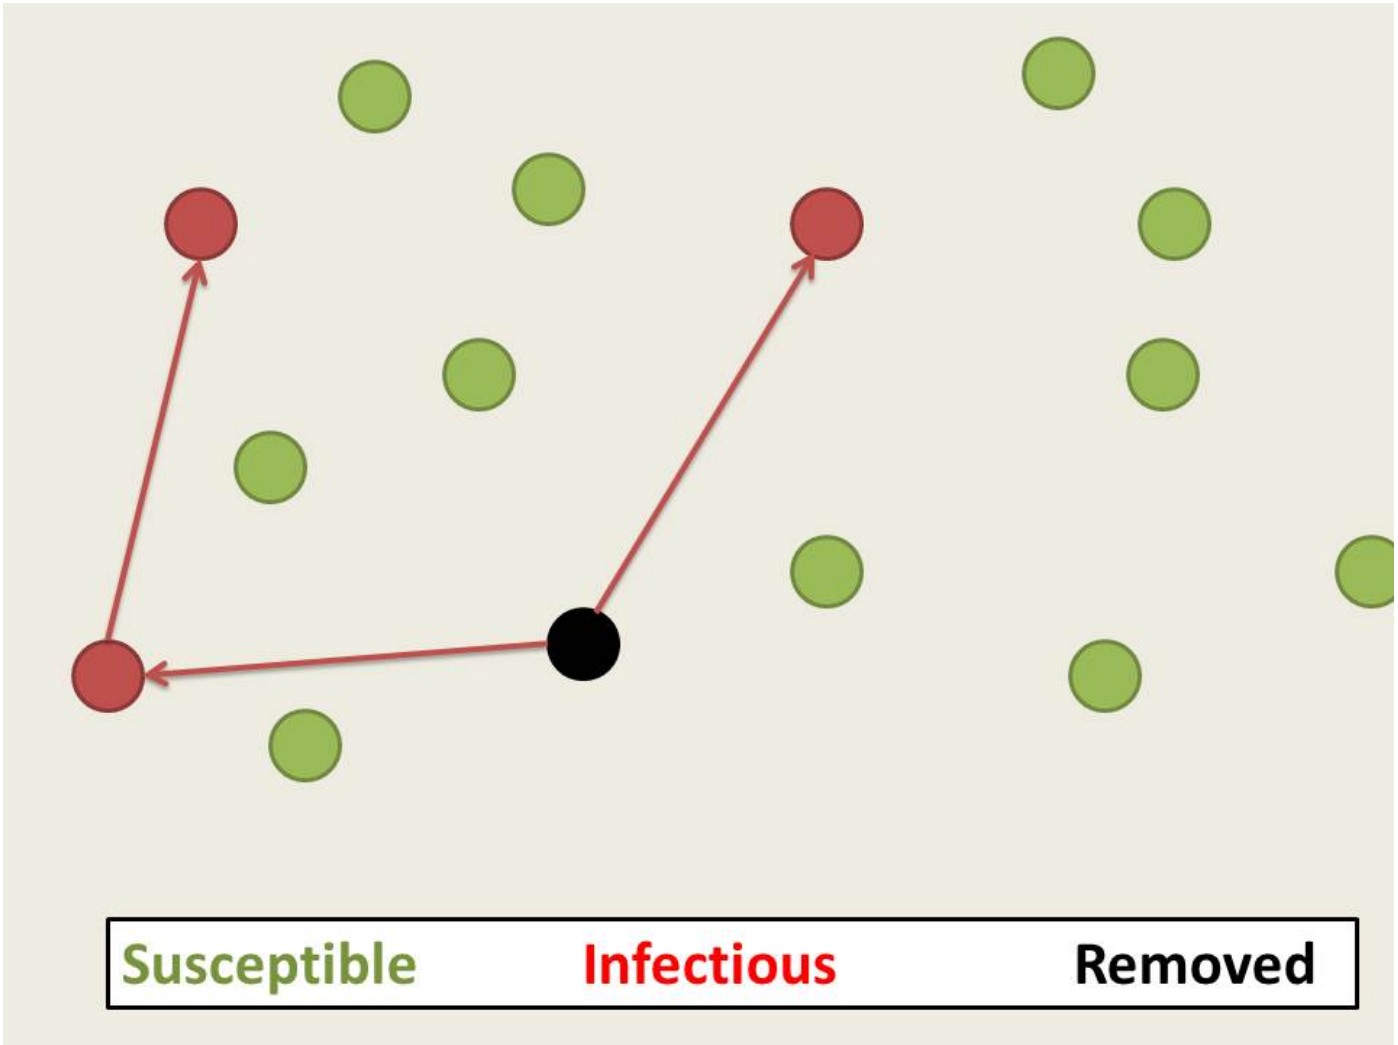

How would an epidemic get started anyway? Example 1

An infectious host is removed.

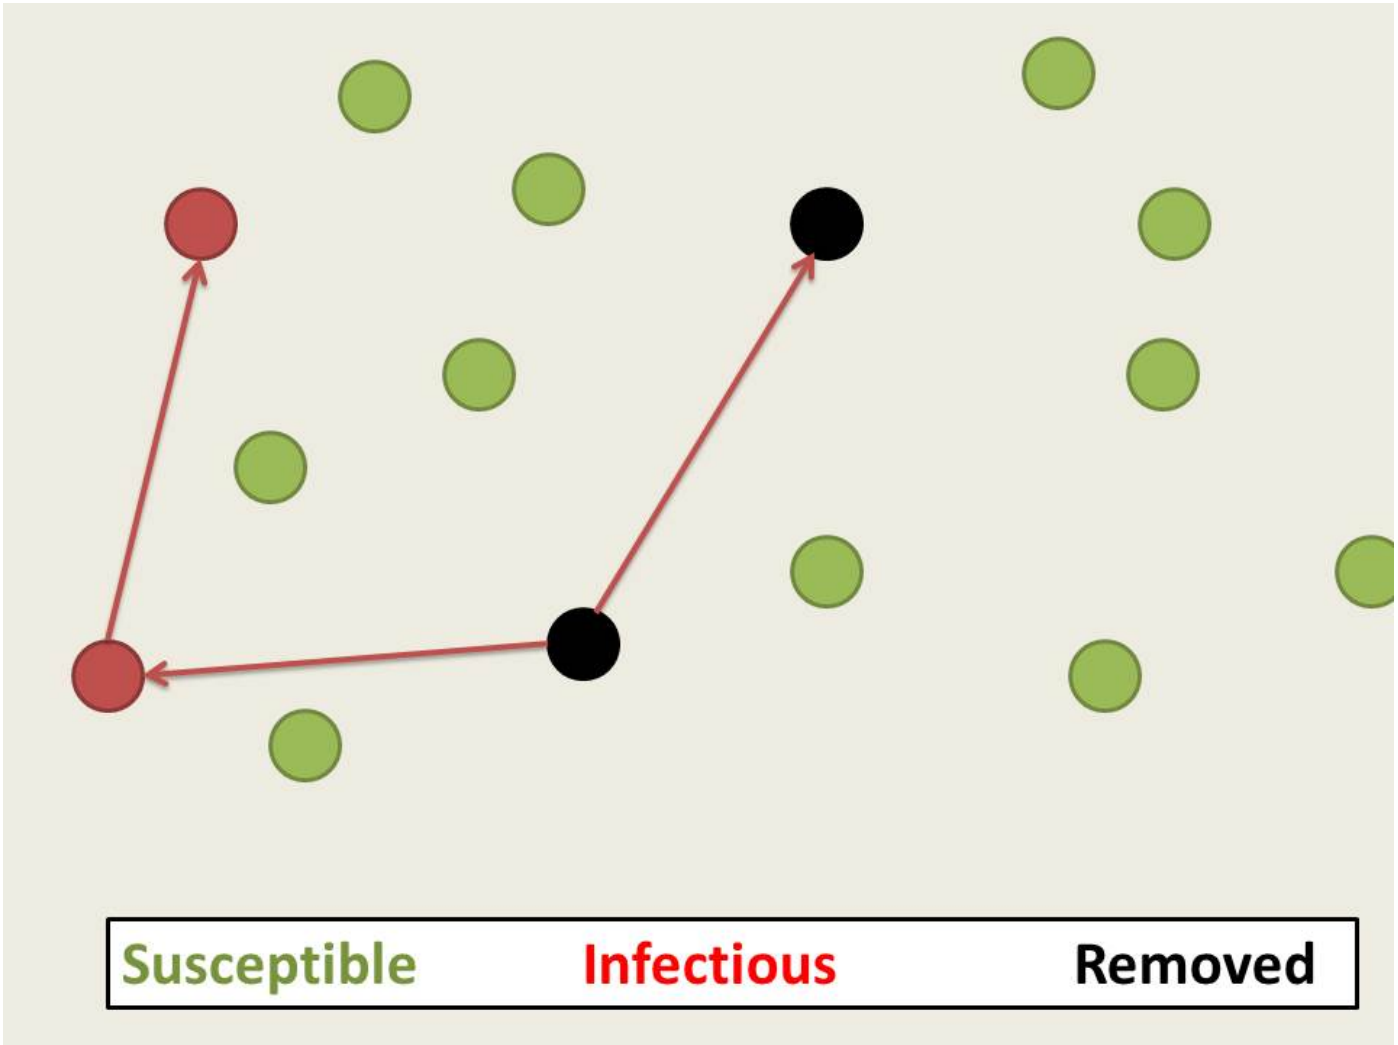

How would an epidemic get started anyway? Example 1

A new infection occurs.

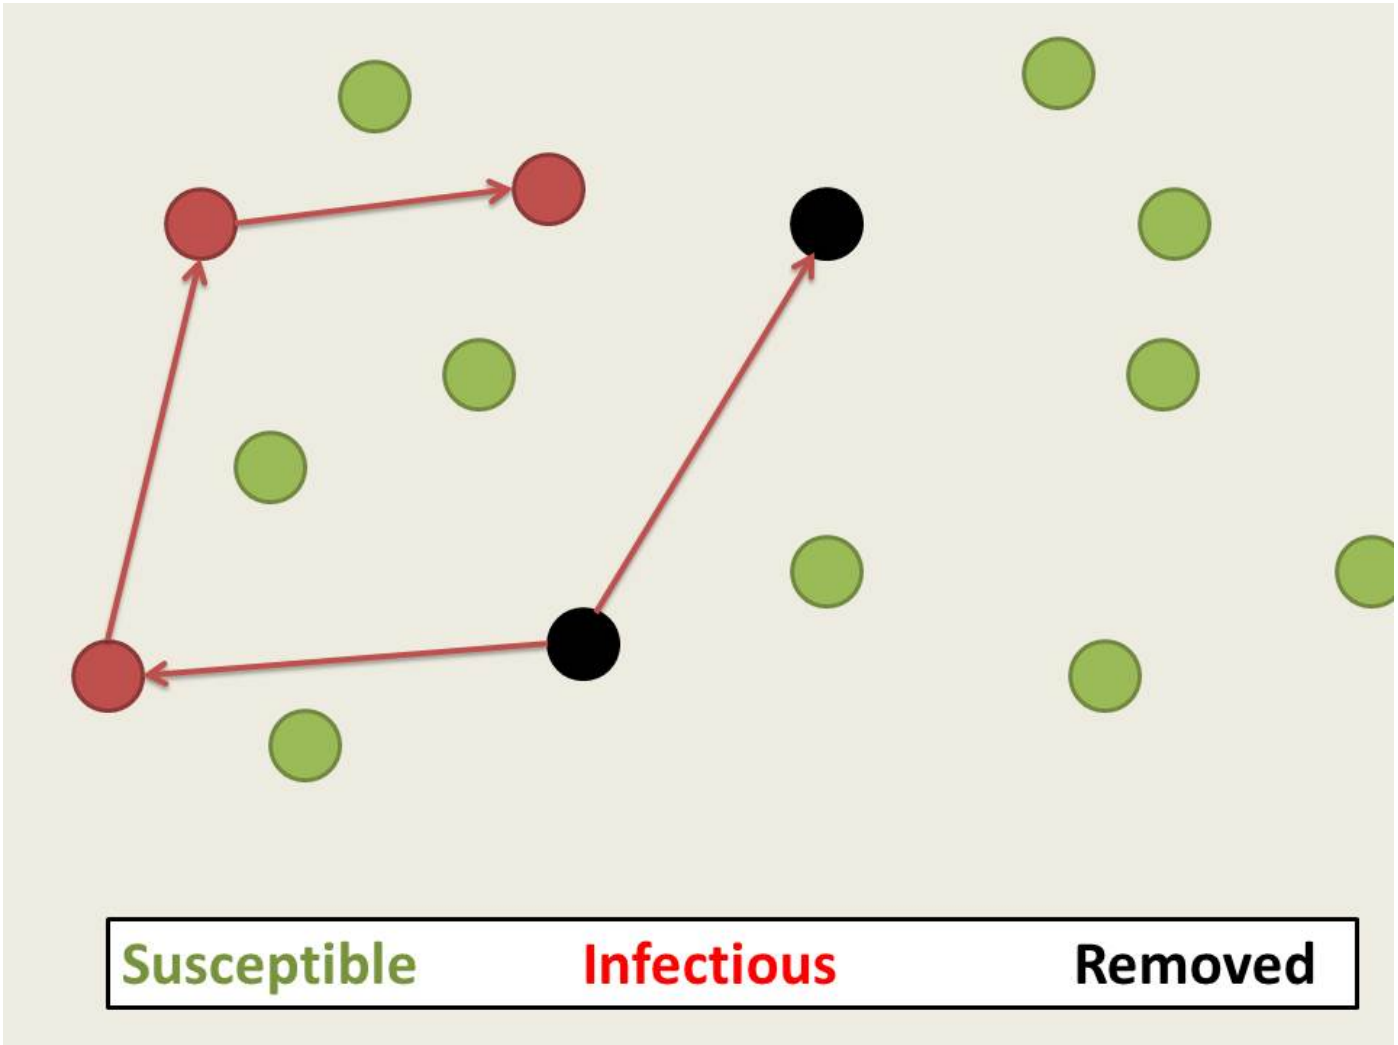

How would an epidemic get started anyway? Example 1

An infectious host is removed.

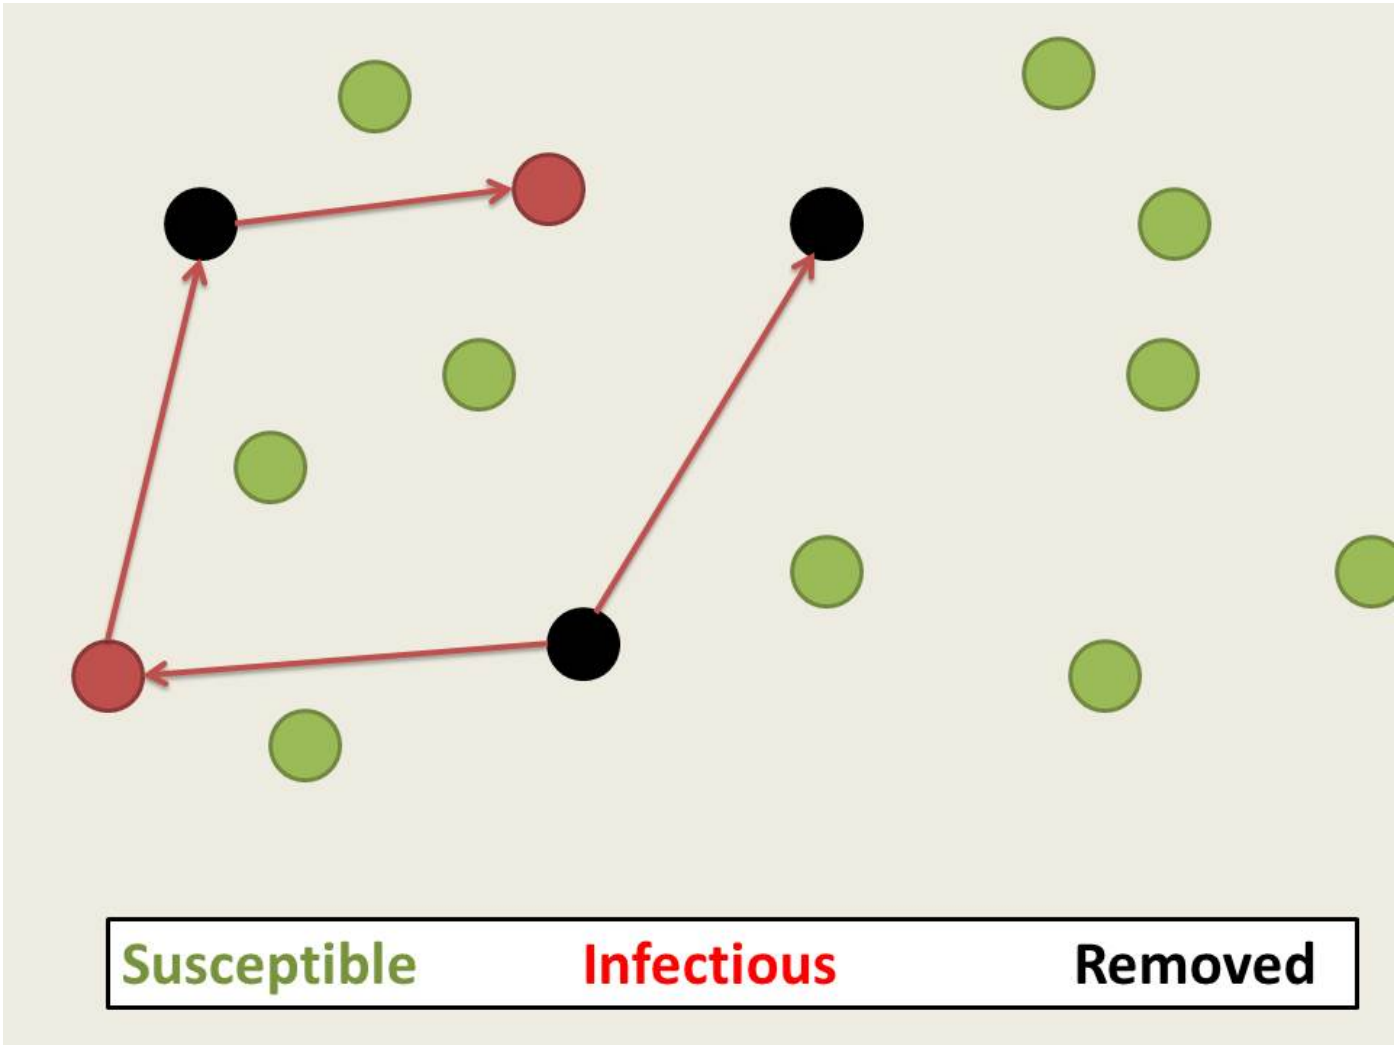

How would an epidemic get started anyway? Example 1

An infectious host is removed.

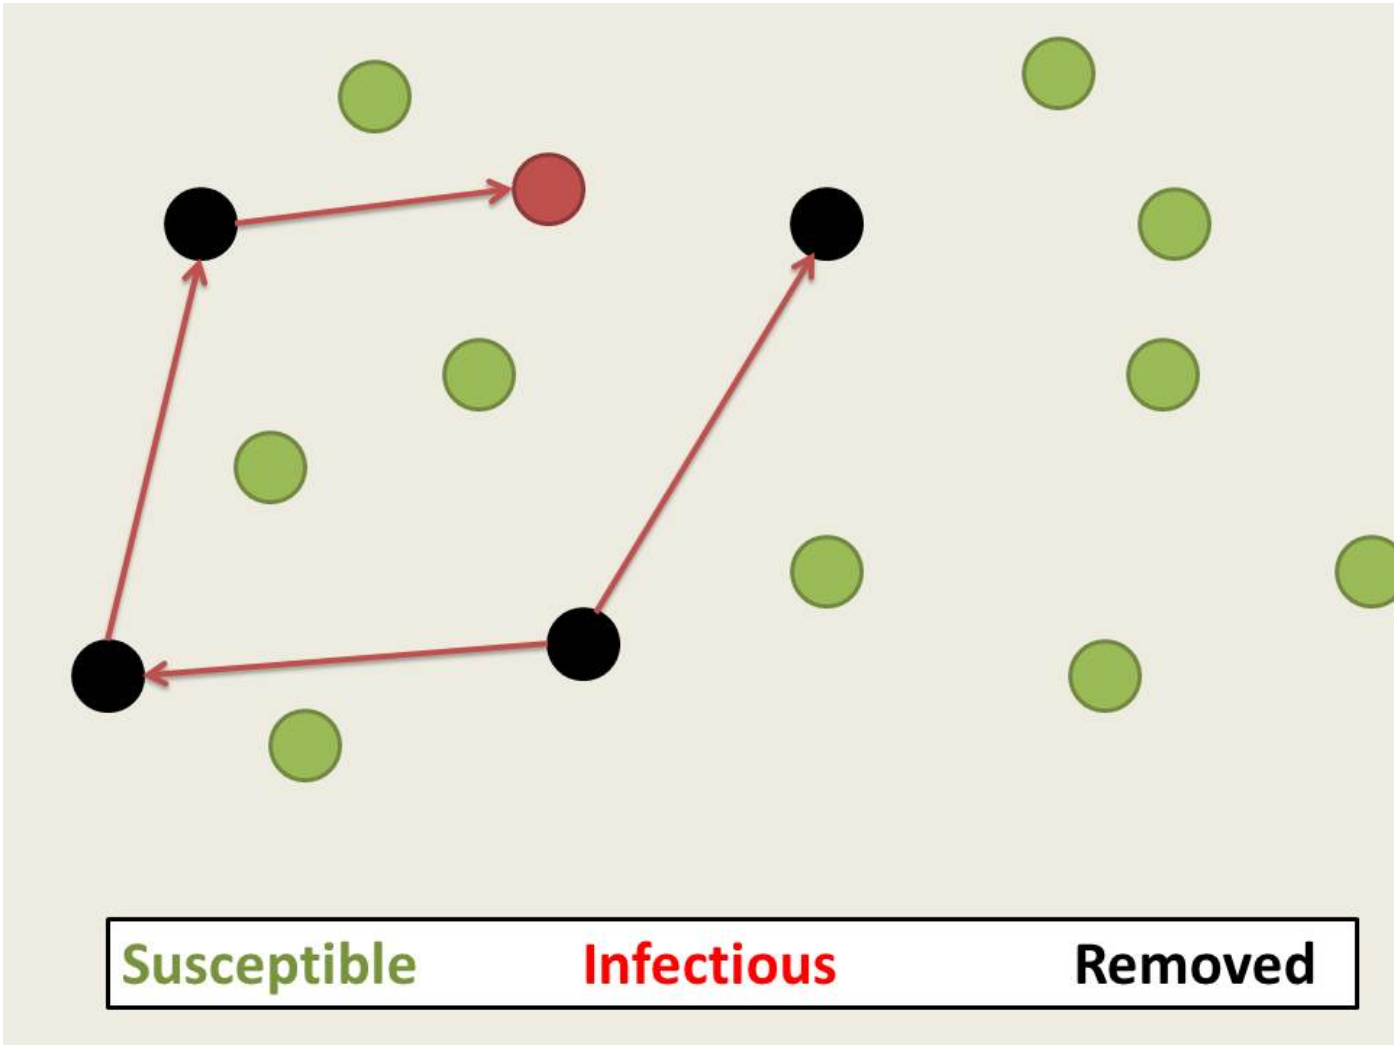

How would an epidemic get started anyway? Example 1

A new infection occurs.

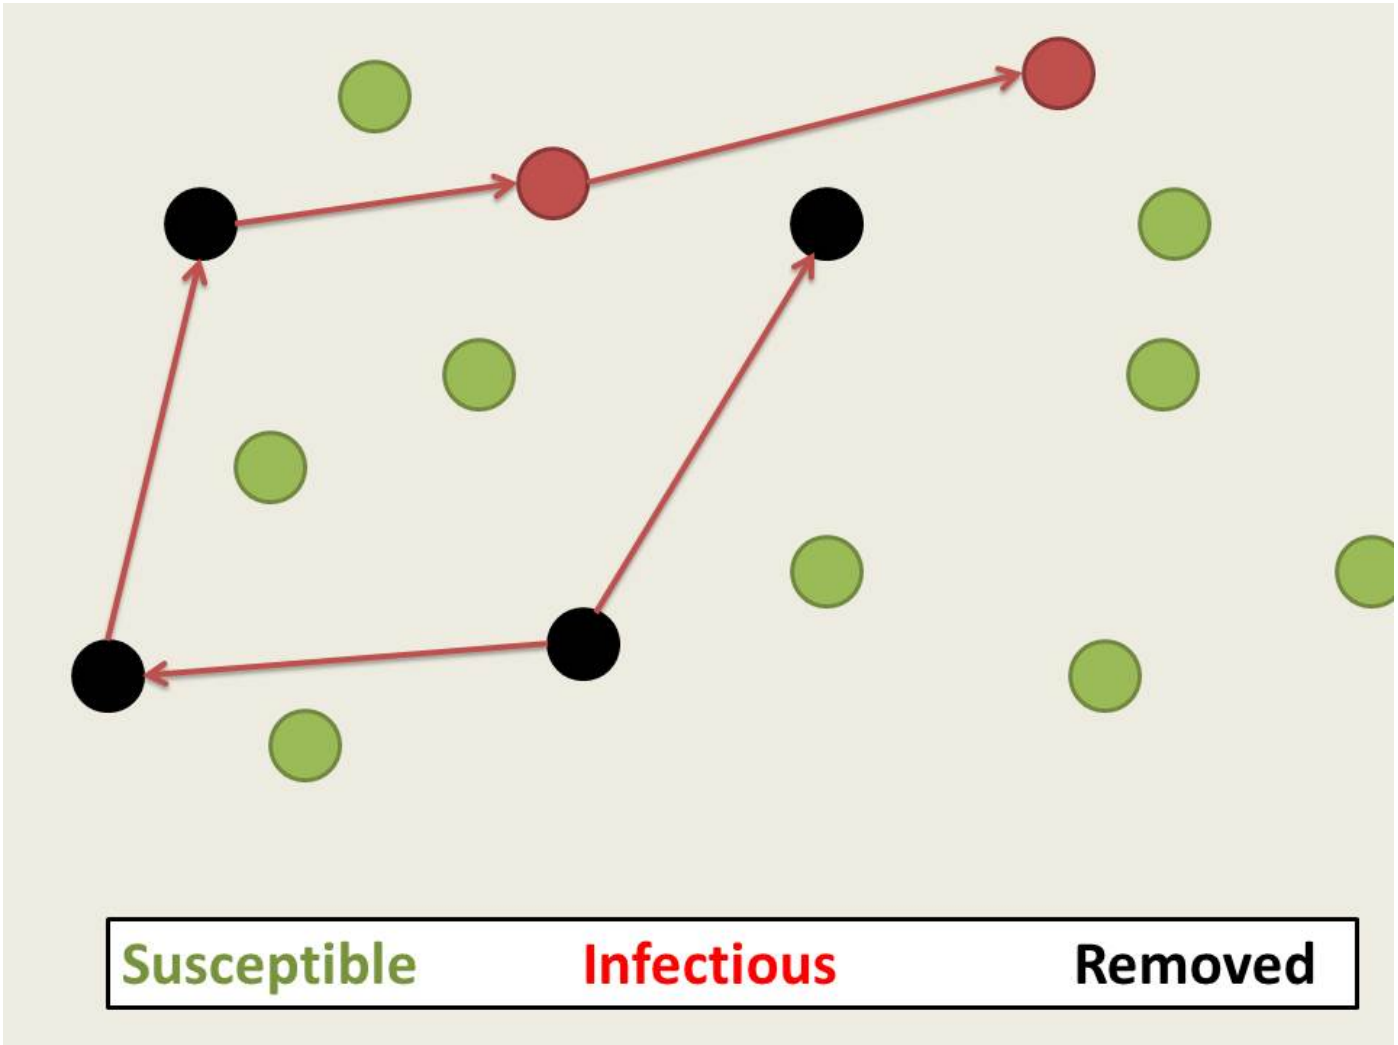

How would an epidemic get started anyway? Example 1

An infectious host is removed.

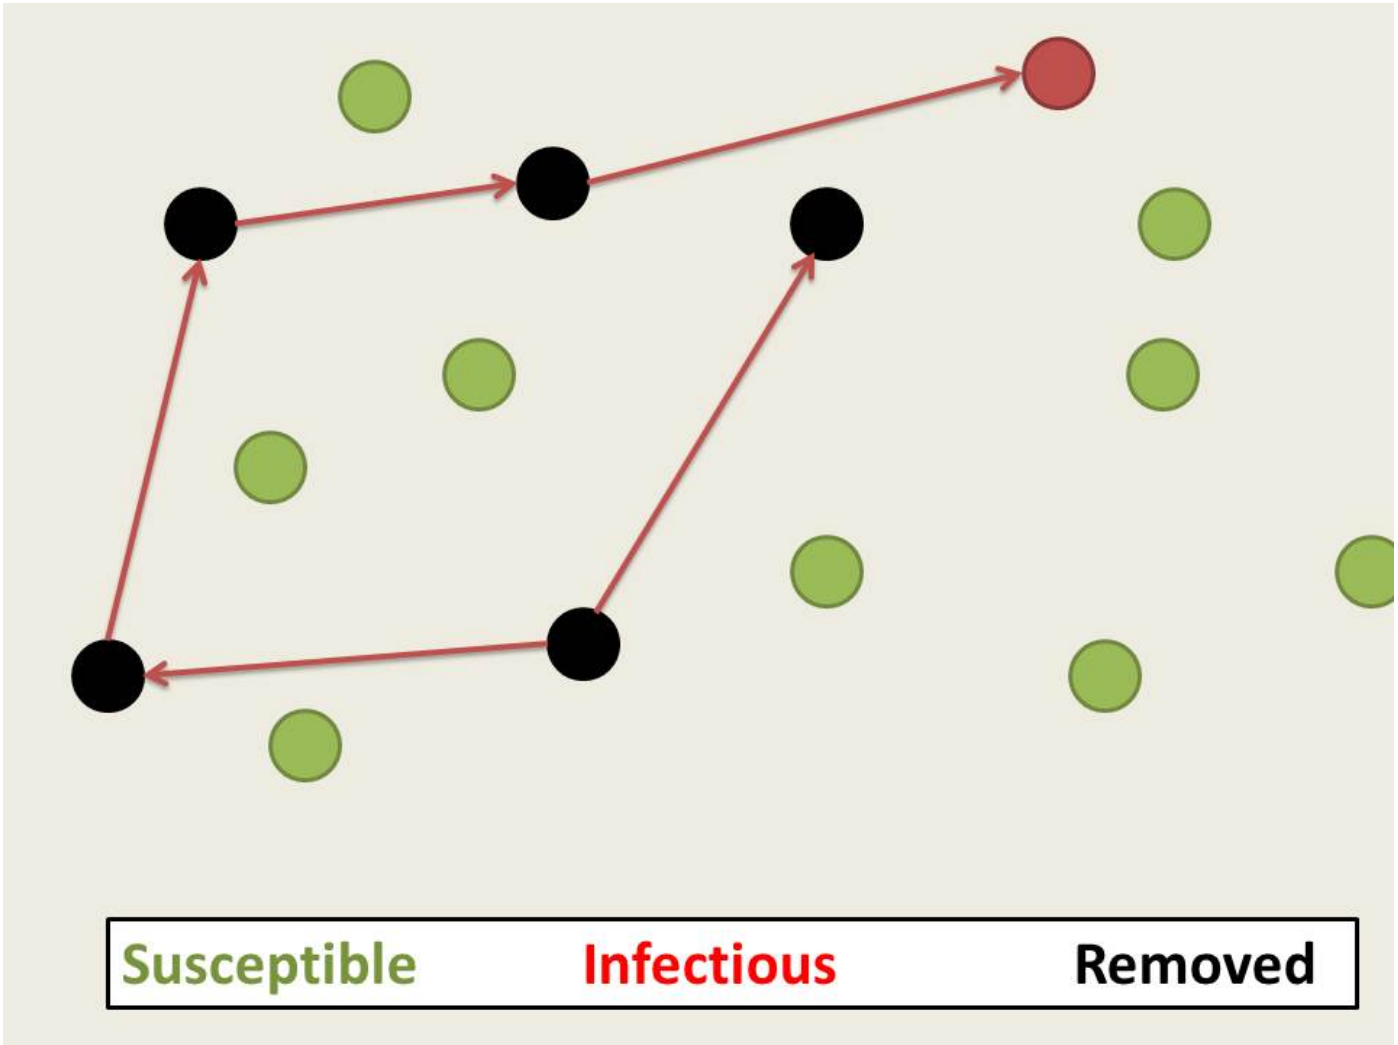

How would an epidemic get started anyway? Example 1

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

No epidemic is observed in this example!

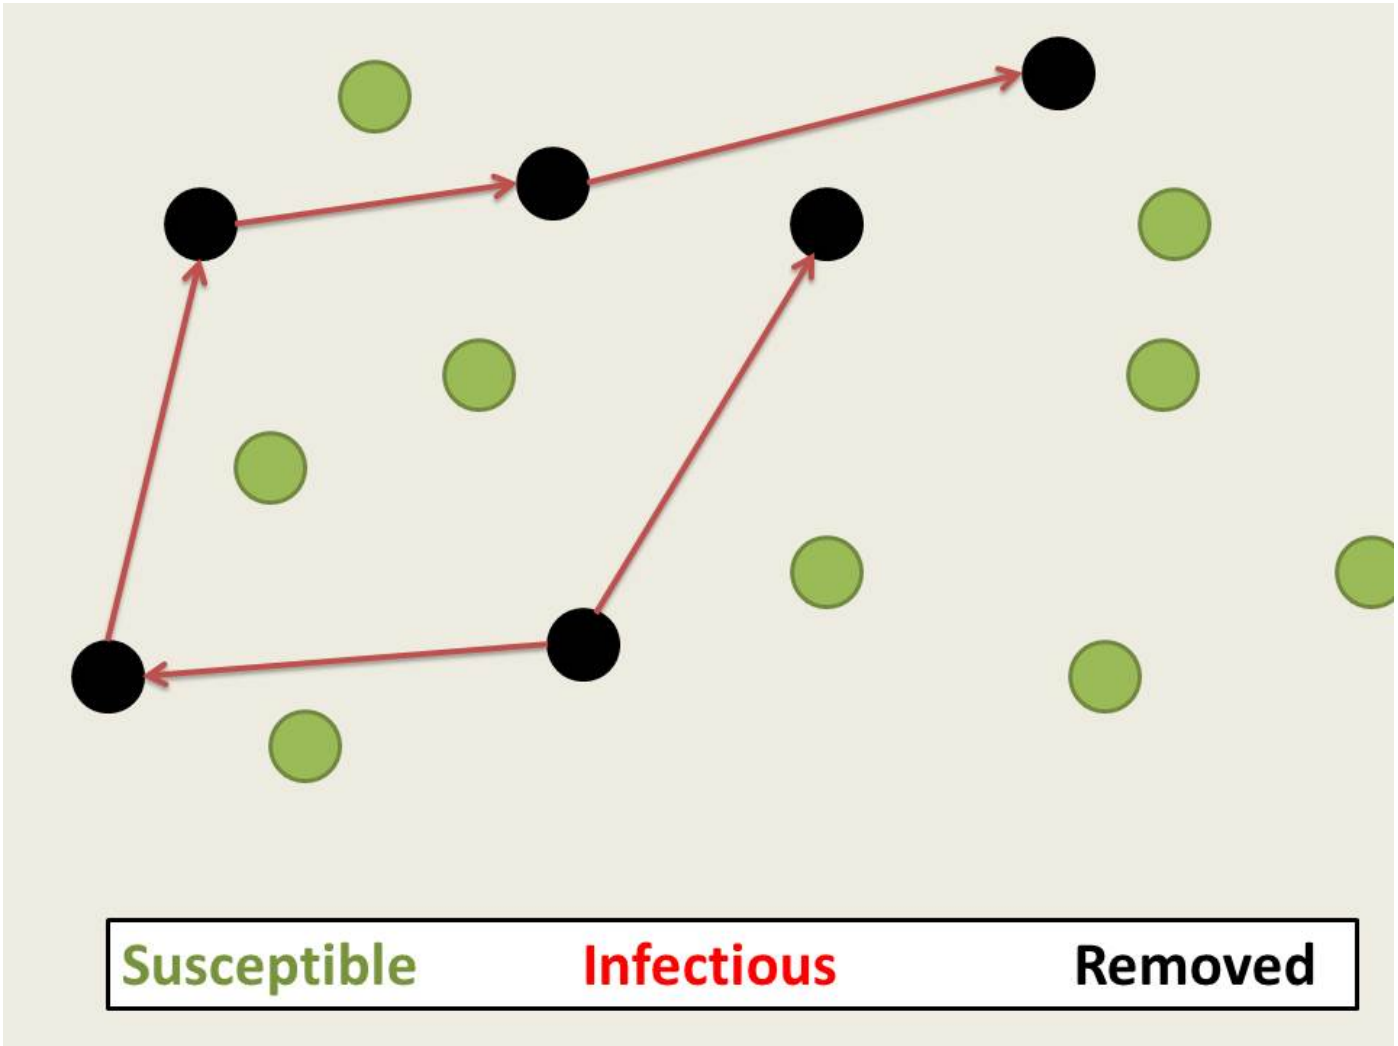

The generations of the infection in Example 1

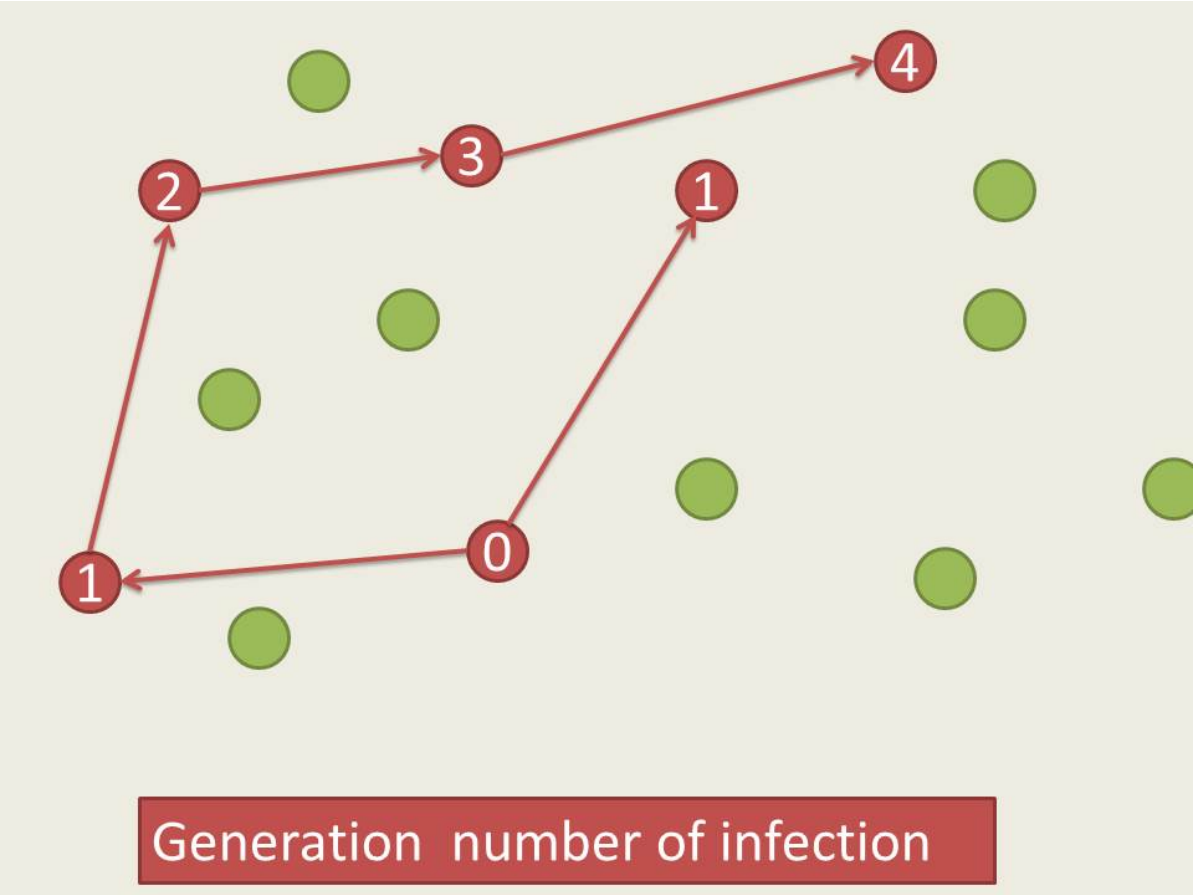

The average number of secondary infections per infectious host in this example is $\frac{2+0+1+1+1+0}{6} = \frac{5}{6}$.

How would an epidemic get started anyway? Example 2

Generations 0 and 1 might look as in Example 1.

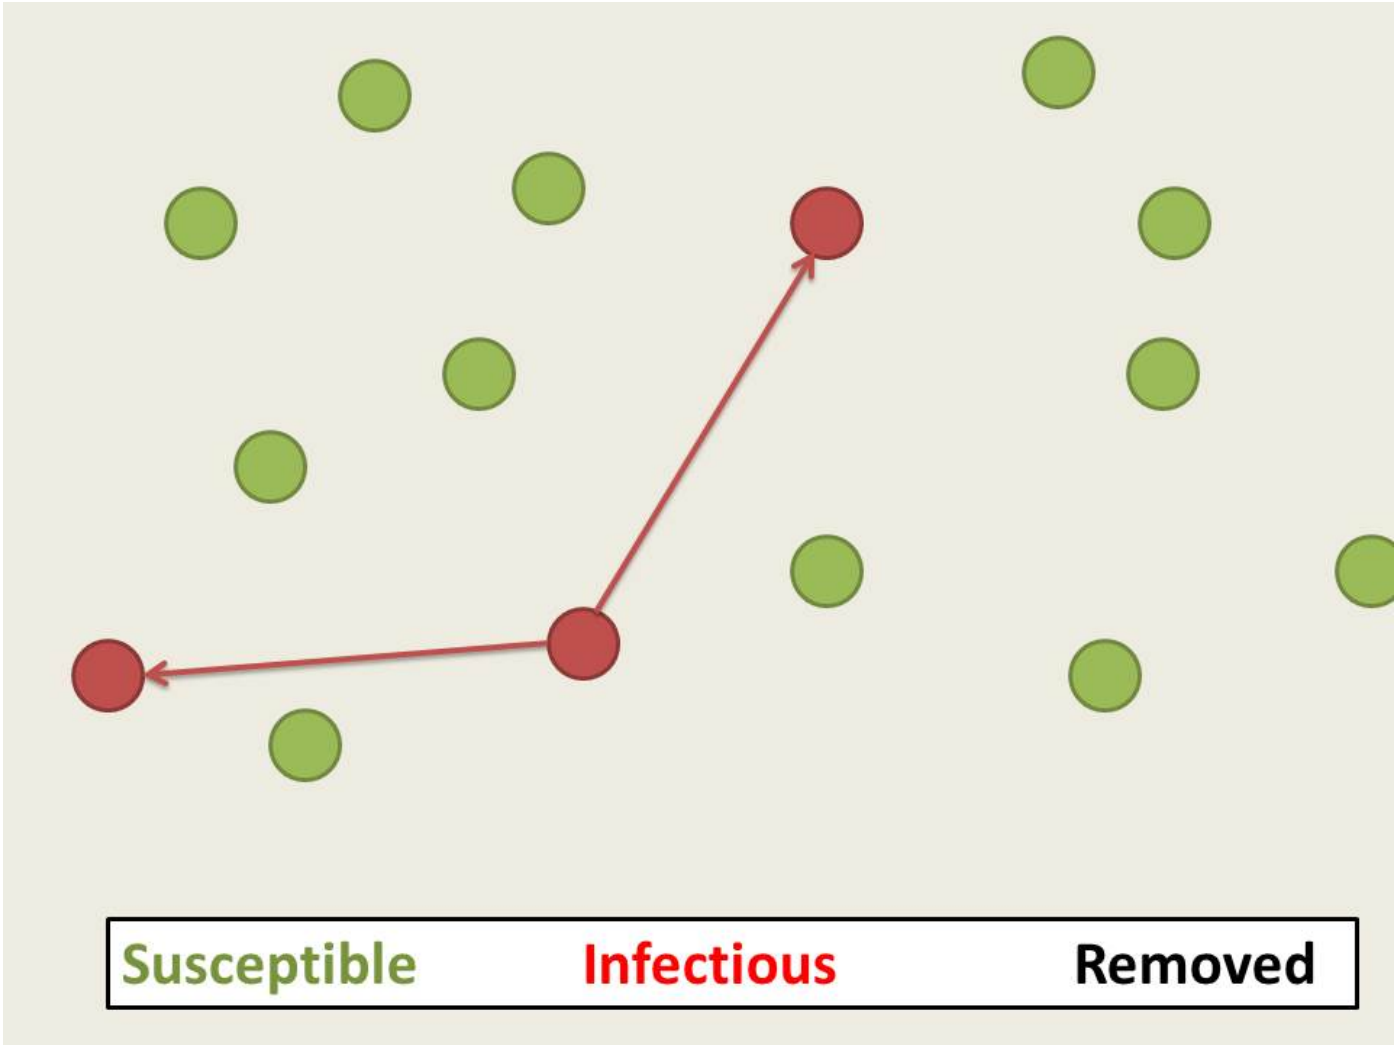

How would an epidemic get started anyway? Example 2

A new infection occurs.

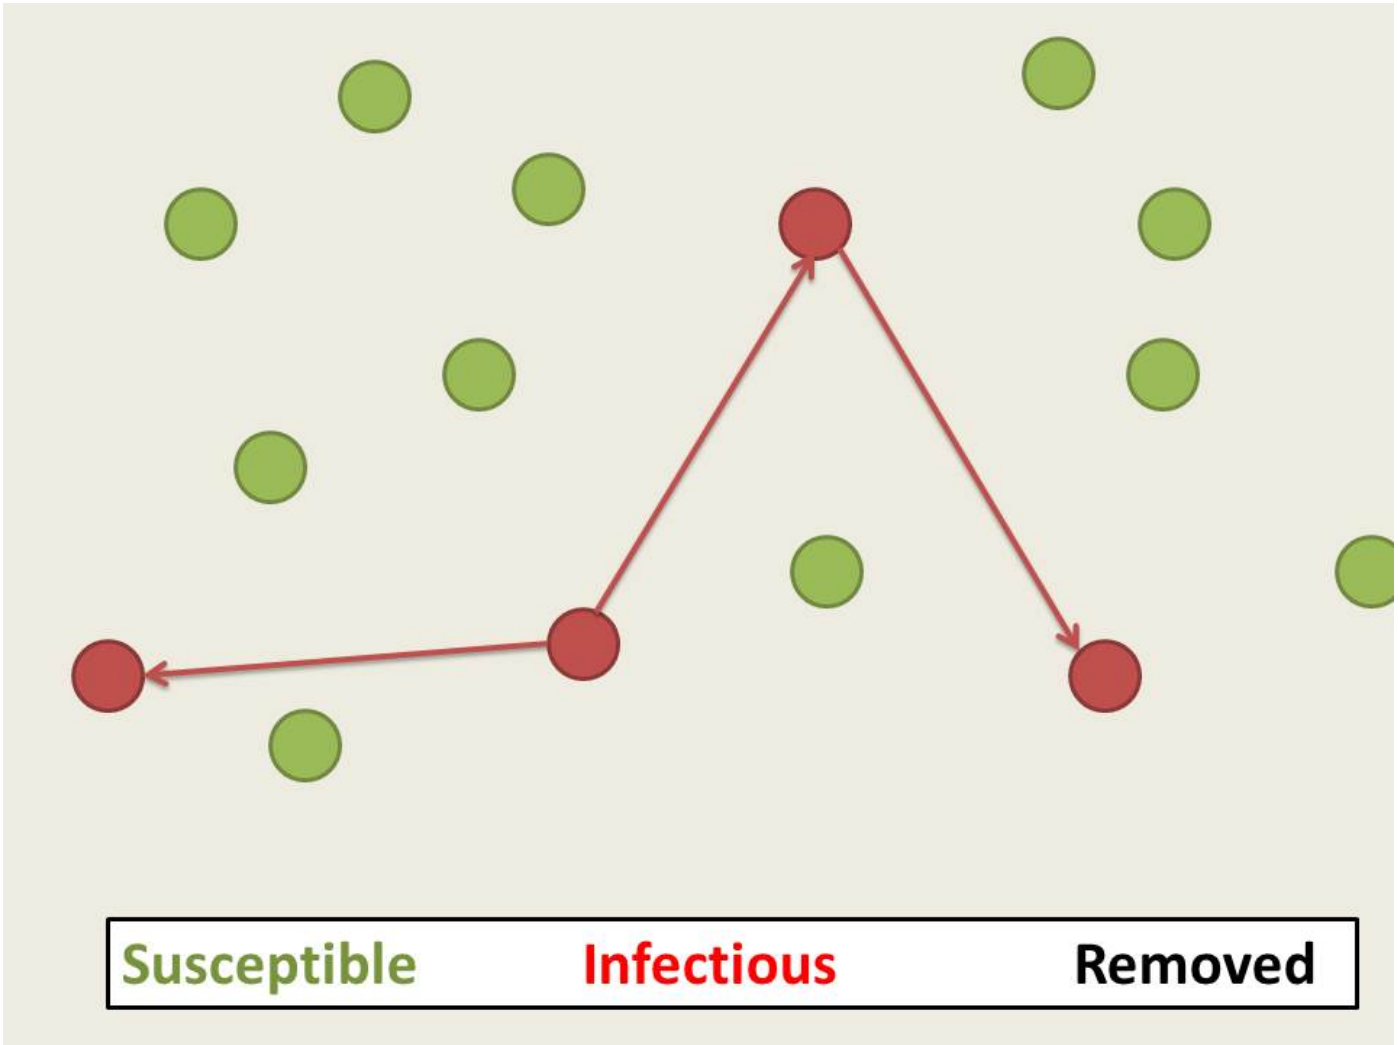

How would an epidemic get started anyway? Example 2

An infectious host is removed.

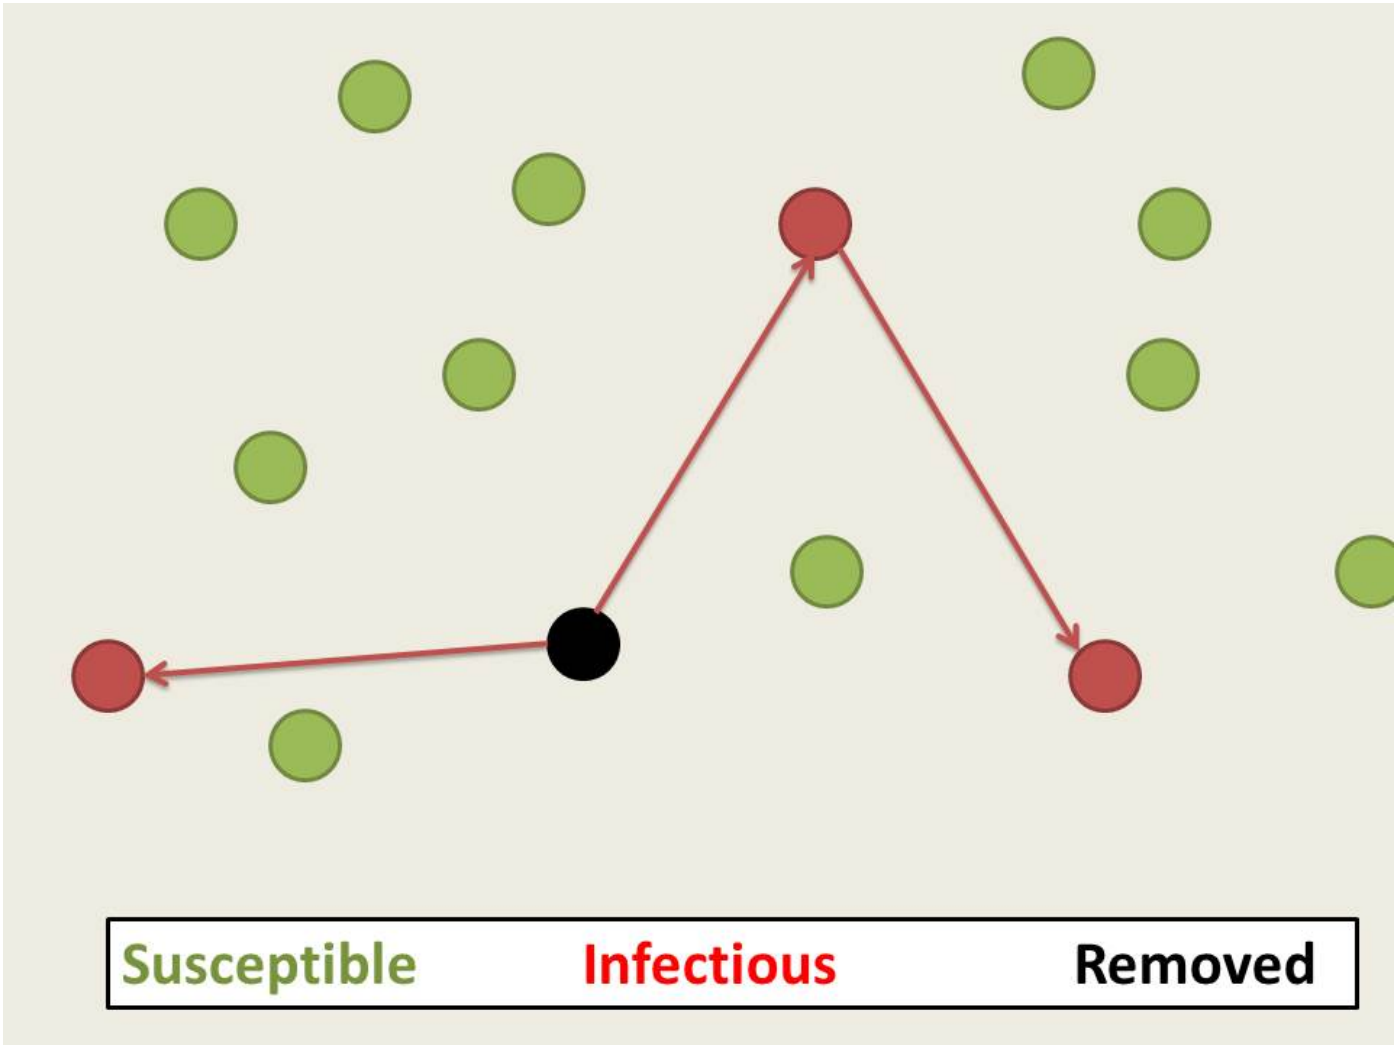

How would an epidemic get started anyway? Example 2

A new infection occurs.

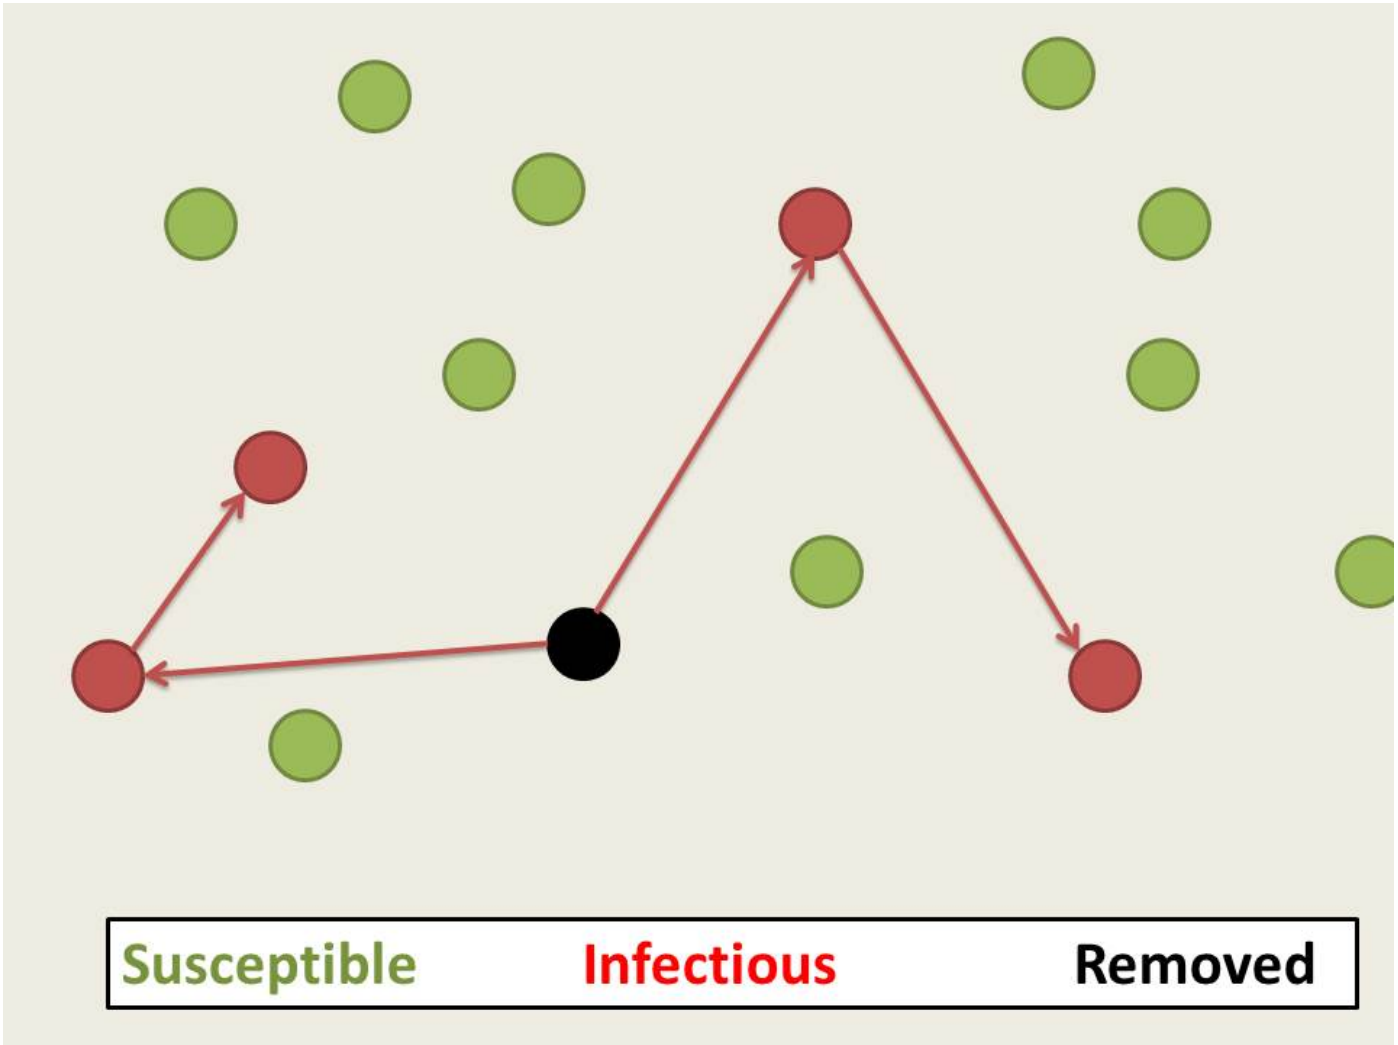

How would an epidemic get started anyway? Example 2

A new infection occurs.

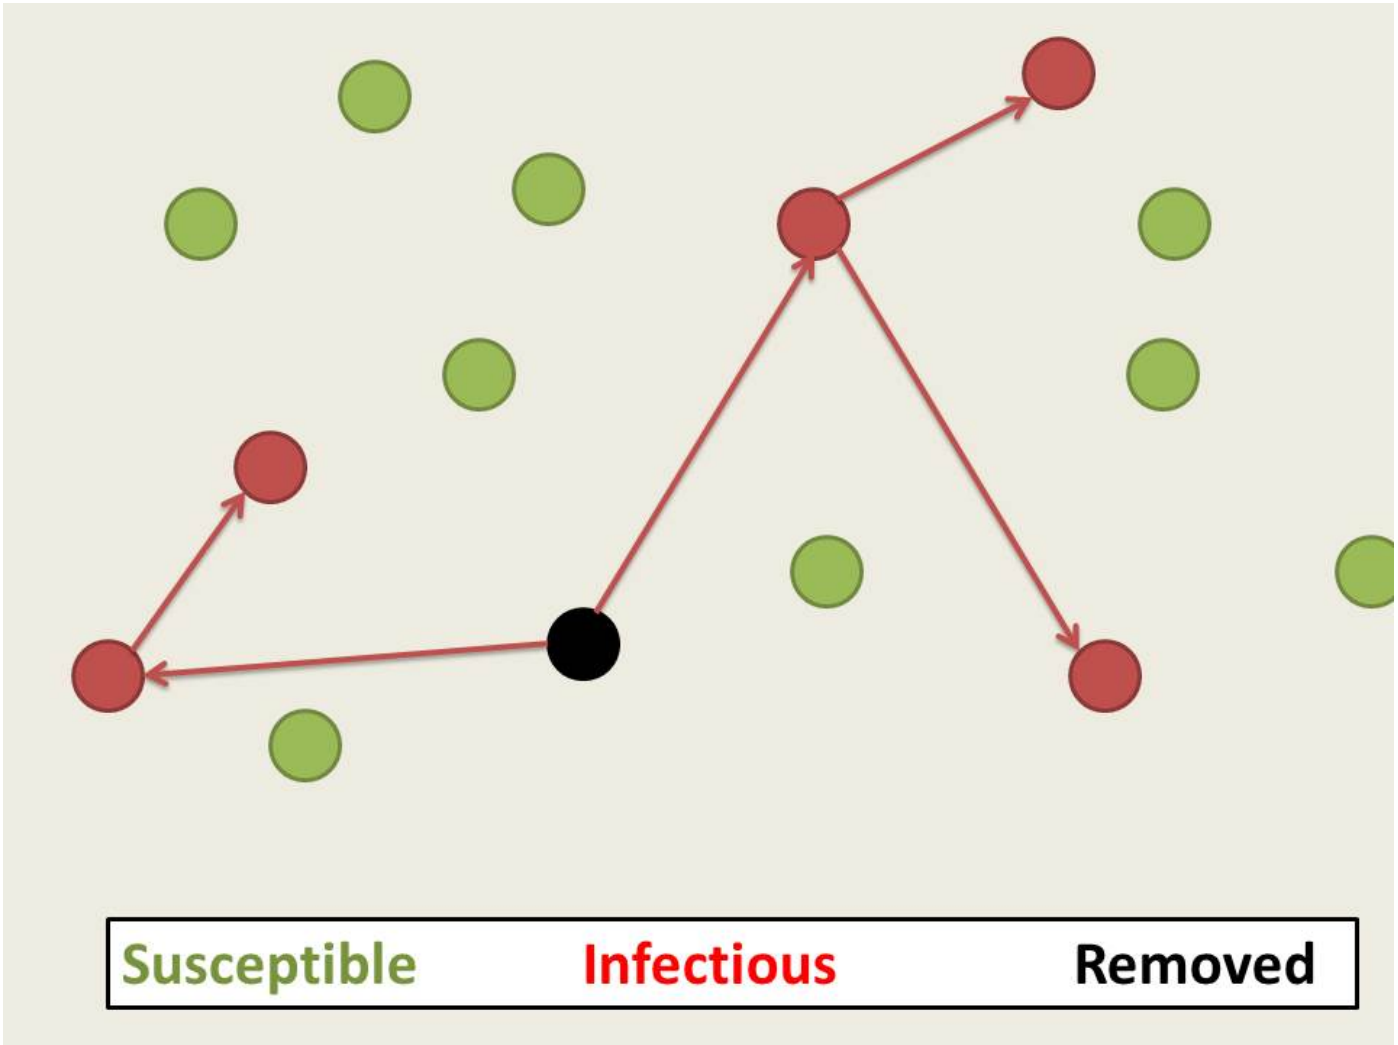

How would an epidemic get started anyway? Example 2

An infectious host is removed.

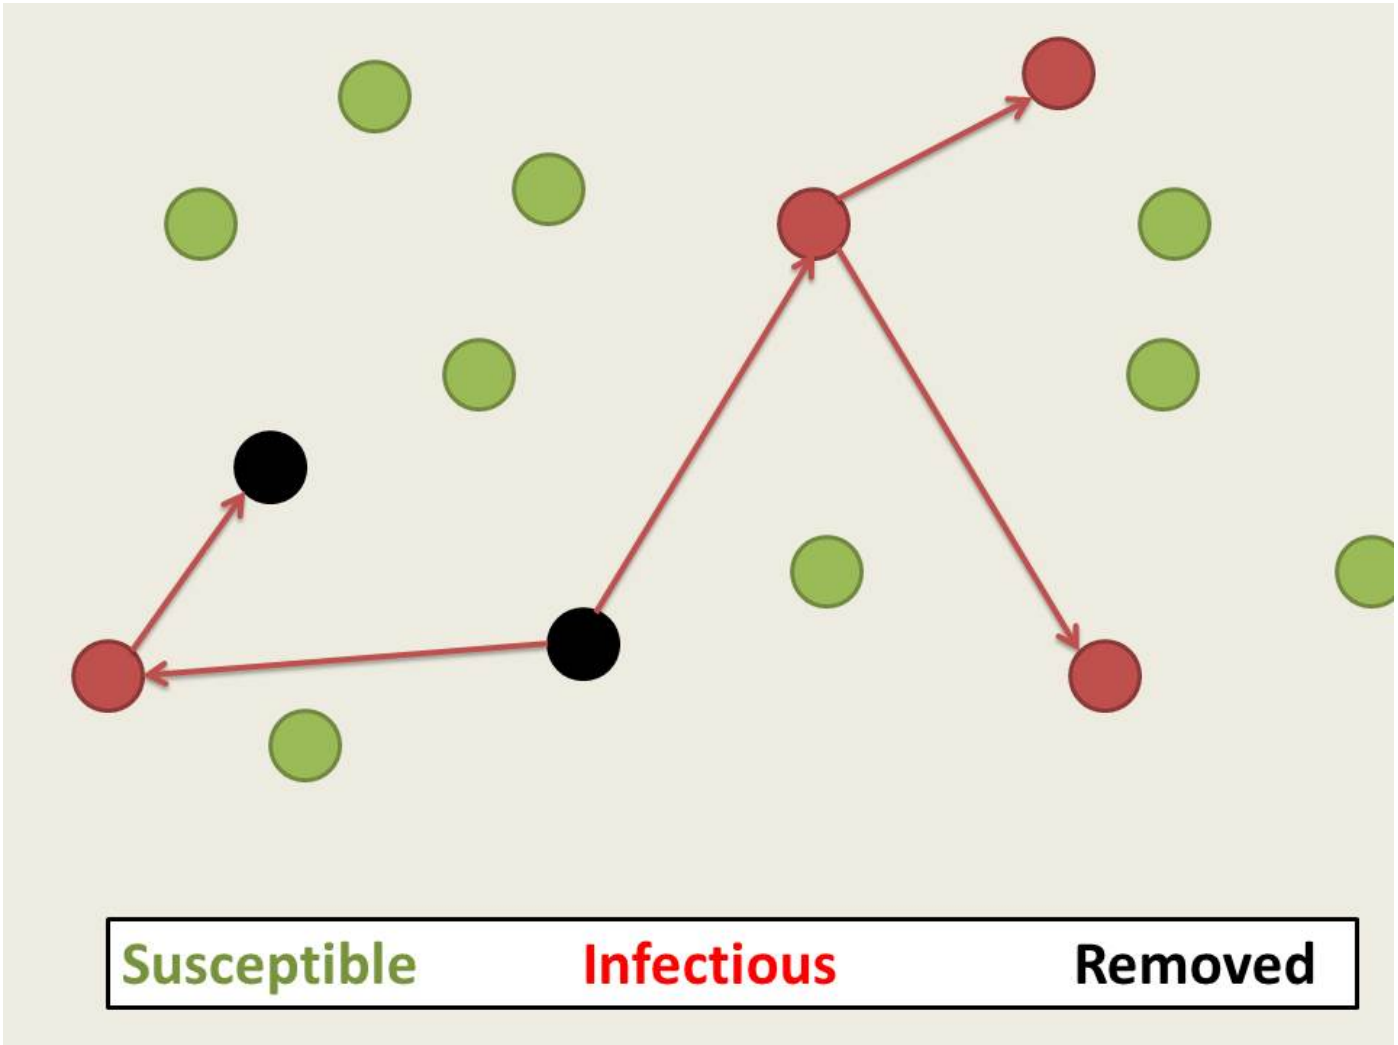

How would an epidemic get started anyway? Example 2

A new infection occurs.

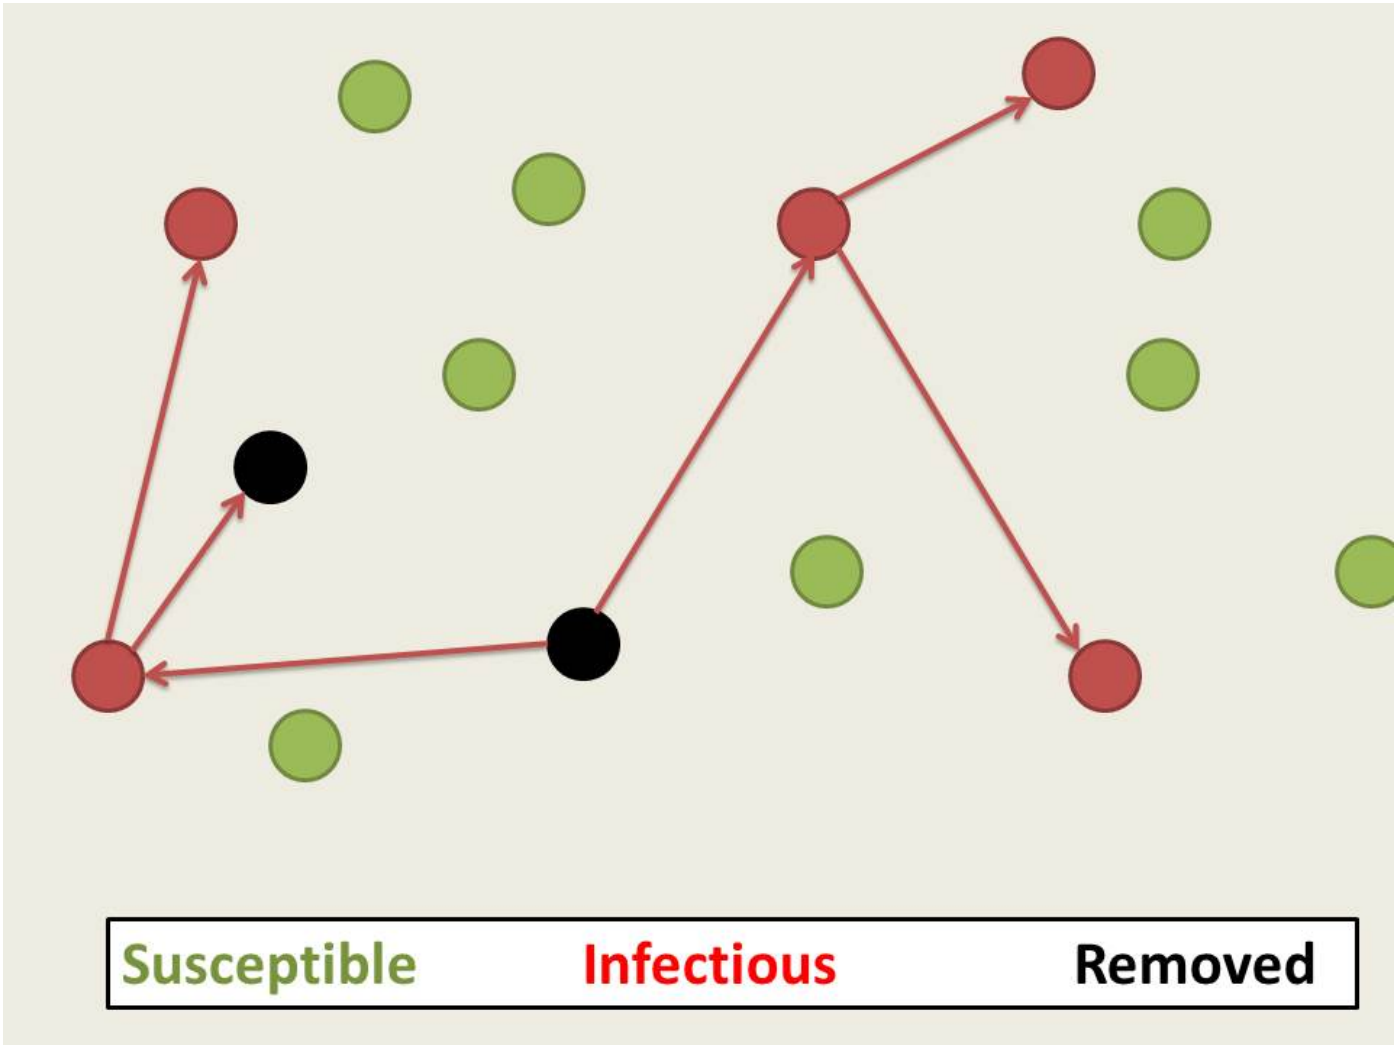

How would an epidemic get started anyway? Example 2

A new infection occurs.

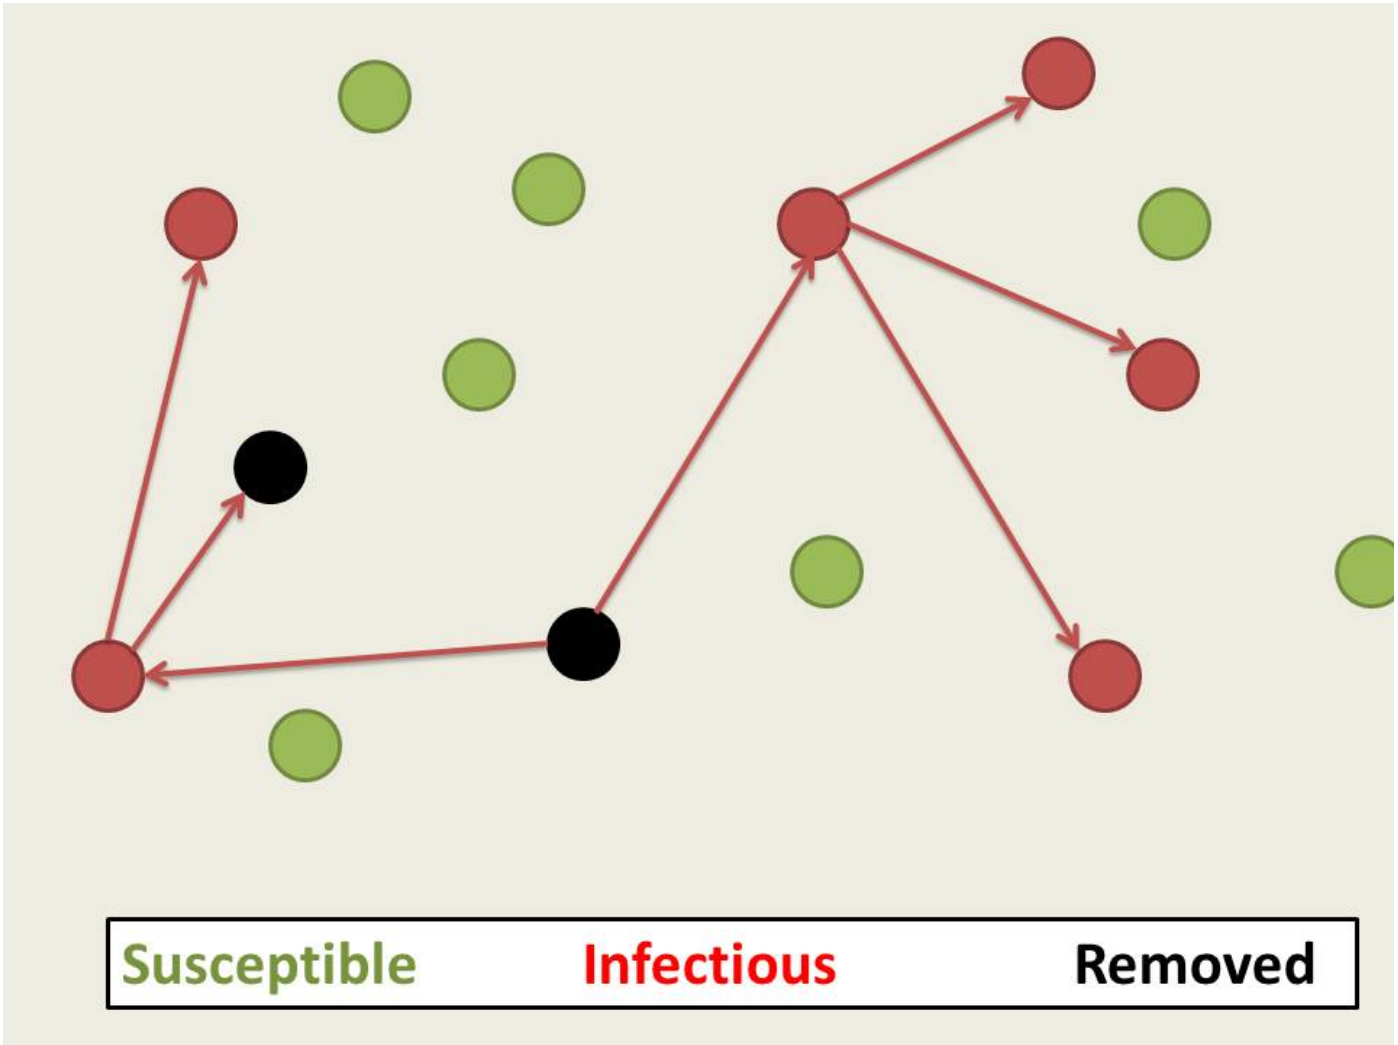

How would an epidemic get started anyway? Example 2

A new infection occurs.

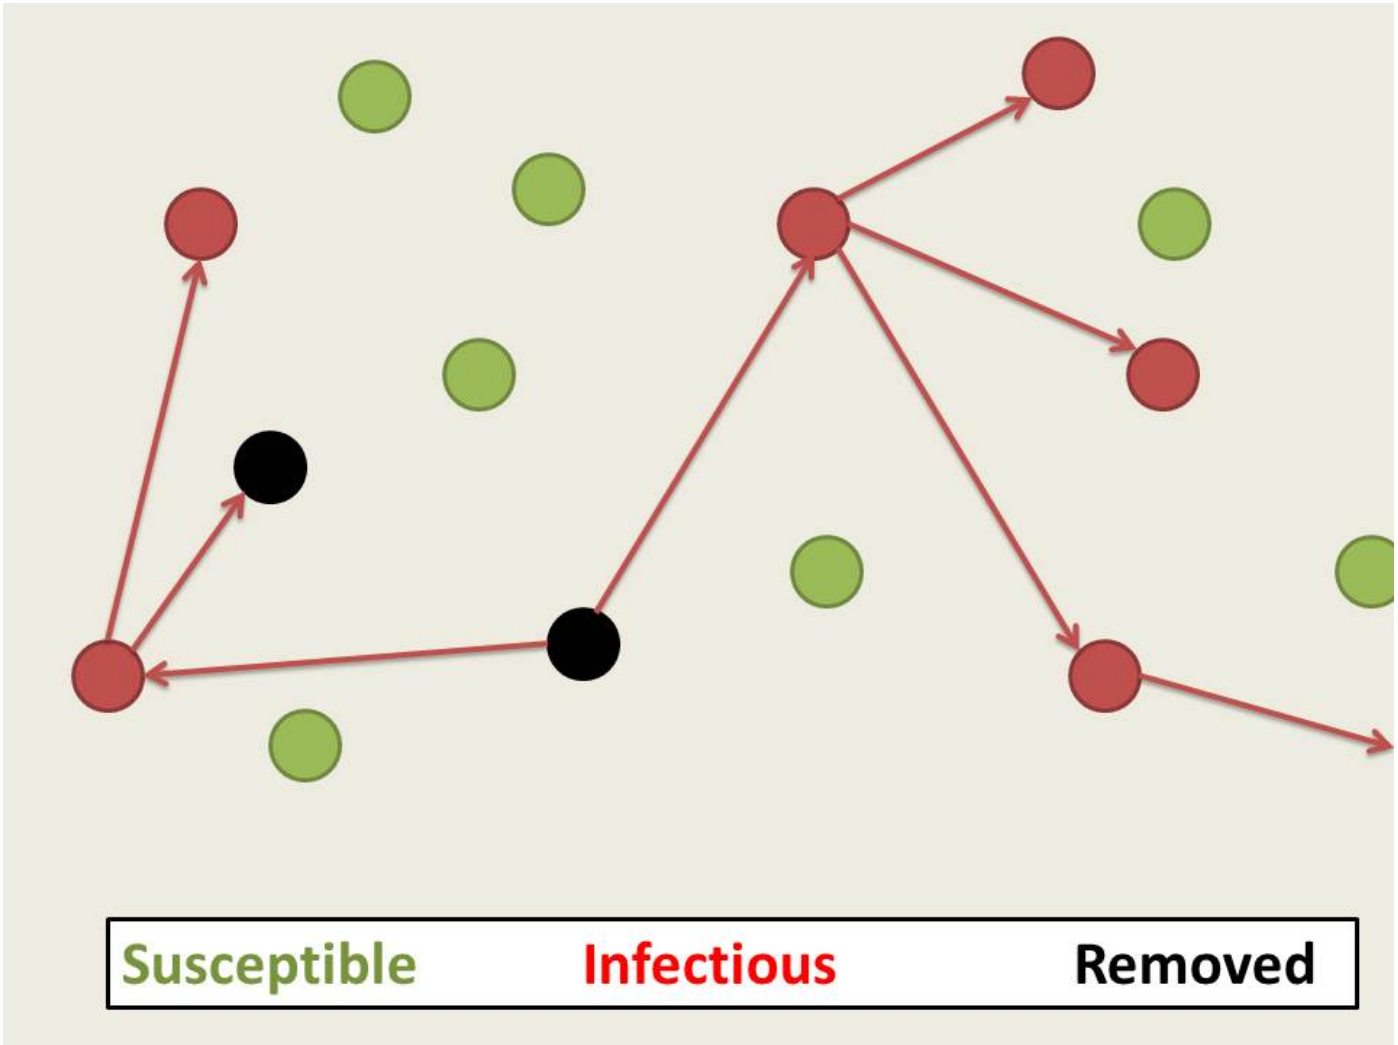

How would an epidemic get started anyway? Example 2

An infectious host is removed.

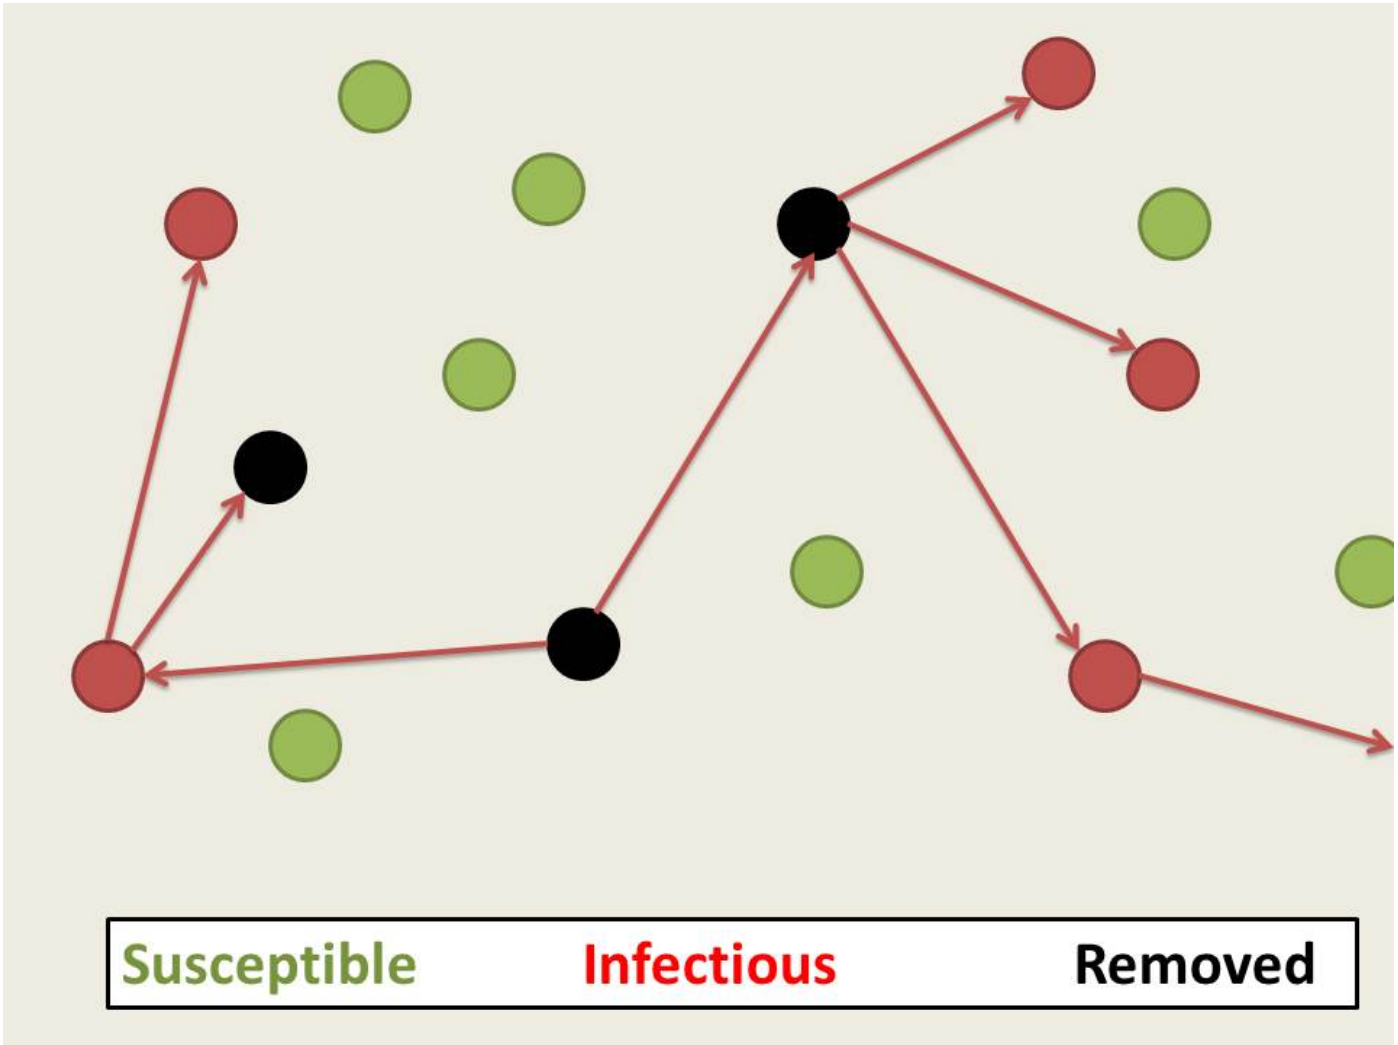

How would an epidemic get started anyway? Example 2

A new infection occurs.

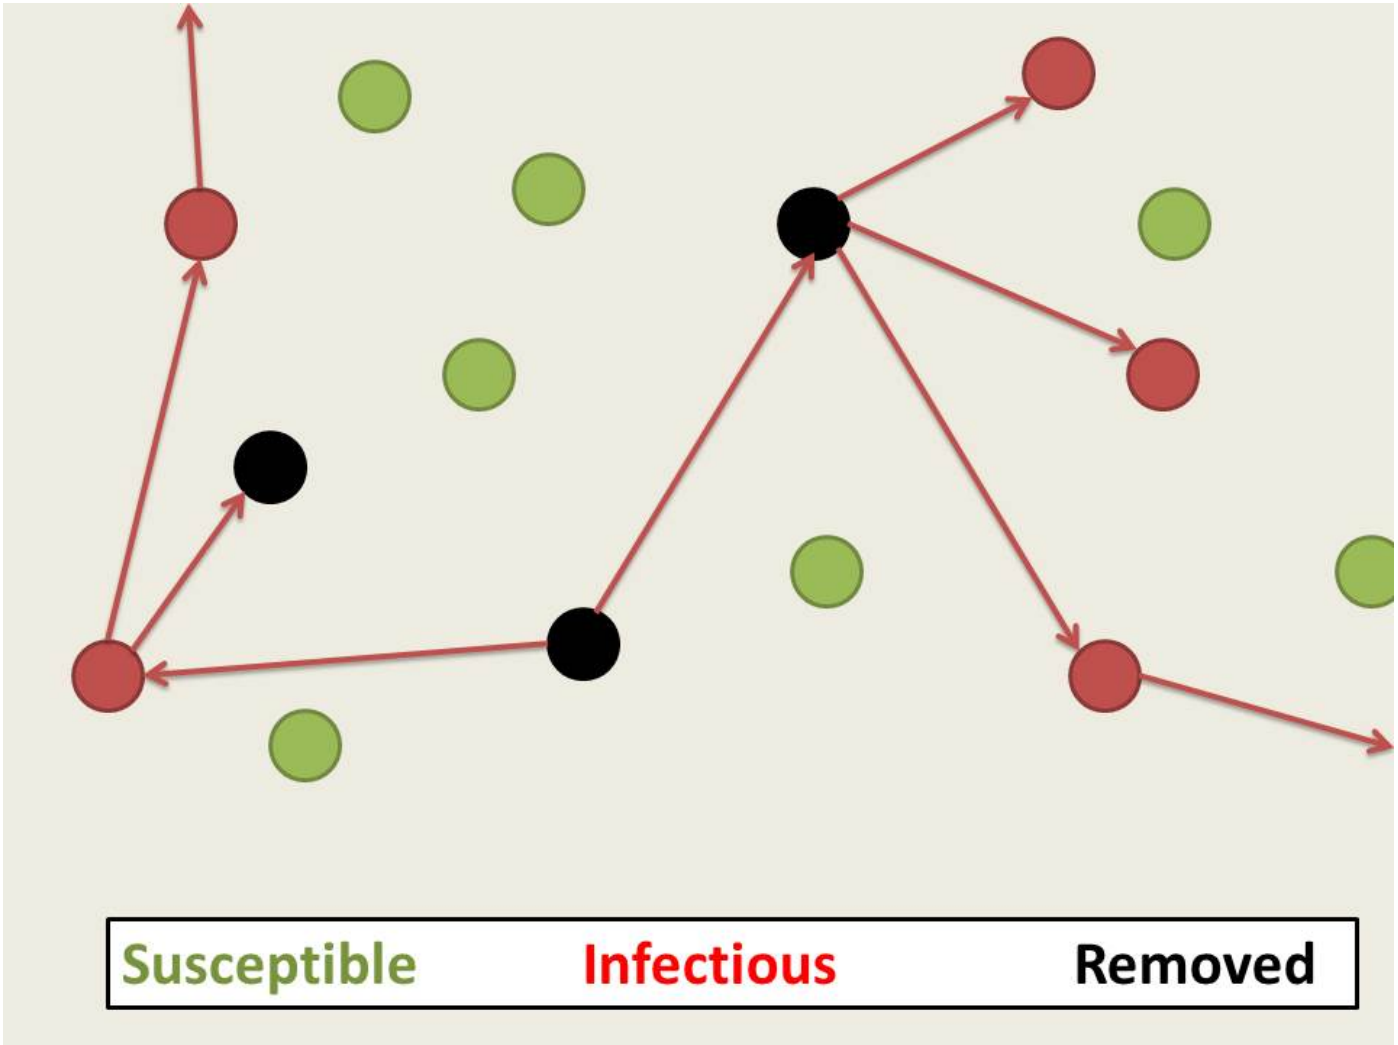

How would an epidemic get started anyway? Example 2

An infectious host is removed.

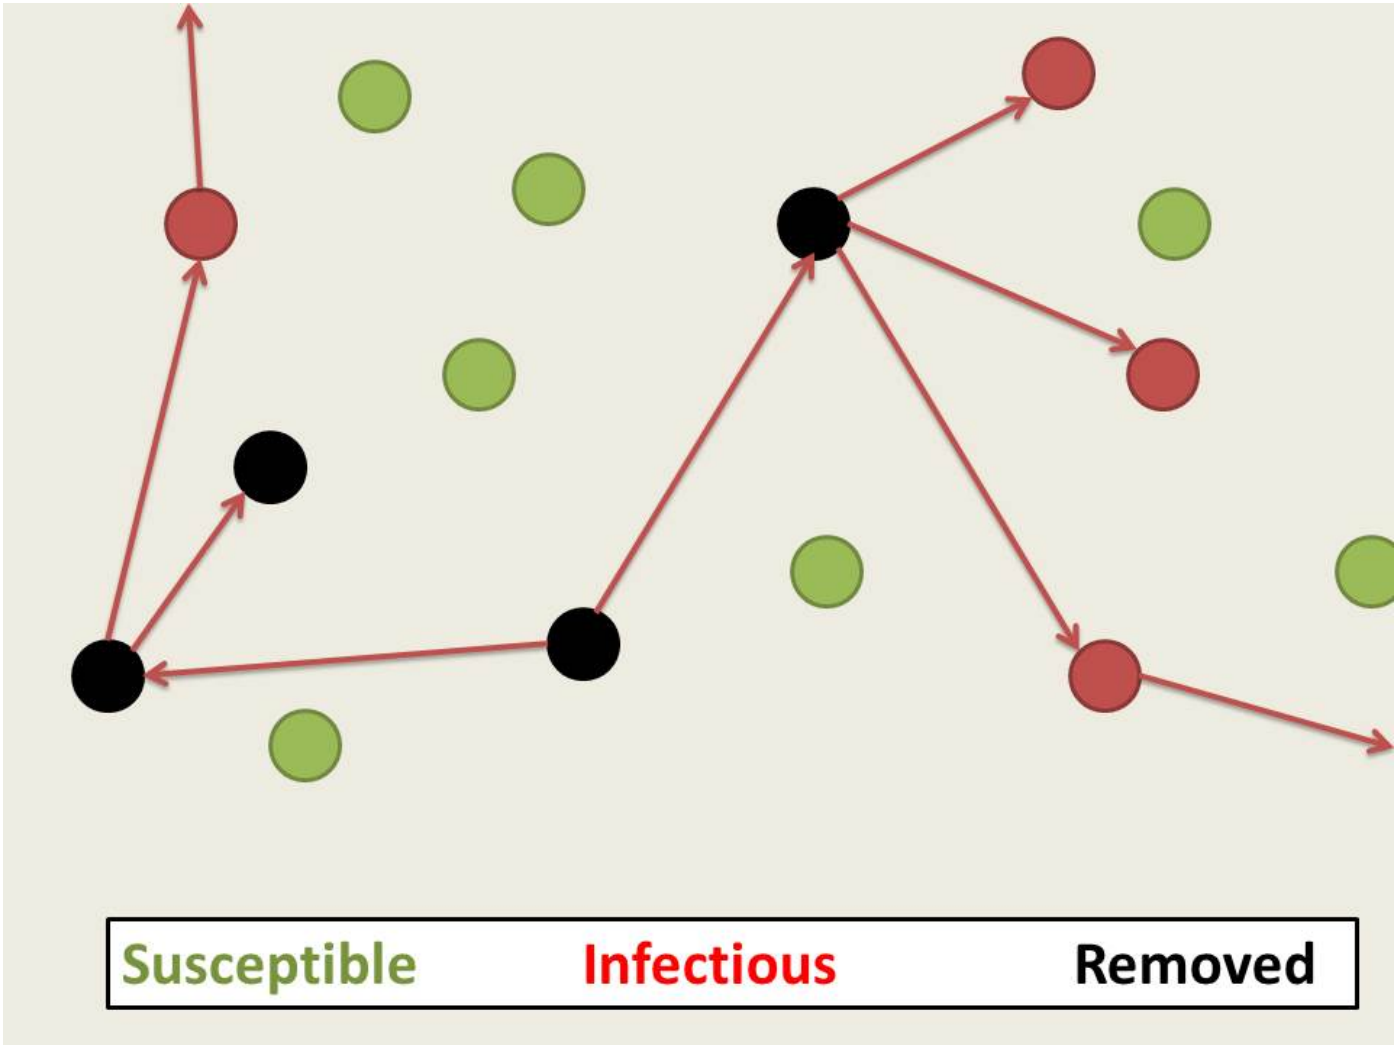

How would an epidemic get started anyway? Example 2

A new infection occurs.

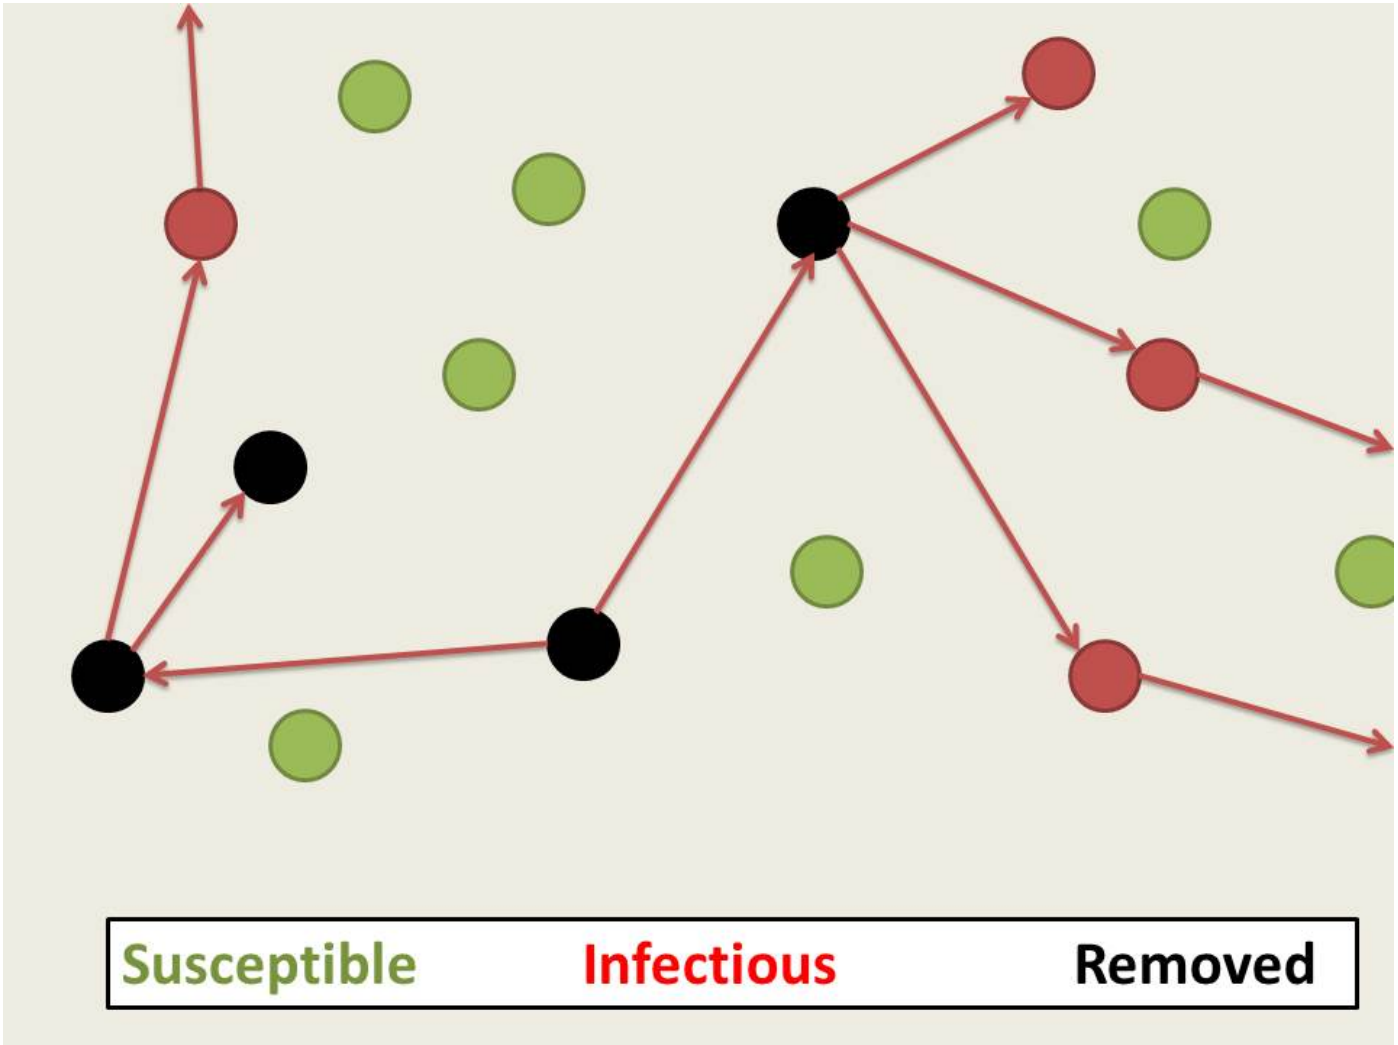

How would an epidemic get started anyway? Example 2

An infectious host is removed.

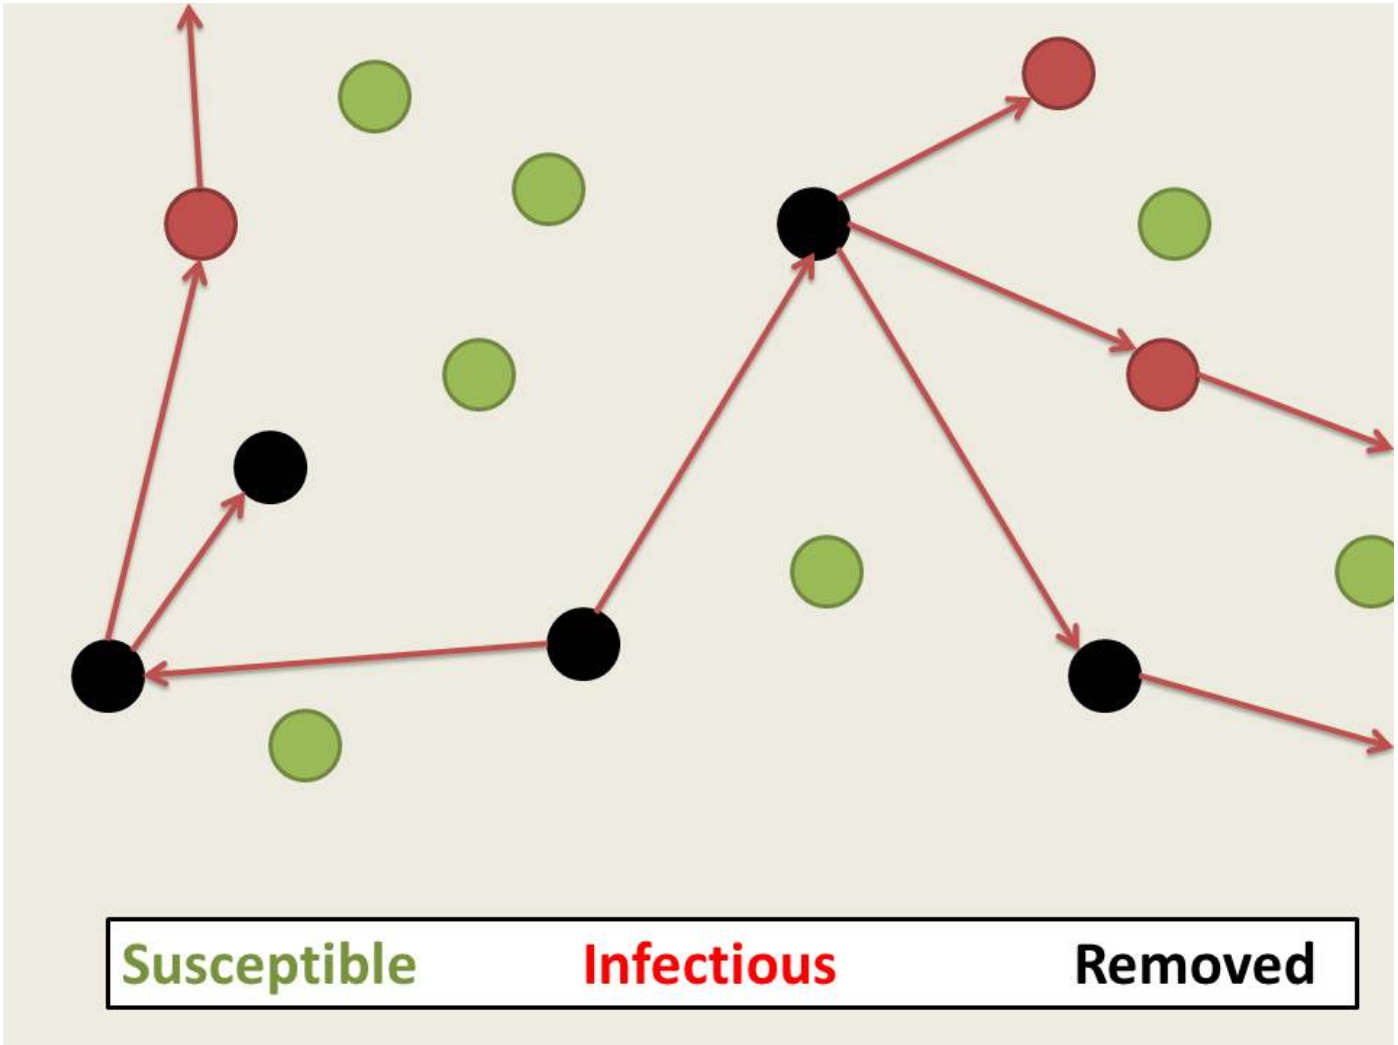

How would an epidemic get started anyway? Example 2

A new infection occurs.

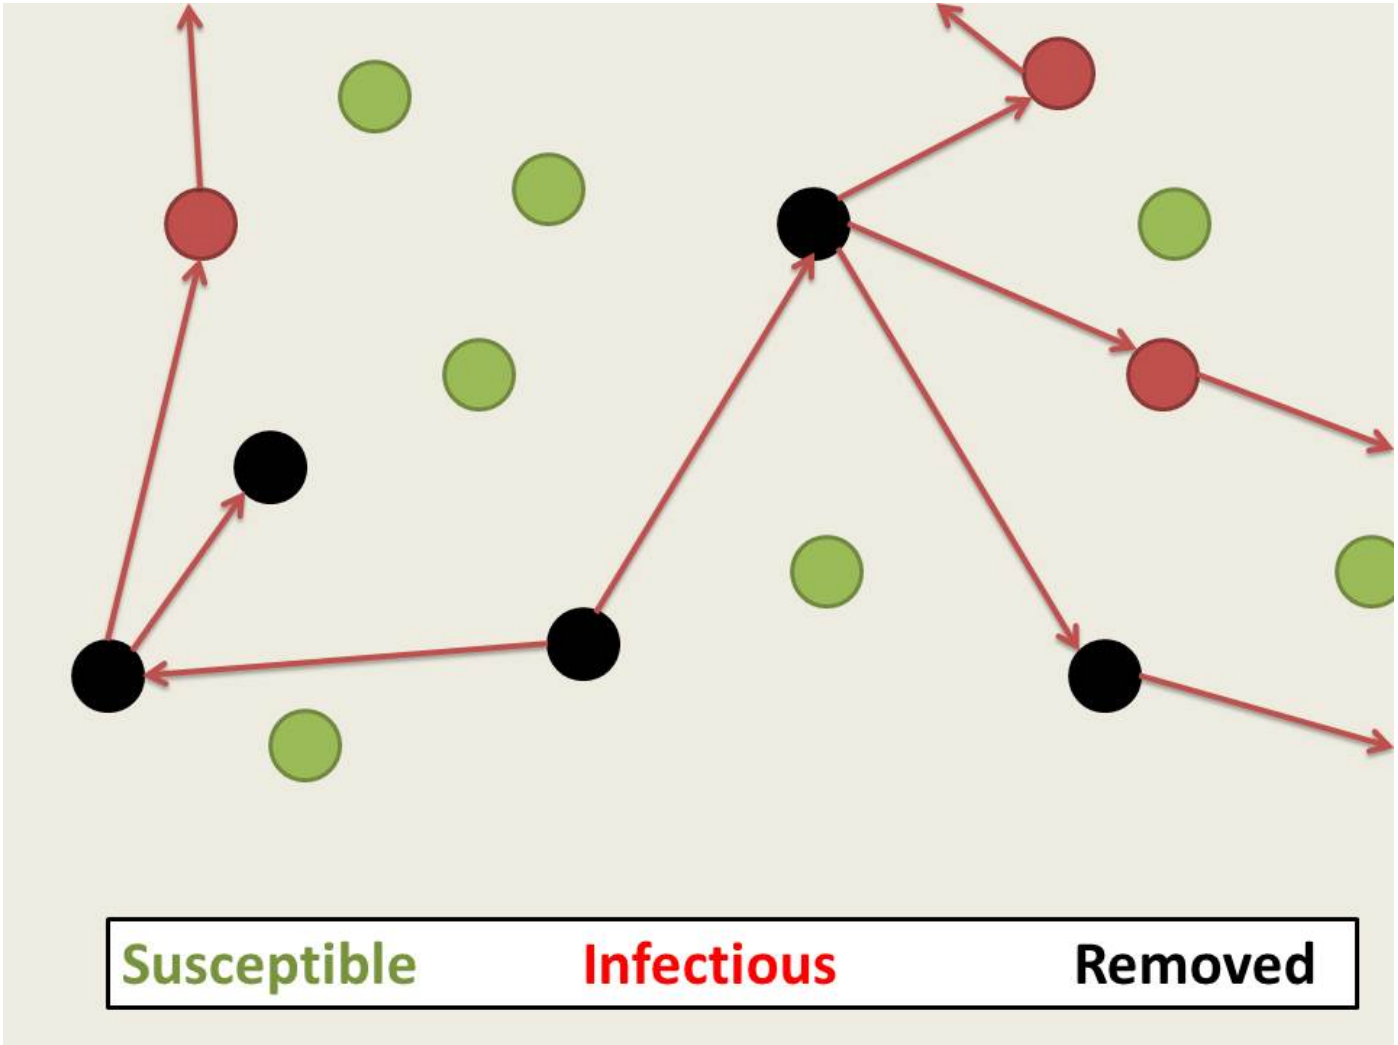

How would an epidemic get started anyway? Example 2

A new infection occurs.

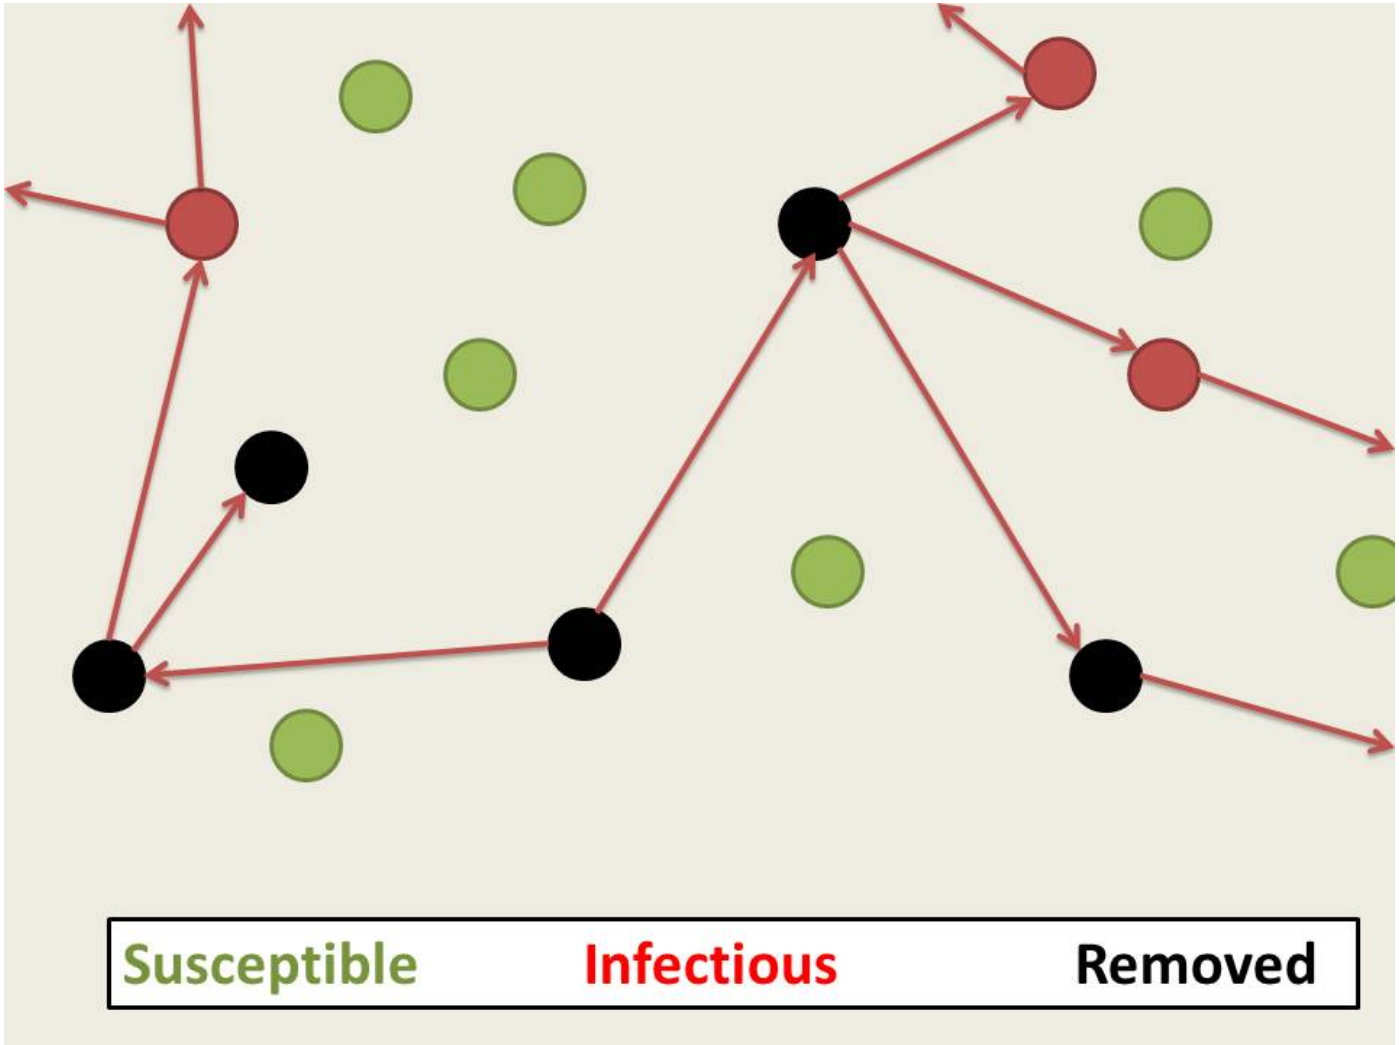

How would an epidemic get started anyway? Example 2

A new infection occurs.

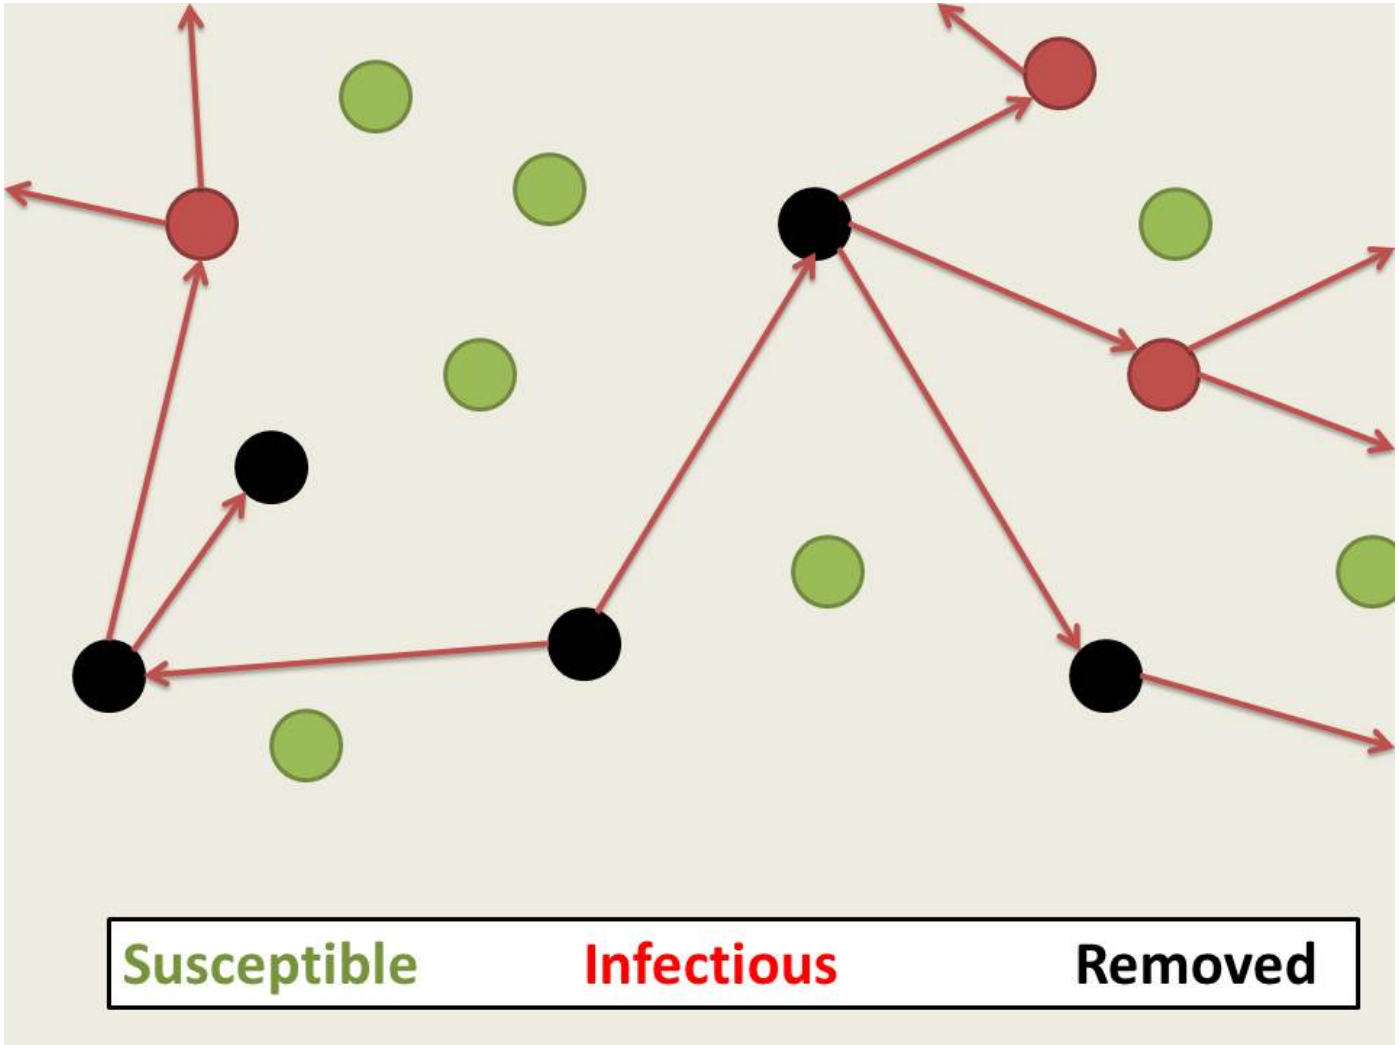

How would an epidemic get started anyway? Example 2

An infectious host is removed.

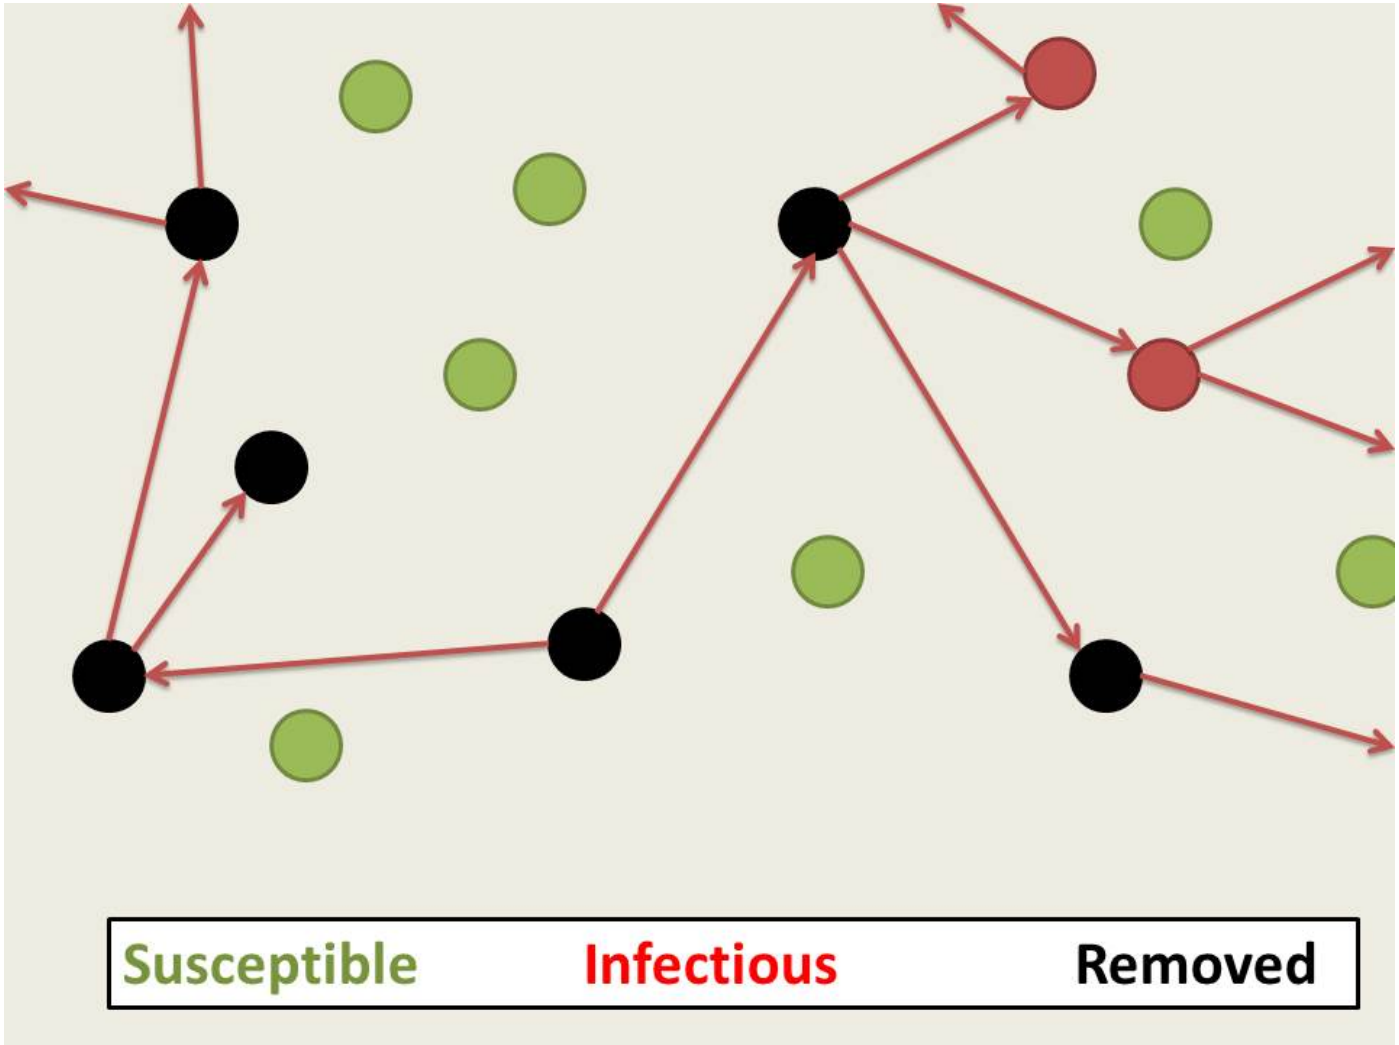

How would an epidemic get started anyway? Example 2

A new infection occurs.

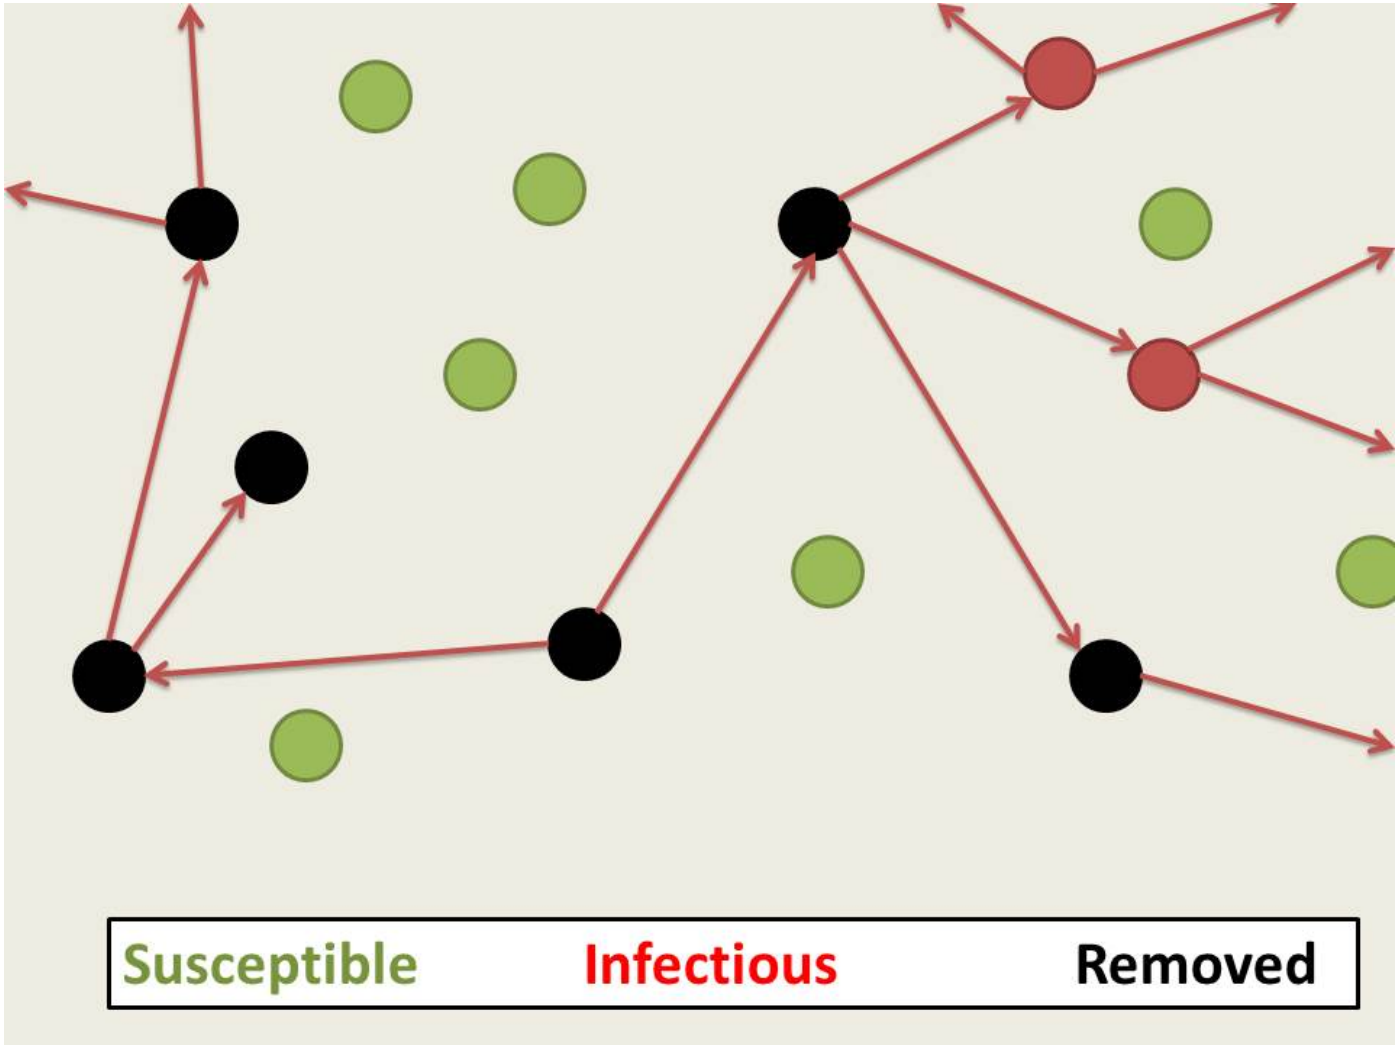

How would an epidemic get started anyway? Example 2

A new infection occurs.

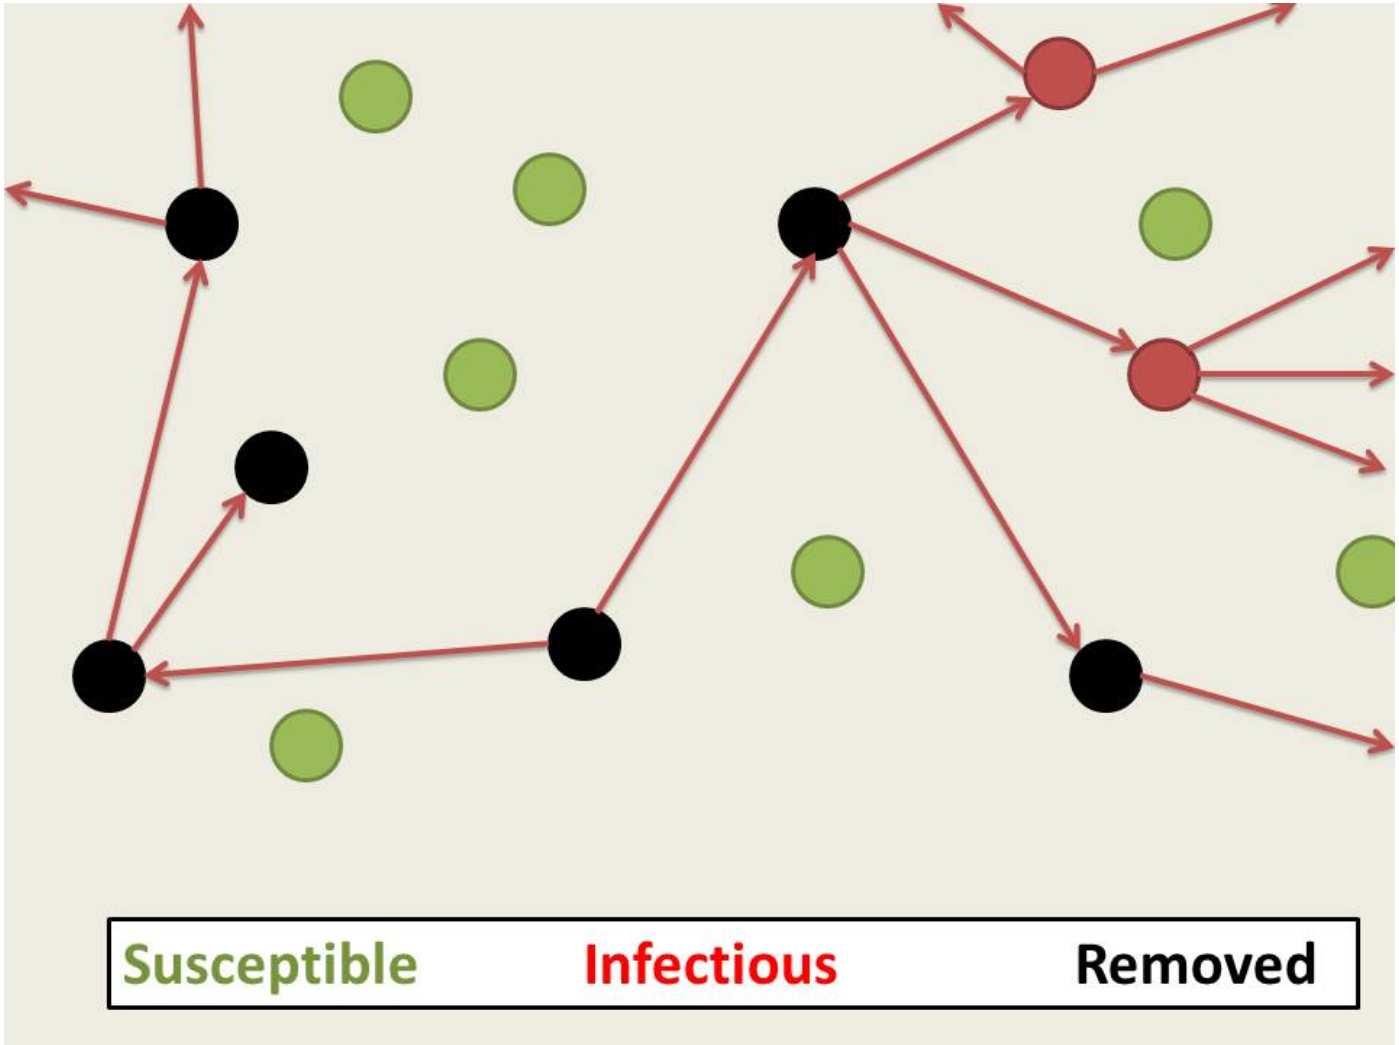

How would an epidemic get started anyway? Example 2

An infectious host is removed.

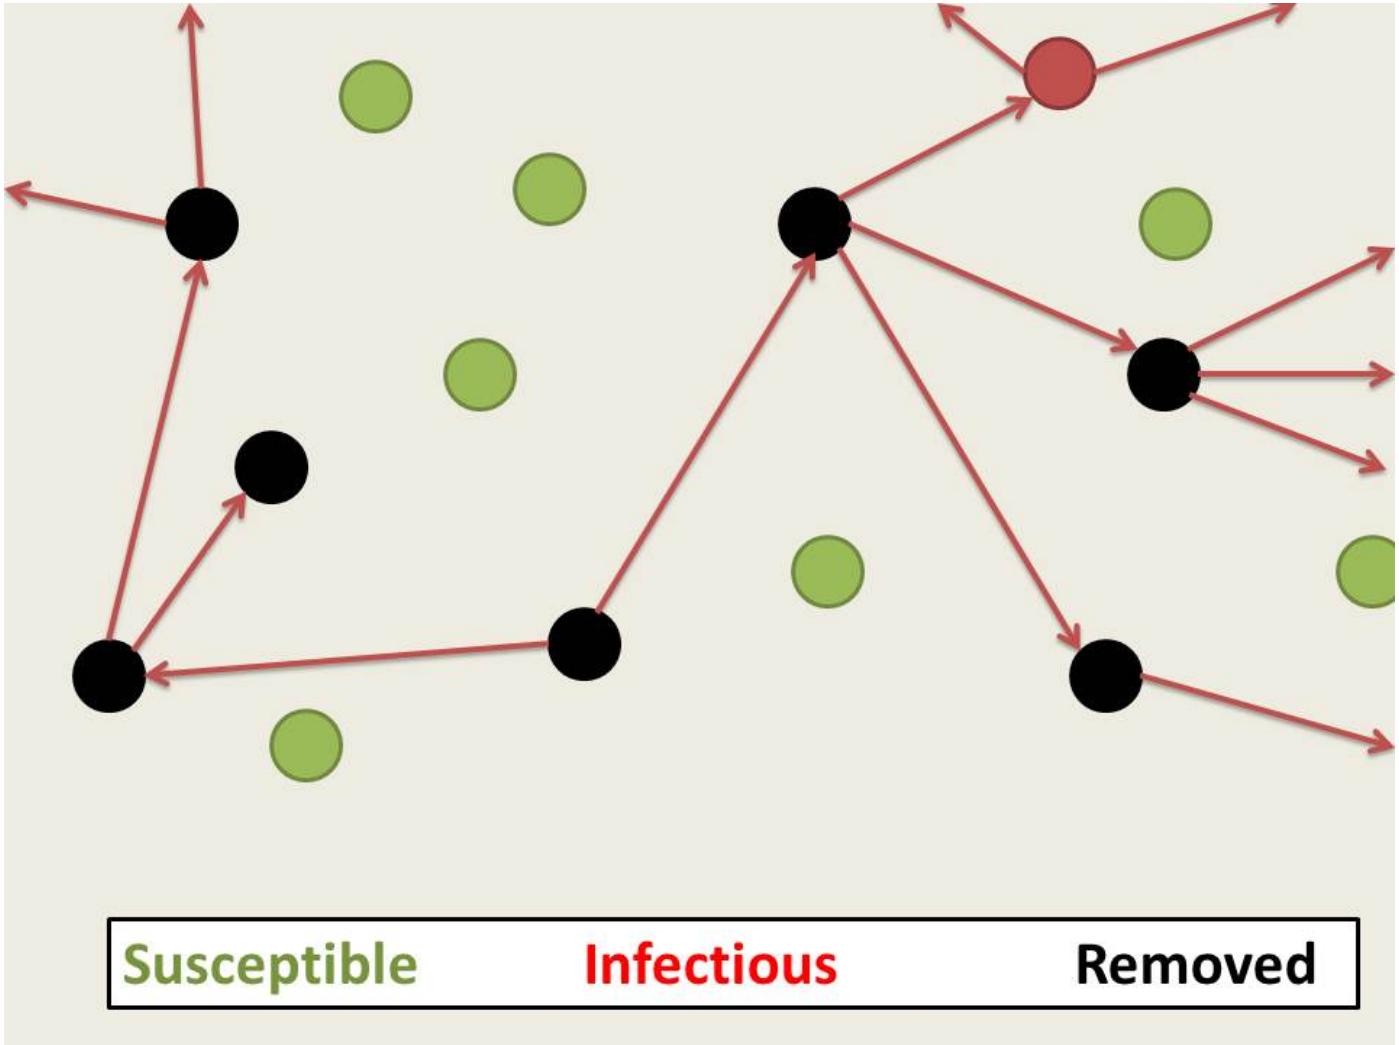

How would an epidemic get started anyway? Example 2

An infectious host is removed.

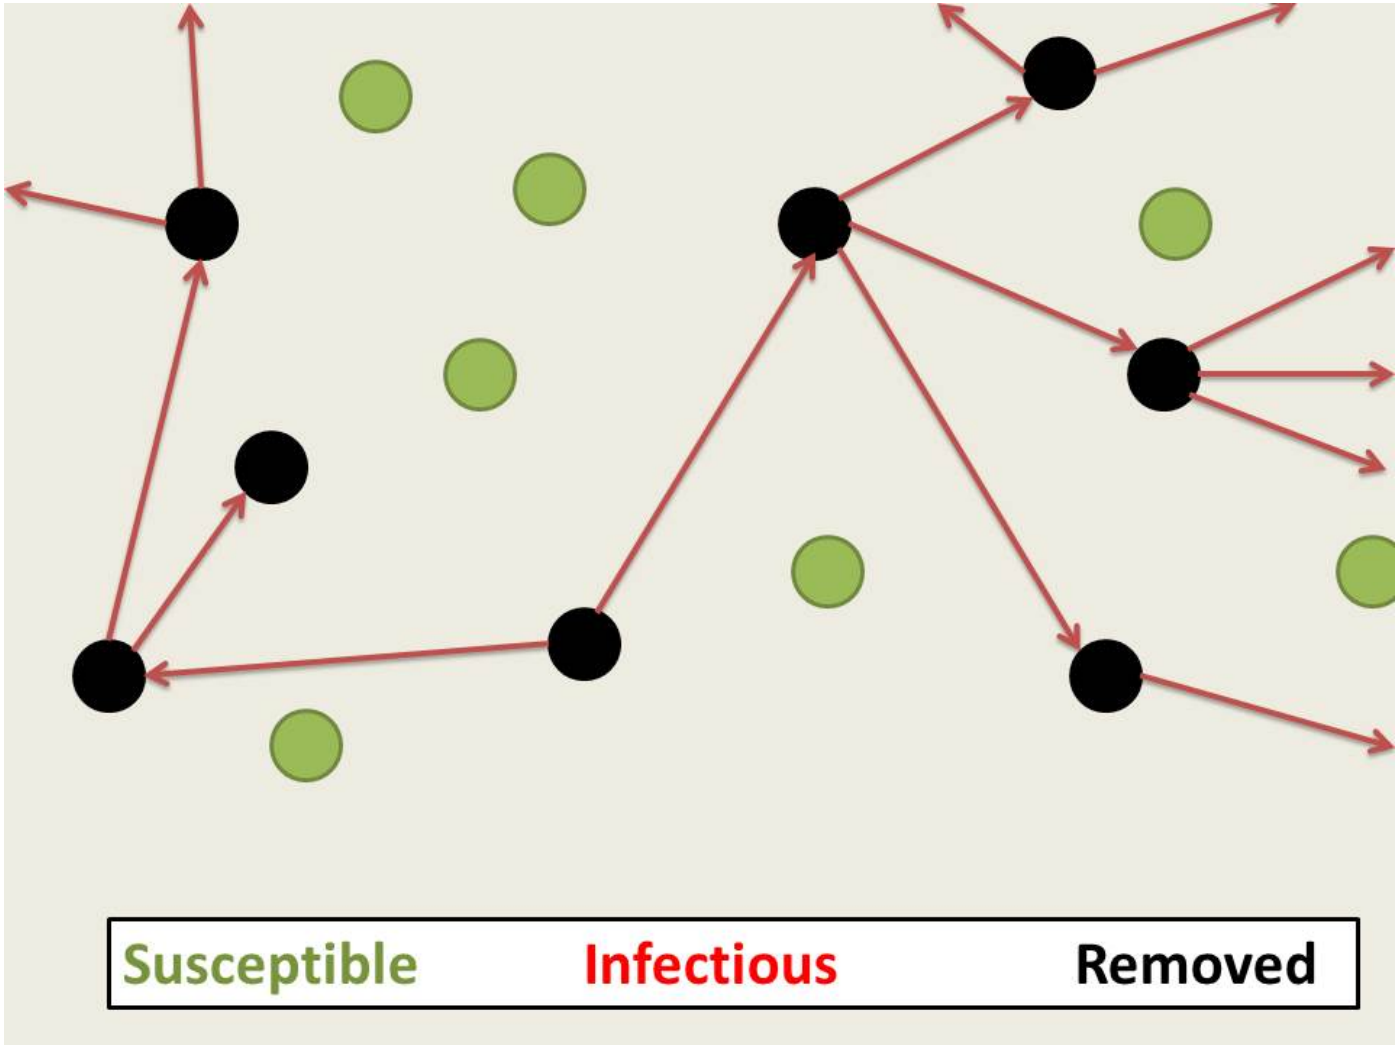

Some generations of the infection in Example 2

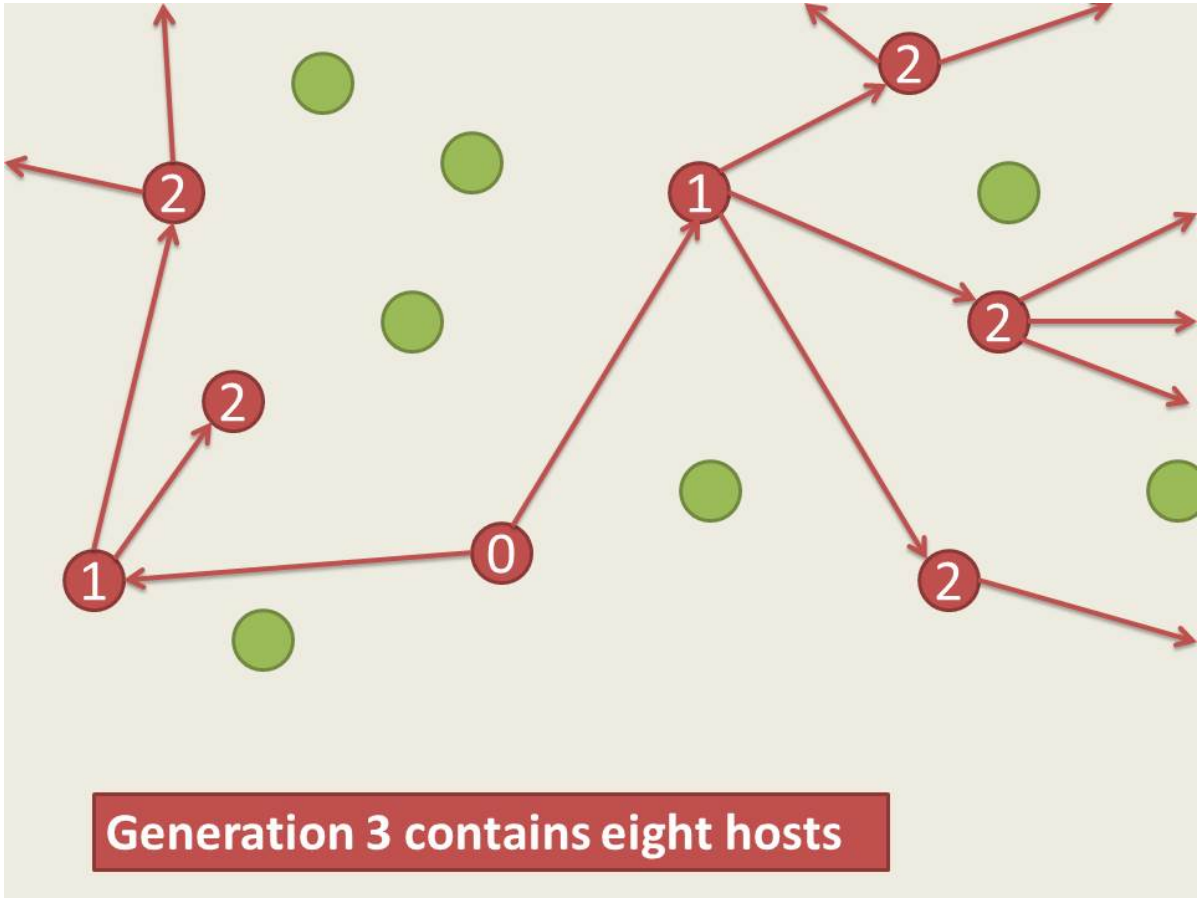

The average number of secondary infections per infectious host in generations 0 to 2 in this example is $\frac{2+2+0+2+3+2+3+1}{8} = \frac{15}{8}$.

Does Example 2 indicate the start of an epidemic?

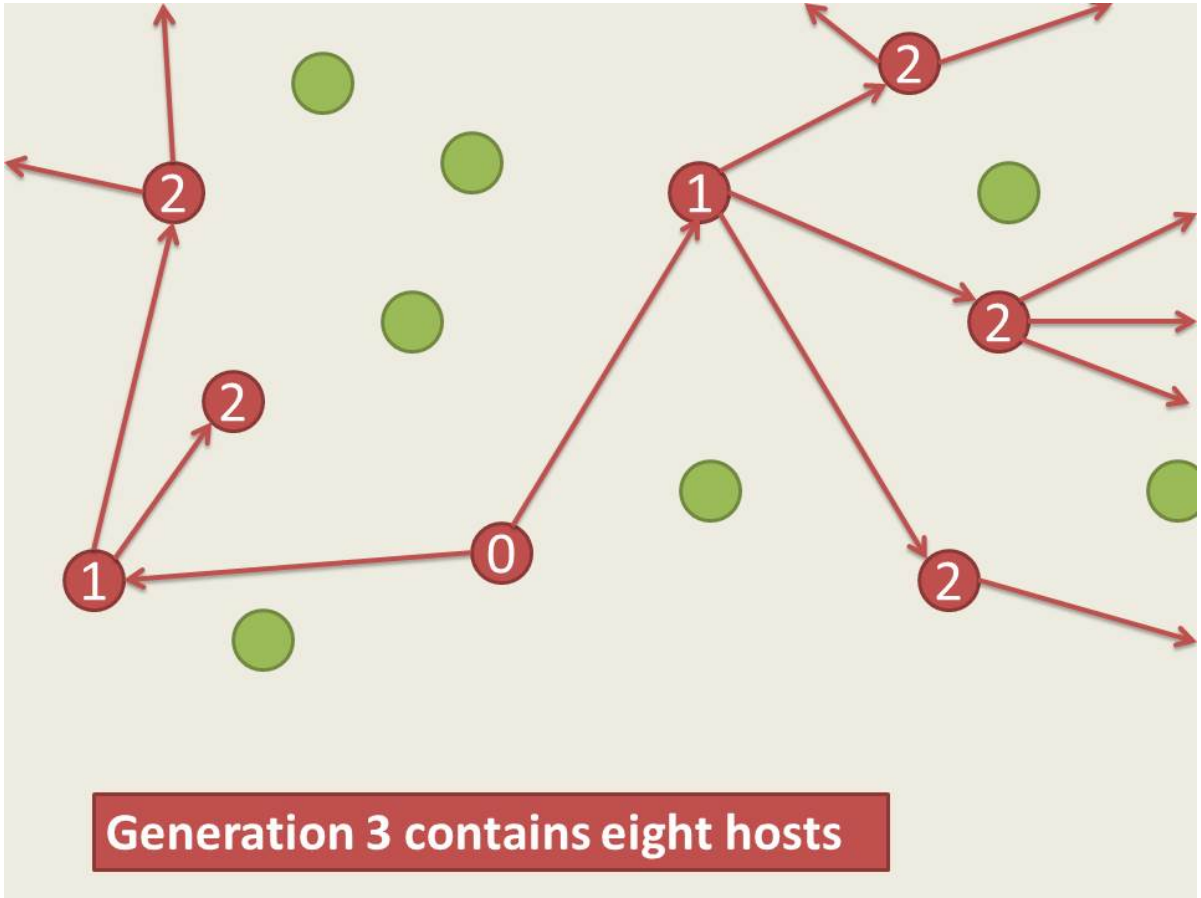

Most likely. From generation 0 to generation 3 the number of infectious hosts has increased by a factor of 8, and one might expect similar increases in subsequent generations.

What makes the difference?

Definition

The expected number of secondary infections that will be caused by a single infectious host that is introduced into a large and entirely susceptible population is denoted by R_0 and called the basic reproductive ratio or basic reproductive number.

If $R_0 \ll N$ and if we assume uniform mixing of the population, then practically contacts of infectious hosts during the first few generations will be with susceptibles, and we might assume, as long as k is sufficiently small, that $R_0 \approx R_k$, where R_k denotes the mean number of secondary infections caused by a host in the k-th generation.

Based on this argument, our best guess at R_0 would be $R_0 \approx \frac{5}{6} < 1$ in Example 1 and $R_0 \approx \frac{15}{8} > 1$ in Example 2.

R_0 makes the difference

Theorem

Assume uniform mixing and introduction of a single infected individual into an entirely susceptible population. Assume, moreover, that R_0 does not depend on N.

If $R_0 < 1$, then the expected number of individuals that eventually become infected is bounded by a constant that depends only on R_0 but not on N, and the disease is predicted to quickly die out.

If $R_0 > 1$, then with probability > 0 an epidemic whose final size is at least a fraction $F(1, N) > 0$ that depends only on R_0 will occur.

"Proof": Under the assumption the expected number of infecteds in generation k satisfies $E(g_k) = R_0 R_1 \dots R_{k-1} \leq R_0^k$, since $R_k \leq R_0$. Thus $E(\lim_{t\to\infty} N - S(t)) \leq \sum_{k=0}^{\infty} R_0^k = \frac{1}{1-R_0}$.

If $R_0 > 1$, then $E(g_k) \approx R_0^k$ for small k. More generally, $R_k \geq 1$ until a significant fraction of susceptibles move to the I- or R-compartments; an epidemic will occur with positive probability.

Which factors might determine R_0 and, more generally, the disease dynamics?

- The pattern of mixing between susceptibles and infectives. We would like to know, for a given susceptible host and a given time interval, the probability distribution of the number and intensity of contacts that this host will have with infectives.
- The probability that a given contact between a susceptible and infective individual at time $T_I^i + t$ results in a "successful" (from the point of view of the disease agent) transmission.
- The distribution of times $T_I^i - T_E^i$ during which a host resides in E and $T_R^i - T_I^i$ during which a host resides in I.

This may be too much to ask (the biologists) for.

Can you think of a population of real hosts for which it would be possible to collect all the relevant data?

Even if we could have all the data, the resulting model would likely be intractable. We need to make simplifying assumptions.

We have seen extreme simplifications

The compartment-based ODE models that we saw earlier distill all these features into two parameters α and β. This may be too extreme. In particular:

- ODE models ignore the stochastic nature of disease transmission.
- Compartment models ignore heterogeneities between individual hosts.
- Compartment models are based on the often unrealistic assumption of uniform mixing between individual hosts.

Let us now try to develop a modeling framework that is capable of incorporating as many potentially relevant details as possible.

Stochastic process models: The basics

Examples 1 and 2 suggest that stuff happens at random times T_I^i of infection and T_R^i of removal of host number i.
One can conceptualize disease dynamics as a stochastic process that moves hosts around between the compartments.

Let us assume a fixed population that consists of hosts that are represented by variables $x_i(t)$, where $i \in \{1, \ldots, n\}$.

At any given time, a r.v. $x_i(t)$ can take values $x_i(t) \in \{S, I, R\}$, depending on the relation of t to T_I^i and T_R^i.

The state of the population at time t is the vector $\vec{x}(t) = (x_1(t), \ldots, x_N(t))$.

What have we swept under the rug so far?

- We ignore demographics, that is births, deaths from unrelated causes, immigration and emigration.
- We assumed $T_I^i = T_E^i$, that is, the time of exposure coincides with the onset of infectiousness.
- Successful transmission is a discrete event that either does or does not happen during a given contact between an infective and susceptible host. This ignores the possibility of multiple below-threshold exposures adding up to an infection.

How distorting are these simplifying assumptions likely to be?

How could we incorporate the ignored details into our model?

Independence of transition times

A state can change only by a variable x_i changing its state from S to I (at time T_I^i) or from I to R (at time T_R^i).

We will assume that for any given state $\vec{x}(t)$ the relevant variables T_I^i and T_R^i are all independent.

How realistic is this assumption?

The Markov Property

We want our stochastic process model to be reasonably tractable; the Markov Property might help. In other words, given a state $\vec{x}$, we want the conditional distribution of future states $\vec{x}(t + \Delta t)$ given that $\vec{x}(t) = \vec{x}$ to depend only on $\vec{x}$ and Δt. One advantage of the Markov Property is that it allows in some cases for approximations of the model by autonomous ODEs.

The Markov Property implies that for a given state $\vec{x}$ each of the relevant variables $T_I^i = T_I^i(\vec{x})$ and $T_R^i = T_R^i(\vec{x})$ is memoryless.

Thus T_I^i will be exponentially distributed with parameter $\beta_i(\vec{x})$ and T_R^i will be exponentially distributed with parameter $\alpha_i(\vec{x})$.

Surprise: We have built a model!

Now assume you can determine values for the parameters $\alpha_i(\vec{x})$ and $\beta_i(\vec{x})$ for all possible states $\vec{x}$ of the population. Then the above assumptions specify a stochastic model. Let's show how to simulate the process on the computer.

- Choose an initial state $\vec{x} := \vec{x}(0)$.
- For all i with $x_i = I$, randomly and independently choose times T_R^i at which host i will move out of the I-compartment according to an exponential distribution with parameter $\alpha_i(\vec{x})$.
- For all i with $x_i = S$, randomly and independently choose times T_I^i at which host i would move into the I-compartment if the state were to remain unchanged, according to an exponential distribution with parameter $\beta_i(\vec{x})$.
- Determine the smallest time t_{next} at which the next "event" (movement of a host to a different compartment) happens.
- Set $\vec{x} := \vec{x}(t_{next})$ accordingly.
- Repeat until stopping criterion.

Advantages of the supermodel

We can think of the construction we have just presented as a supermodel. It has a lot of attractive features:

- It accounts for the stochastic nature of disease transmission.
- It can be easily explored by simulations for moderately large N.
- It allows for heterogeneities between individual hosts (think of $\alpha_i(\vec{x})$ as reflecting individual strength of the immune system).
- It allows for exploring a variety of mixing patterns between individual hosts (since $\beta_i(\vec{x})$ in general may depend on $\vec{x}$, that is, on which other hosts are infectious at a given time).

Can you think of a scenario where we would want $\alpha_i(\vec{x})$ to actually depend on $\vec{x}$?

A drawback of the Markov Property

The Markov Property may impose some plainly unrealistic features though.

For example, the assumption that $T_R^i - T_I^i$ is a memoryless r.v. is blatantly wrong for most diseases. Recovery times show usually a distribution that peaks at some modal value. For example, you are much more likely to recover during day 7 of a bout of the flu than during day 2, while an exponential distribution would predict the opposite.

How can we modify the model so that the distribution of recovery times becomes more realistic without sacrificing the Markov Property of the process?

Beware of supermodels!

Supermodels are pleasant to contemplate, but notoriously difficult to work with.

Here are some problems ours suffers from:

- There are just too many parameters. For an *SIR*-model we would need $2N3^N$ parameters $\beta_i(\vec{x})$ and $\alpha_i(\vec{x})$.
- Even for relatively small N it is plainly impossible to estimate that many parameters from any kind of data.
- The dimension N of the model is too large to study it analytically.

Can we reduce the number of parameters that need to be estimated from the data?

Let's make some tough choices

- Let us ignore heterogeneities in individual immune response and set $\alpha_i(\vec{x})$ to a fixed α for all i and $\vec{x}$.
- There is only one type of contact, and contacts last an instant rather than having a duration. The probability of a "successful" transmission of the disease during a given contact between an infectious and a susceptible host is a fixed parameter p.
- The time $\tau_{i,j}$ that host number i has to wait for the next contact with host number j after time t is a memoryless r.v. and thus has an exponential distribution with some parameter $\lambda_{i,j}$ that does not depend on $\vec{x}$.

This still allows us to explore a variety of mixing patterns and reduces the number of parameters from a stratospheric $2N3^N$ to a still lofty but more reasonable $\binom{N}{2} + 2$.

What are some potential problems with these assumptions and how could we address them?

The new parameters specify a model

- $\alpha_i(\vec{x}) = \alpha$ for a fixed α for all i and $\vec{x}$.
- The probability of a "successful" transmission of the disease during a given contact between an infectious and a susceptible host is a fixed parameter p.
- The time $\tau_{i,j}$ that host number i has to wait for the next contact with host number j has an exponential distribution with some parameter $\lambda_{i,j}$.

We need to convince ourselves that the new parameters p and $\lambda_{i,j}$ suffice to specify $\beta_i(\vec{x})$ for all i and $\vec{x}$.

For sufficiently small Δt the probability that a successful transmission from infectious host number j to susceptible host number i occurs in the interval $[0, \Delta t]$ can be approximated as $P(\tau_{i,j} \leq \Delta t)p \approx p\lambda_{i,j}\Delta t$,
and the probability that susceptible host number i will become infected during this time interval can be approximated as $P(T_I^i(\vec{x}) < \Delta t) \approx \sum_{\{j:\ x(j)=I\}} p\lambda_{i,j}\Delta t = \beta_i(\vec{x})\Delta t$.

The uniform mixing assumption

The uniform mixing assumption translates into

$\lambda_{i,j} = \lambda$ for some fixed constant λ,

or, equivalently, into

$\beta_i(\vec{x}) = p\lambda \#\{j : \ x(j) = I\} = \beta \#\{j : \ x(j) = I\}$

for the fixed constant $\beta = p\lambda$.

One can also interpret $\beta \#\{j : \ x(j) = I\}$ as the rate at which susceptible host number i acquires an infection; it is called the force of infection.

For large N, this allows us to approximate our stochastic process models by the ODE models that we encountered earlier.

Basic ODE versions of the *SI*-, *SIR*- and *SIS*-models

***SI*-model:**

$$\frac{dS}{dt} = -\beta SI$$
$$\frac{dI}{dt} = \beta SI$$

***SIS*-model:**

$$\frac{dS}{dt} = -\beta IS + \alpha I$$
$$\frac{dI}{dt} = \beta IS - \alpha I$$

***SIR*-model:**

$$\frac{dS}{dt} = -\beta IS$$
$$\frac{dI}{dt} = \beta IS - \alpha I$$
$$\frac{dR}{dt} = \alpha I$$

The uniform mixing assumption

The uniform mixing assumption translates into

$\lambda_{i,j} = \lambda$ for some fixed constant λ,

or, equivalently, into

$\beta_i(\vec{x}) = p\lambda \#\{j : \ x(j) = I\} = \beta \#\{j : \ x(j) = I\}$

for the fixed constant $\beta = p\lambda$.

We get a reduction to only two parameters.

But when would the the uniform mixing assumption be realistic?

Mixing may be nearly uniform if hosts move around a lot relative to the size of the habitat, encounter each other rarely, and there is no social structure.

How to model more realistic mixing patterns?

In populations with a well-defined social or territorial structure though, some pairs of individuals will have contact relatively frequently (think of co-workers or neighbors in human populations), while other pairs of individuals will almost certainly never encounter each other (think of your likelihood to ever meet the Supreme Leader of North Korea).

We can approximate the latter situation by assuming the existence of a contact network which determines whether it is even possible that the disease can be transmitted between two given hosts.

The nature of the required contact, and thus the relevant contact network, may depend on the particular disease. Think of the flu *vs.* a computer virus *vs.* a sexually transmitted disease.

Mathematical structures for modeling contact networks: graphs

A graph is an ordered pair $G = (V, E)$, where V denotes the set of vertices, or nodes, and the set E of edges of G is a subset of the set of unordered pairs of nodes.

A contact network can be modeled as a graph whose vertices are the individual hosts in the population, and an edge between two hosts signifies an above-threshold probability of a relevant contact between these two hosts.

One can then make the simplifying assumption that disease transmission can occur only between two hosts that are represented by adjacent nodes, that is, endpoints of a common edge, and study the possible or likely dynamics of the disease on the network.

An example of a contact network

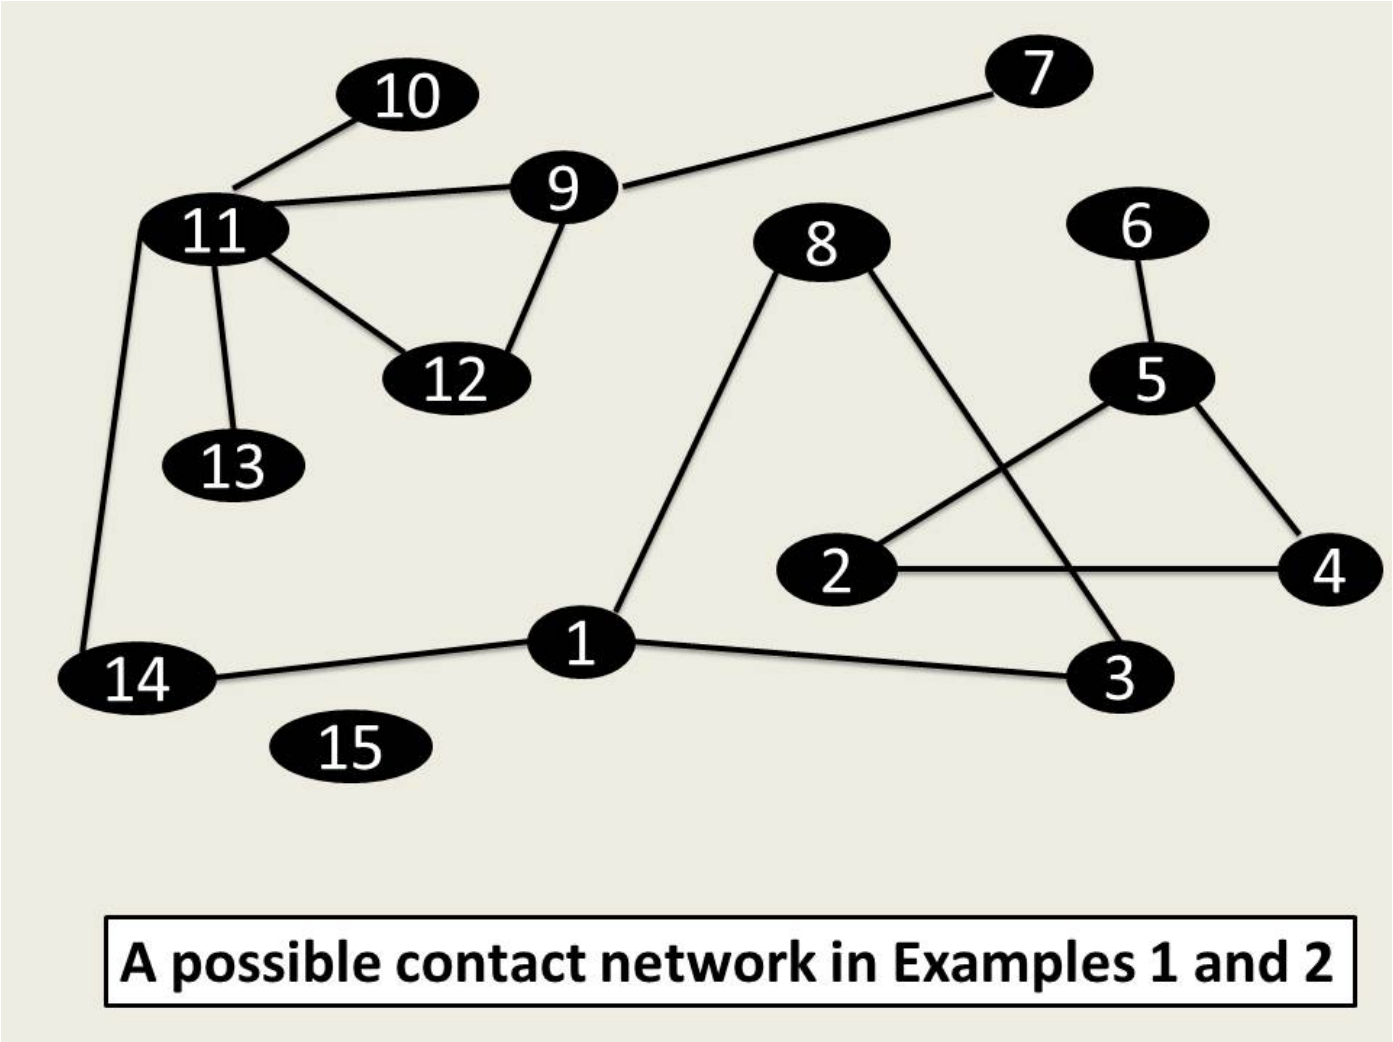

A possible contact network in Examples 1 and 2

Could Example 1 occur on this contact network?

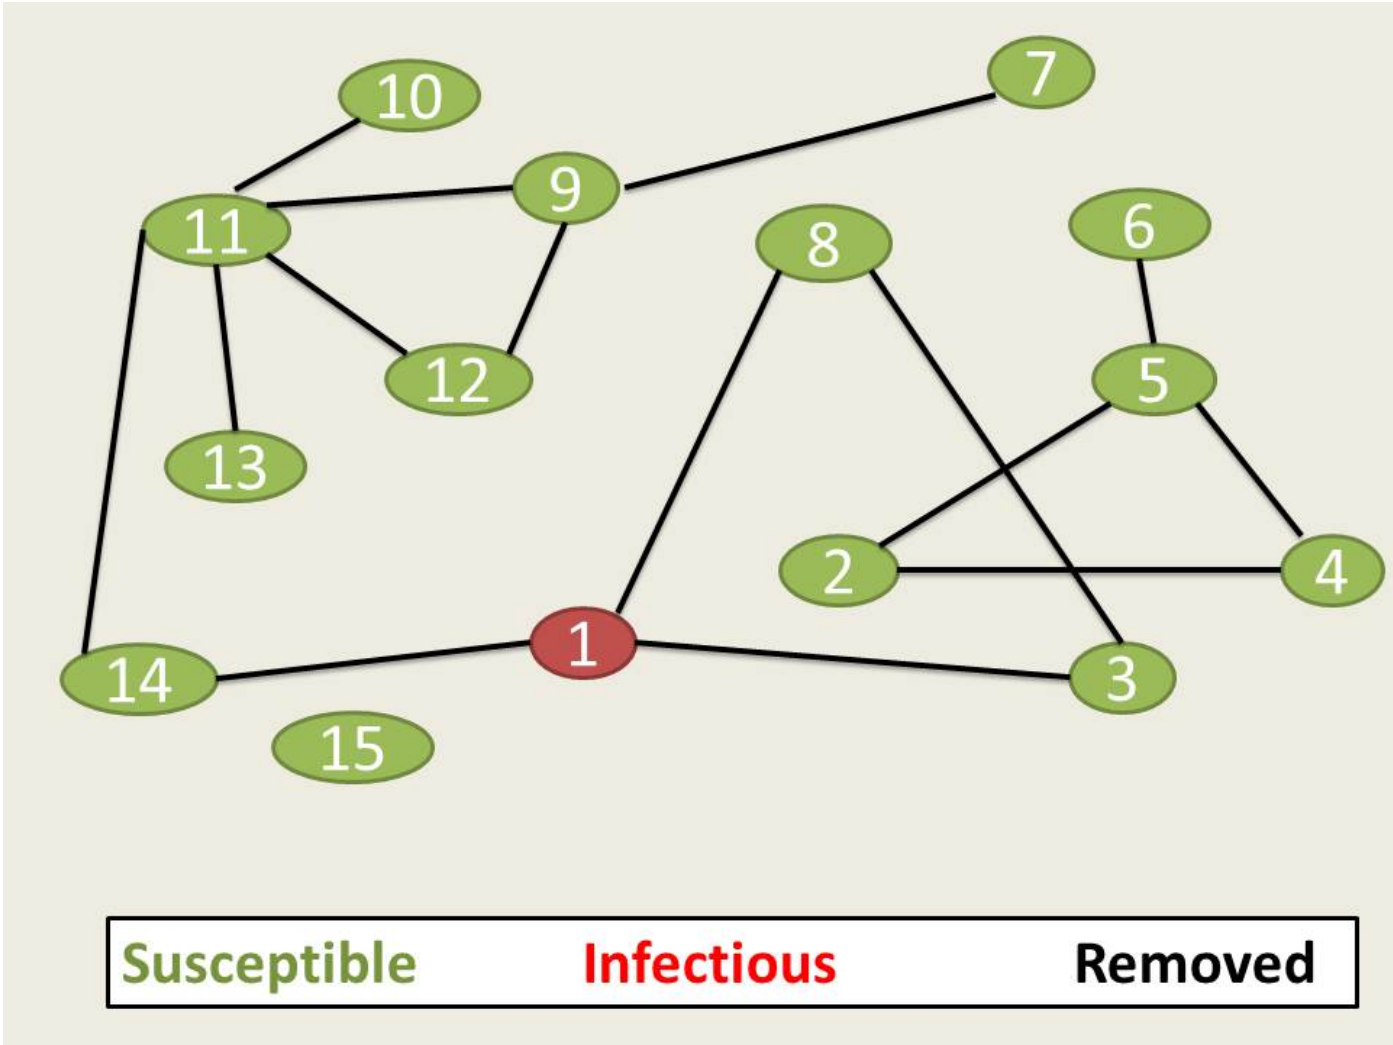

Could Example 1 occur on this contact network?

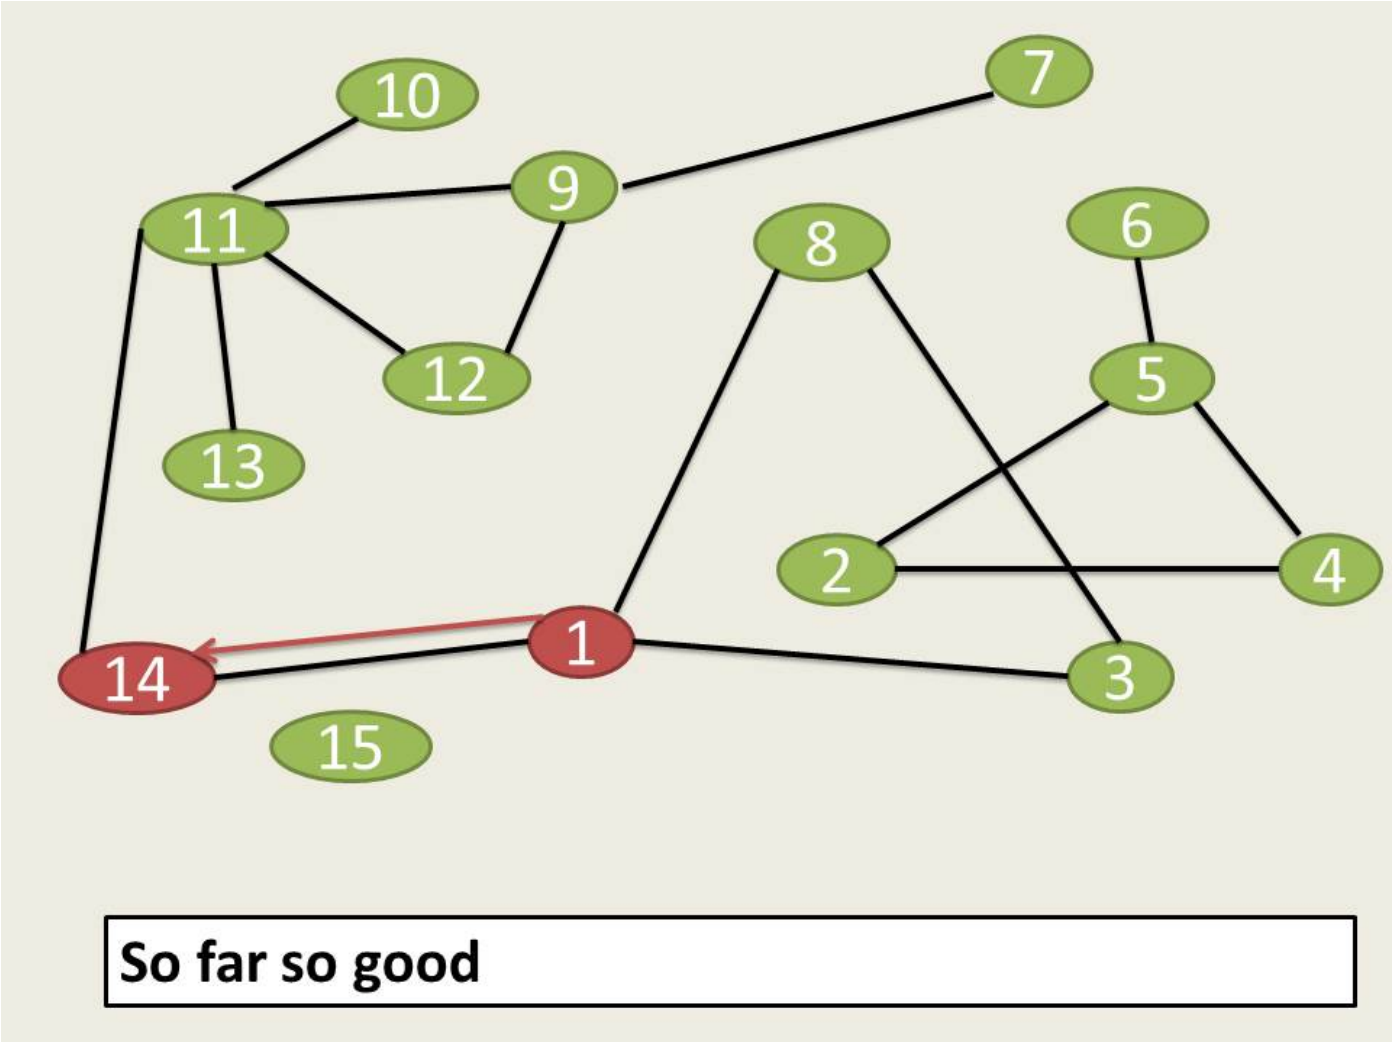

Could Example 1 occur on this contact network?

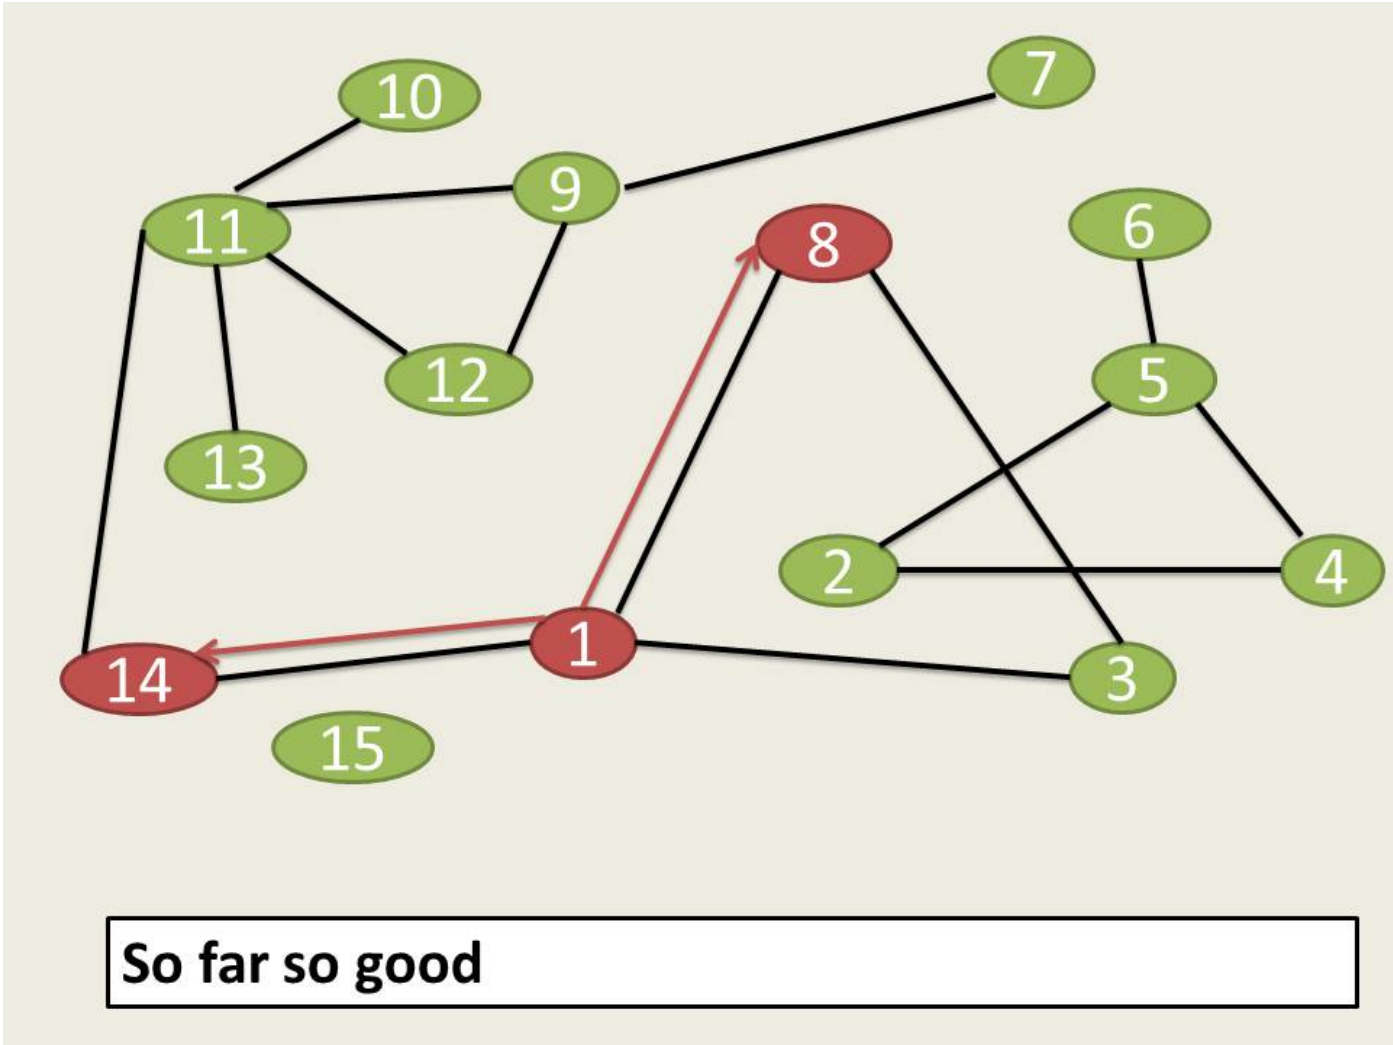

Could Example 1 occur on this contact network?

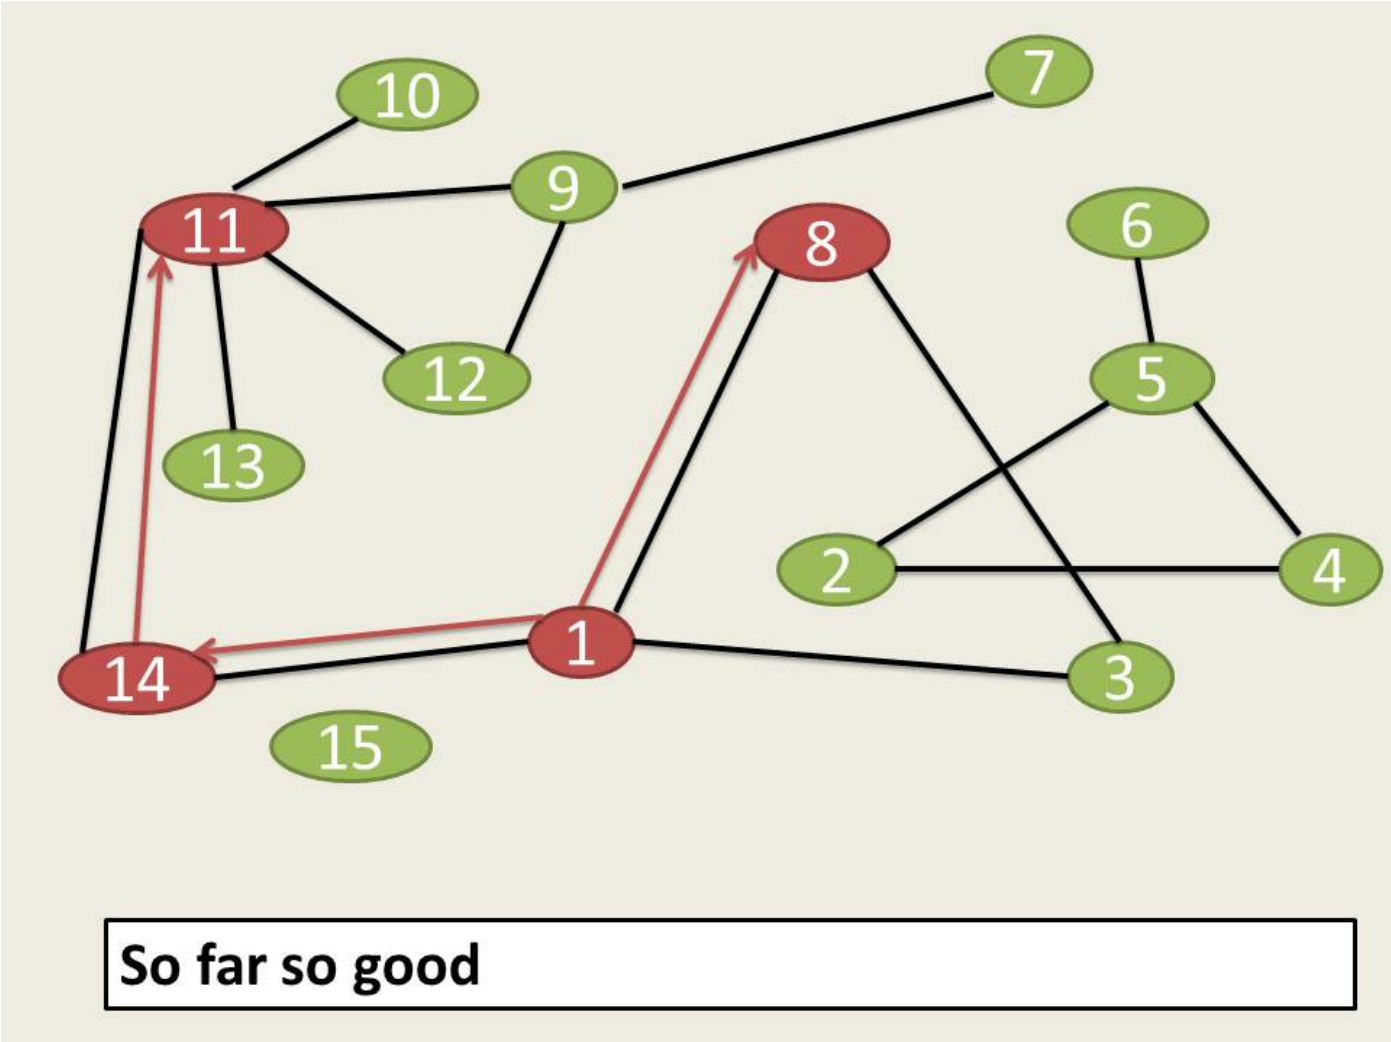

Could Example 1 occur on this contact network?

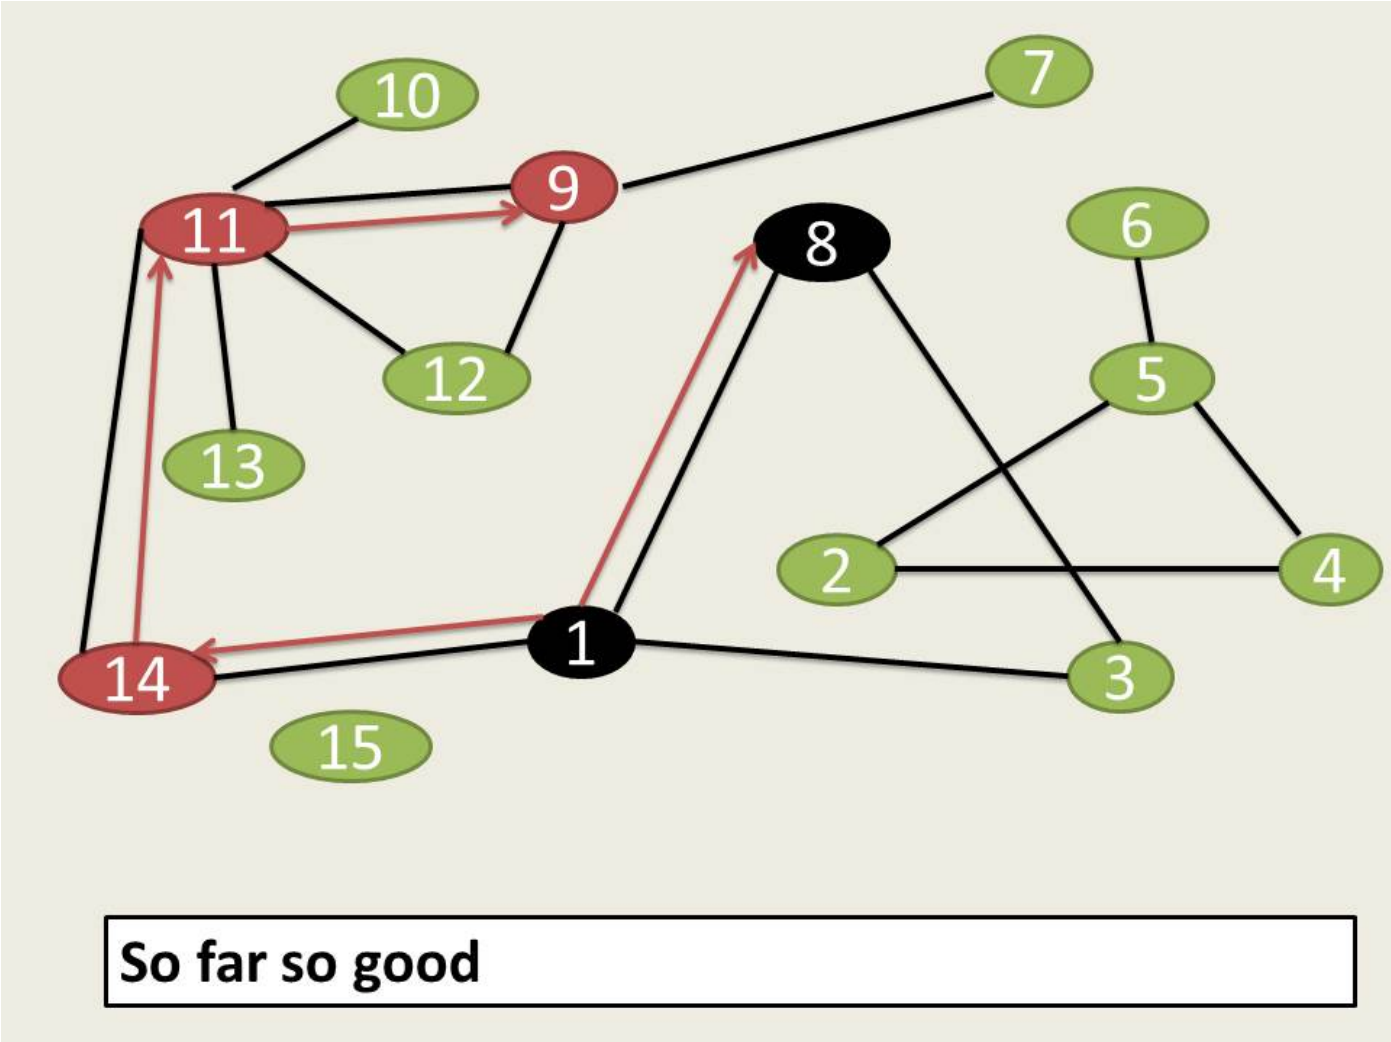

Could Example 1 occur on this contact network?

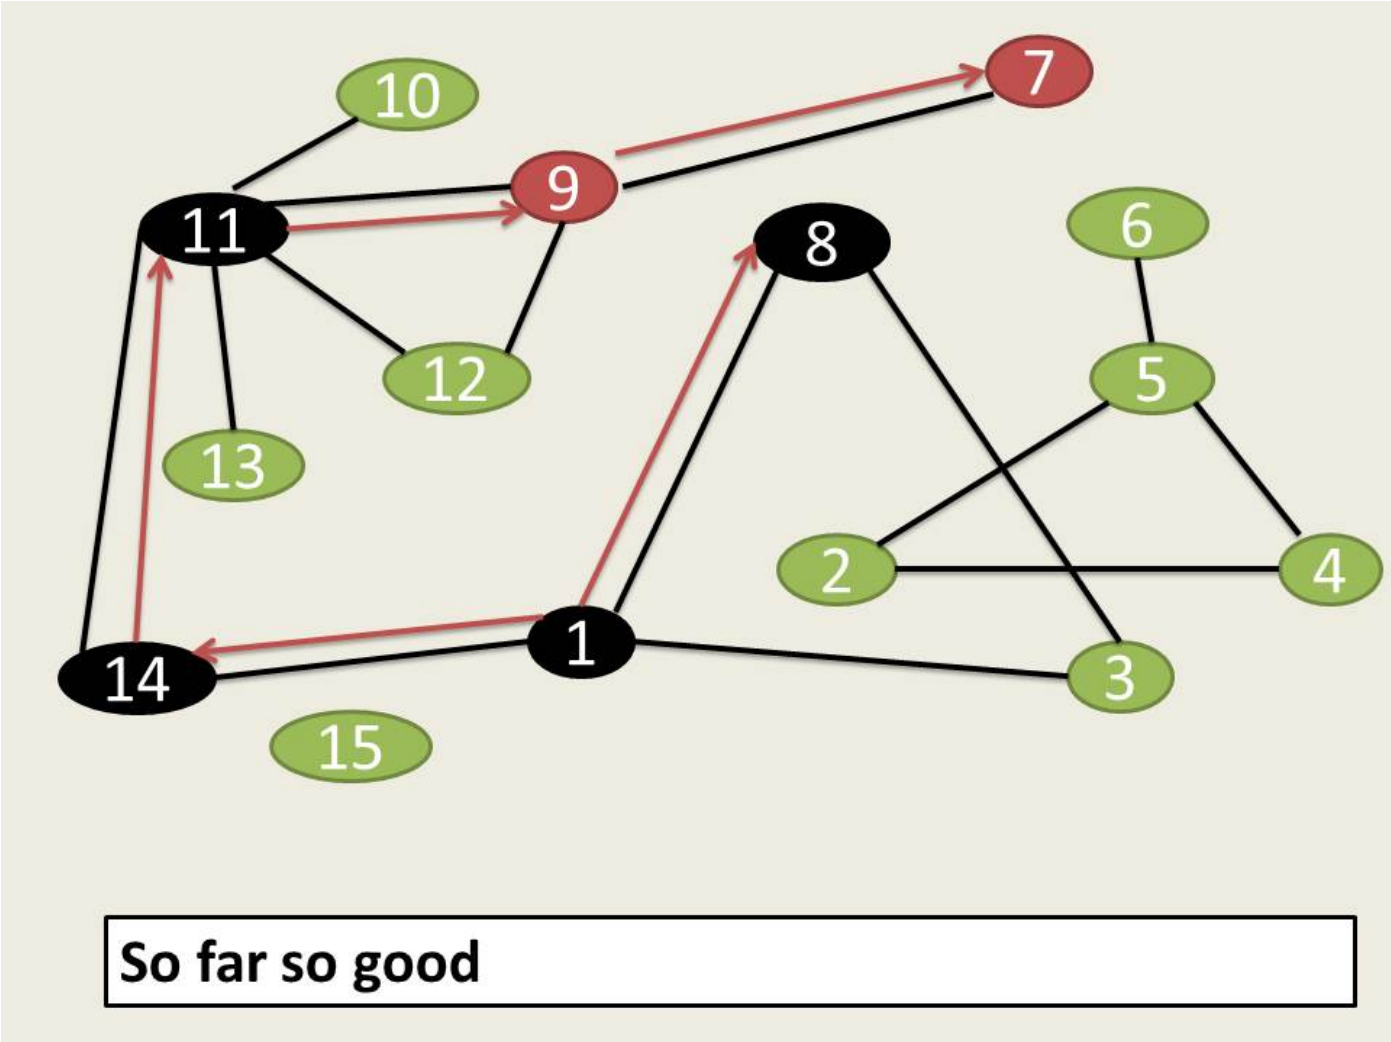

Could Example 1 occur on this contact network?

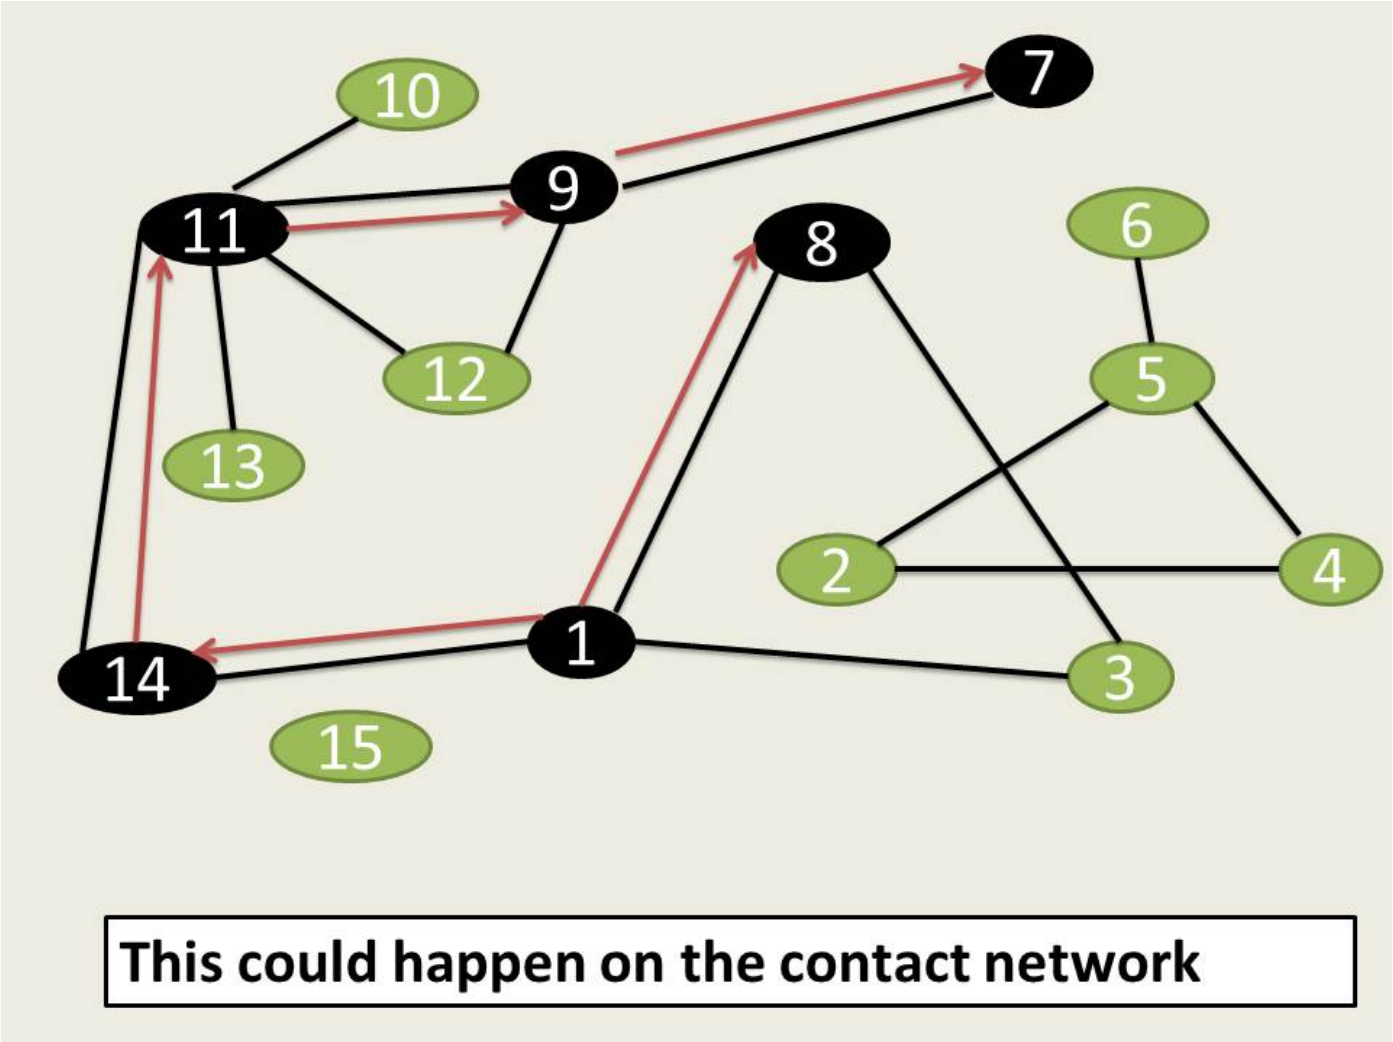

Could Example 2 occur on this contact network?

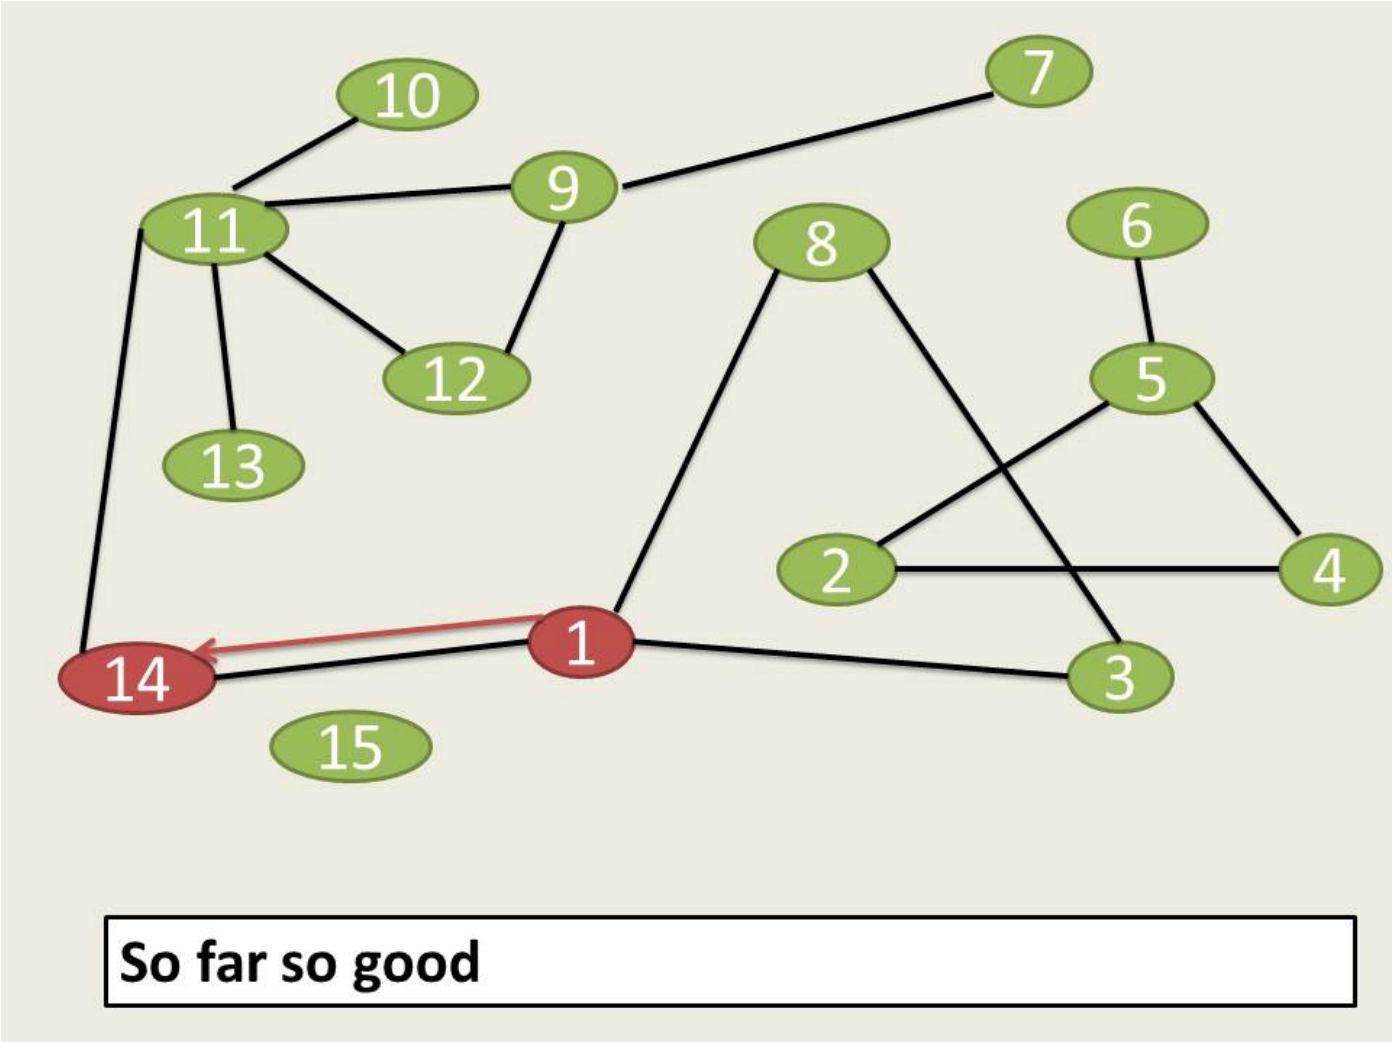

Could Example 2 occur on this contact network?

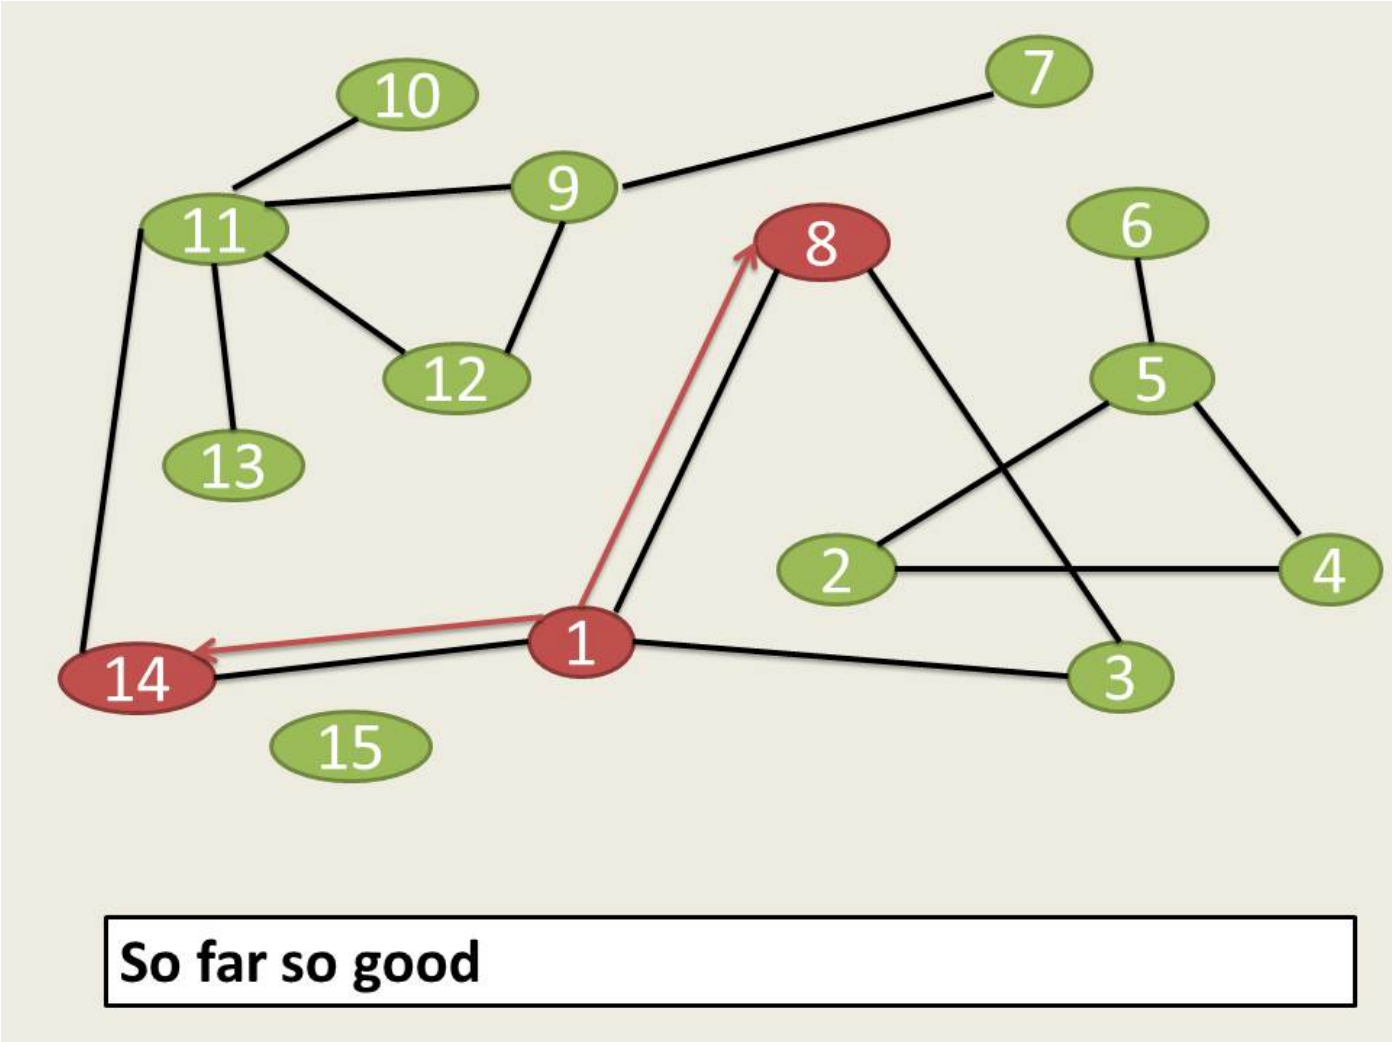

Could Example 2 occur on this contact network?

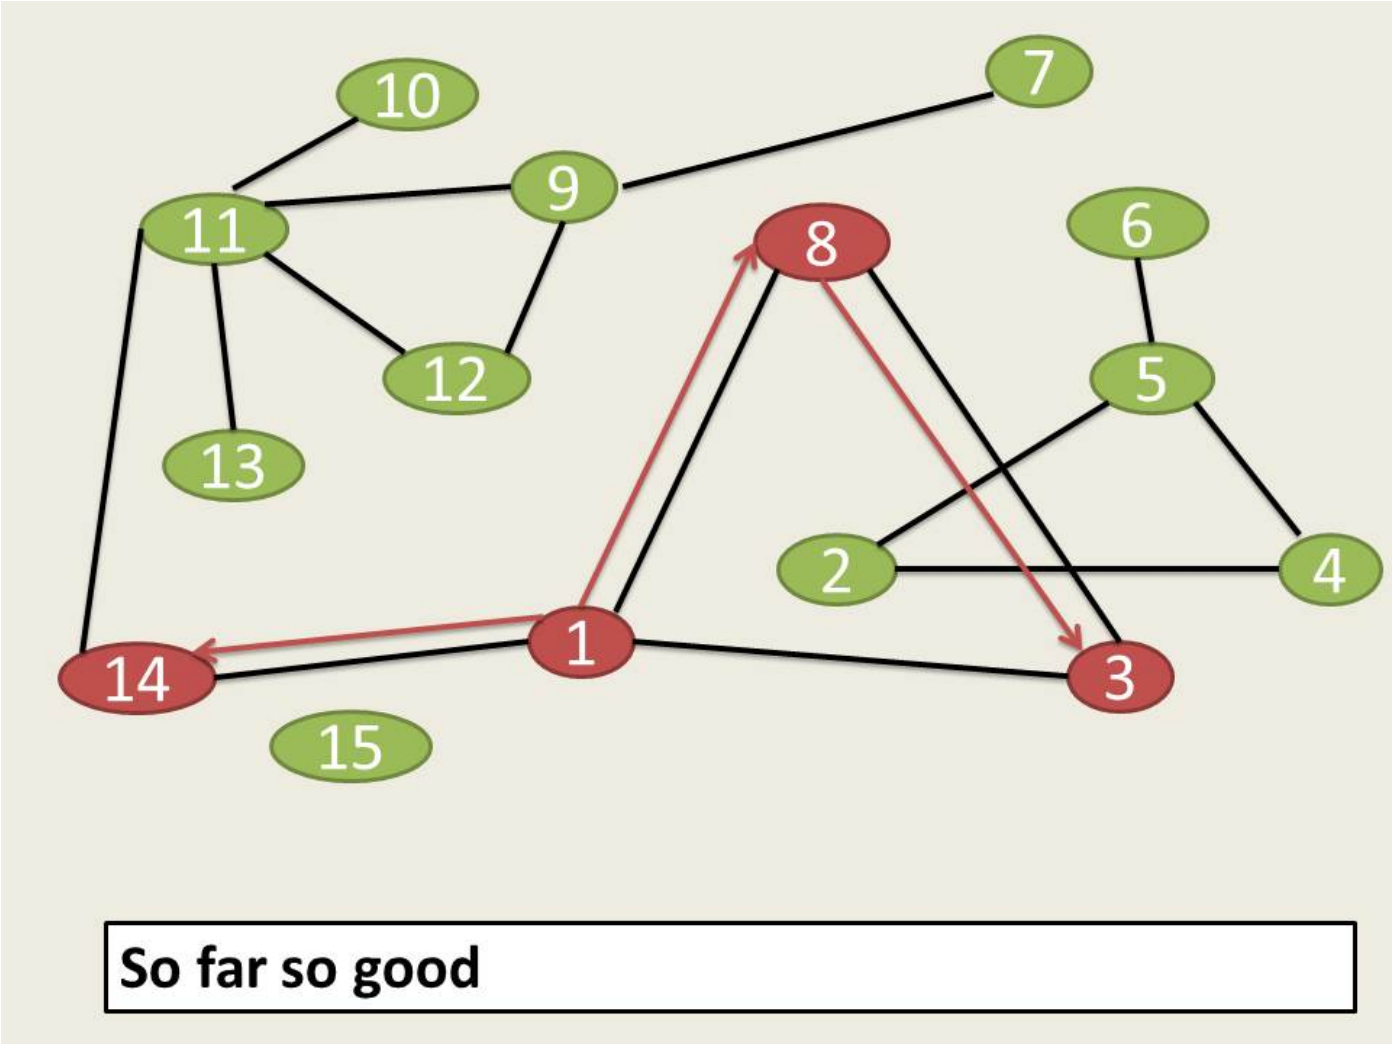

Could Example 2 occur on this contact network?

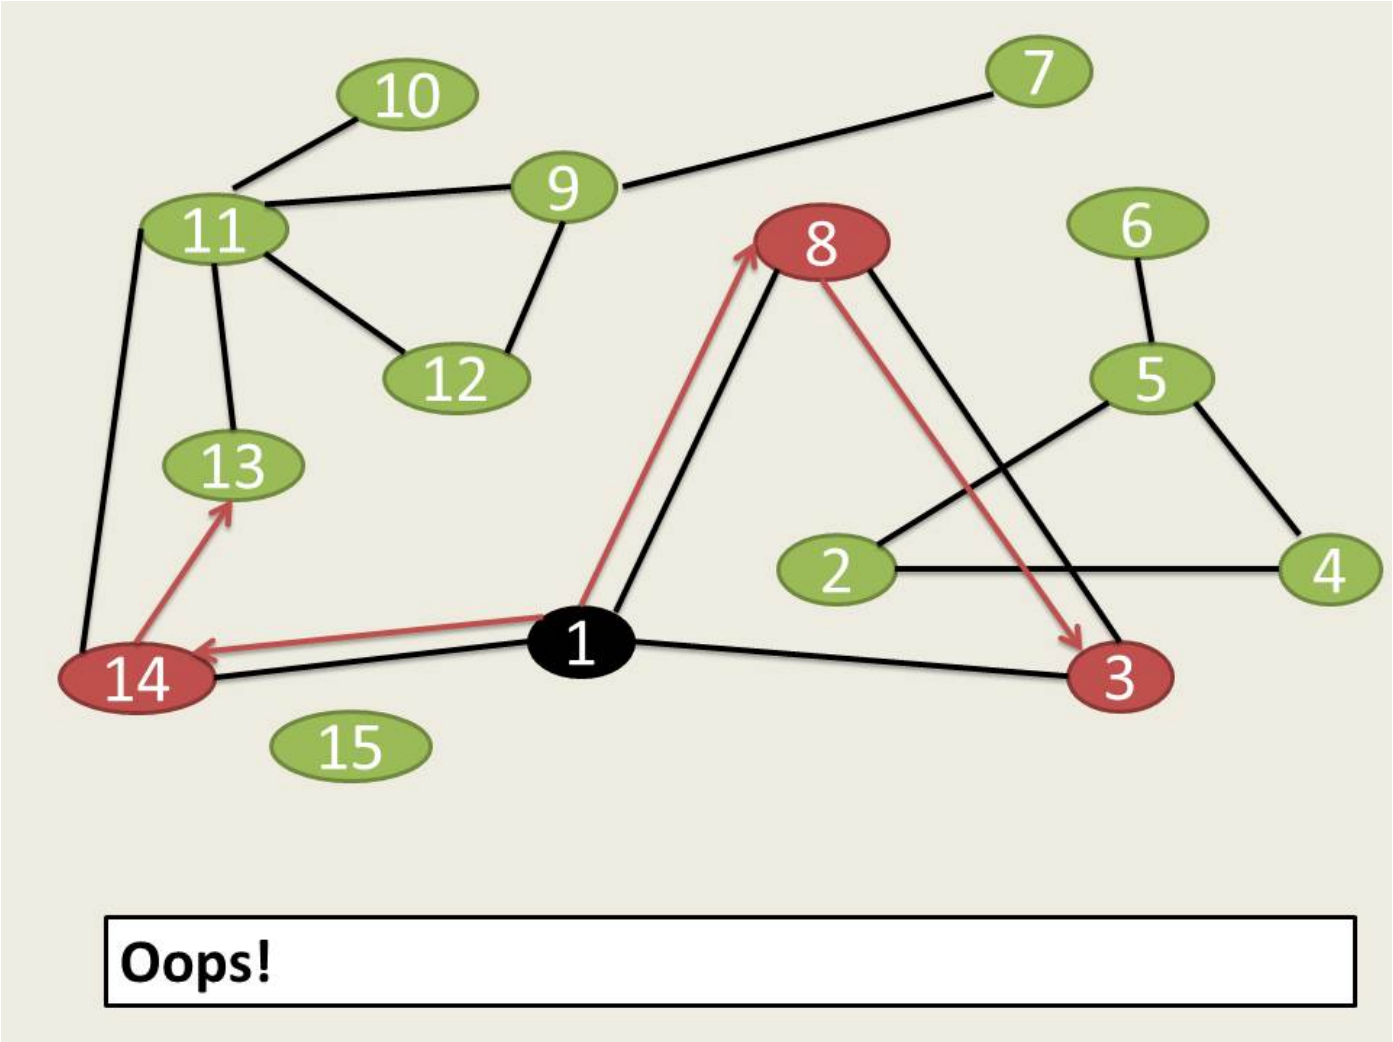

Example 2 could not occur on this contact network

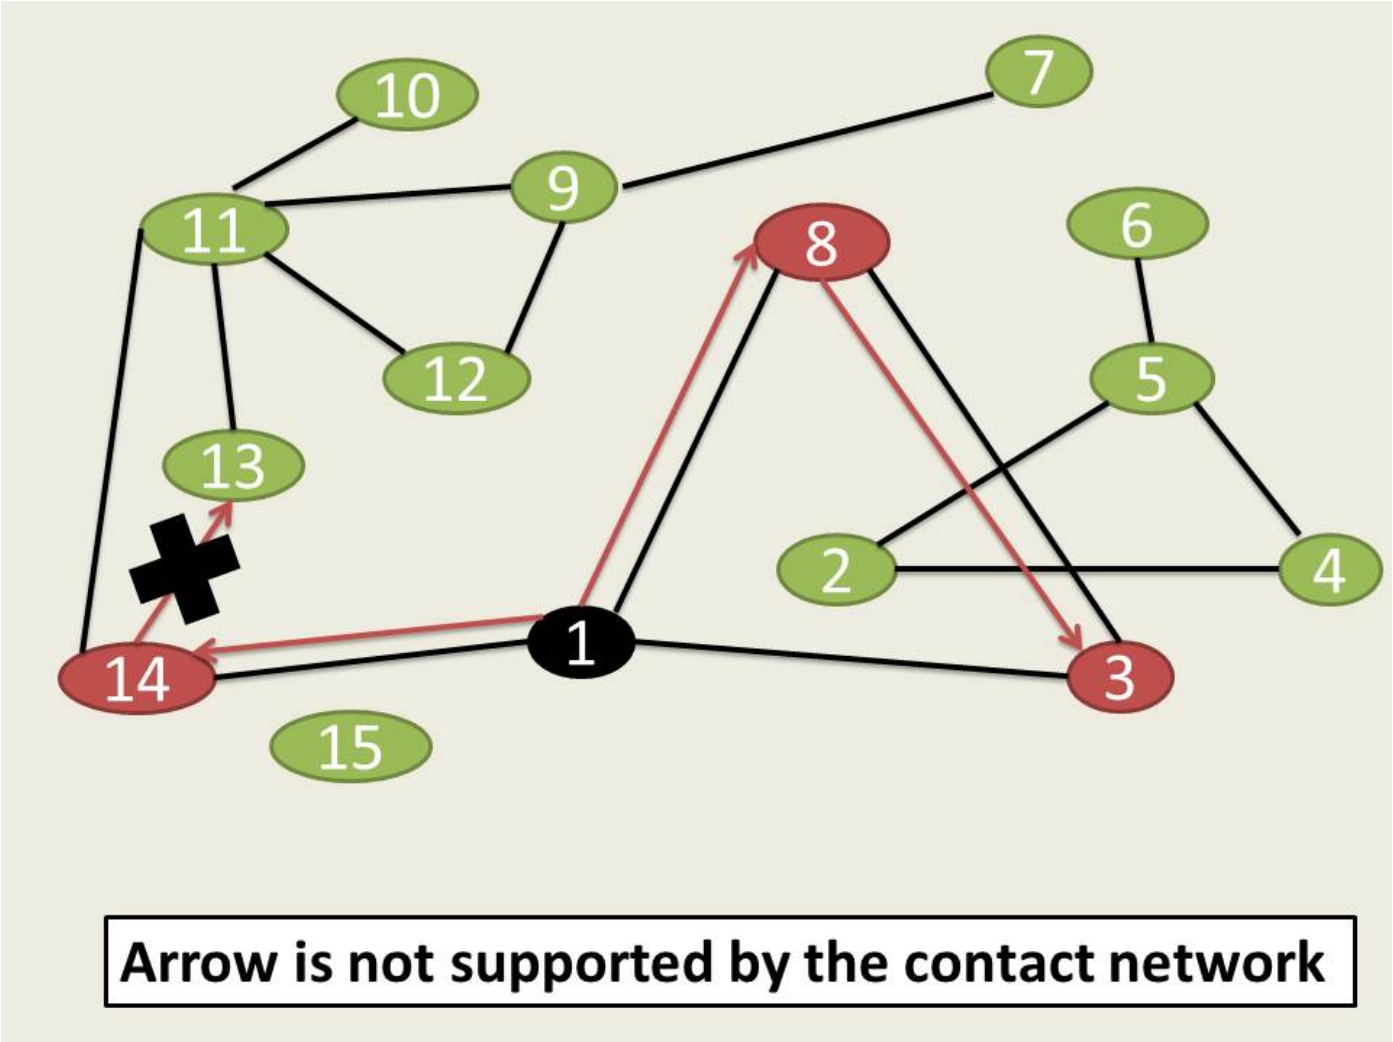

One last simplifying assumption

Let us assume that for some fixed constant λ we have $\lambda_{i,j} = \lambda$ whenever $\{i, j\}$ is an edge in the contact network, and $\lambda = 0$ otherwise.

This assumption strictly speaking does not reduce the number of parameters, but it allows us to define stochastic process models purely in terms two real parameters α and β and the contact network, which is a discrete structure.

Stochastic process models for disease transmission on networks

Ingredients:

- Specification of the type of model (*SI*, *SIR*, or *SIS*).
- A graph G with N vertices that represents the contact network.
- A parameter α that represents the removal rate.
- A parameter $\beta = p\lambda$ that specifies the rate at which a given susceptible host acquires infections from a given adjacent infectious host.

The process will then be modeled as described above, with

$$\beta_i(\vec{x}(t)) = \beta \#\{j : x_j(t) = I \ \& \ \{i,j\} \in E\}.$$

But how to model the network?

A major problem is that we usually have only very limited knowledge of the actual contact network. There are basically two ways of building mathematically meaningful models of the underlying networks.

- In some cases the network may have a very special structure that can be determined from data.
- Alternatively, we can assume that the network is randomly drawn from a probability distribution with certain parameters. The values of these parameters should be chosen in such a way that they favor networks with properties that conform to whatever data we have about the actual network.
- Popular choices for the second alternative are Erdős-Renyi random graphs or scale-free networks that can be randomly generated according to the preferential attachment model. In the hands-on part we will explore several types of random graphs.

What is to be gained from network models of disease transmission?

While the kind of network models we defined here are more realistic than compartment models, they still are based on a lot of simplifying assumptions. But they can give us important insights:

- They can make predictions about the probability of an epidemic that are not available from ODE models.
- They may point to features of the contact network that significantly influence the outcome of an epidemic. This gives some guidance about what kind of data we need to collect in order to be able to make reasonably accurate predictions.
- They may allow us to discern cases when a compartment-based model is inadequate or, alternatively, guide our choice of the parameters for compartment-based models.
- They can inform the design of effective control measures when the uniform mixing assumption is inadequate.

Exhibit A: the uniform mixing assumption revisited

The uniform mixing assumption corresponds to the case where G is the complete graph that contains all possible edges.

Assume an SIR model. To estimate R_0, consider a state $\vec{x}(0)$ with exactly one infectious host, number j, and all other hosts being susceptible.

Then the expected number of transmissions over a small time interval of length Δt from host j to a given host i is $\approx \beta \Delta t$, thus the expected overall number of secondary infections caused by host j over this interval is $\approx \beta(N-1)\Delta t$.

This formula applies as long as the interval is contained in $(0, T_R^j)$.

Expected values add up, so if we partition $(0, T_R^j)$ into small subintervals of length Δt each, we can deduce that the expected overall number of secondary infections caused by host j satisfies

$$R_0 \approx \beta(N-1)\Delta t \frac{E(T_R^j)}{\Delta t} = \frac{\beta(N-1)}{\alpha}.$$

When is this approximation valid?

I haven't been very clear about all the assumptions needed for this approximation to be a good one.

We will explore in the hands-on session what extra assumptions are needed here and why.

Exhibit A: compare this with the ODE model

$$\frac{dS}{dt} = -\beta IS$$
$$\frac{dI}{dt} = \beta IS - \alpha I = I(\beta S - \alpha)$$
$$\frac{dR}{dt} = \alpha I$$

If one infectious individual is introduced into an otherwise susceptible population, then $S = N - 1$. For $\beta < \frac{\alpha}{N-1}$, the model predicts a decrease in I; for $\beta > \frac{\alpha}{N-1}$, an initially exponential increase of this variable. Reasoning backwards from the theorem about R_0, we see that $R_0 = \frac{\beta(N-1)}{\alpha}$, the same as for the stochastic process model. This assumes population size expressed by the number of individuals; for units of, say, thousands of individuals we get $R_0 \approx \frac{\beta N}{\alpha}$. R_0 it is often reported in these forms.

But the ODE model predicts that for $R_0 > 1$ an epidemic will always occur for this initial condition, while the stochastic nature of transmission implies that this happens only with a positive probability < 1.

The network model allows us to determine this probability.

Exhibit B: When is R_0 a good predictor?

The inequality $R_0 > 1$ is supposed to predict a positive probability of an epidemic, with initially exponential increase in the number of infected hosts.

This theorem was based on the uniform mixing assumption.

When does this prediction fail?
If it fails, what may be a better predictor?

We will explore this issue with random networks and the following two special networks.

Special case 1: $L \times M$ rectangular grids

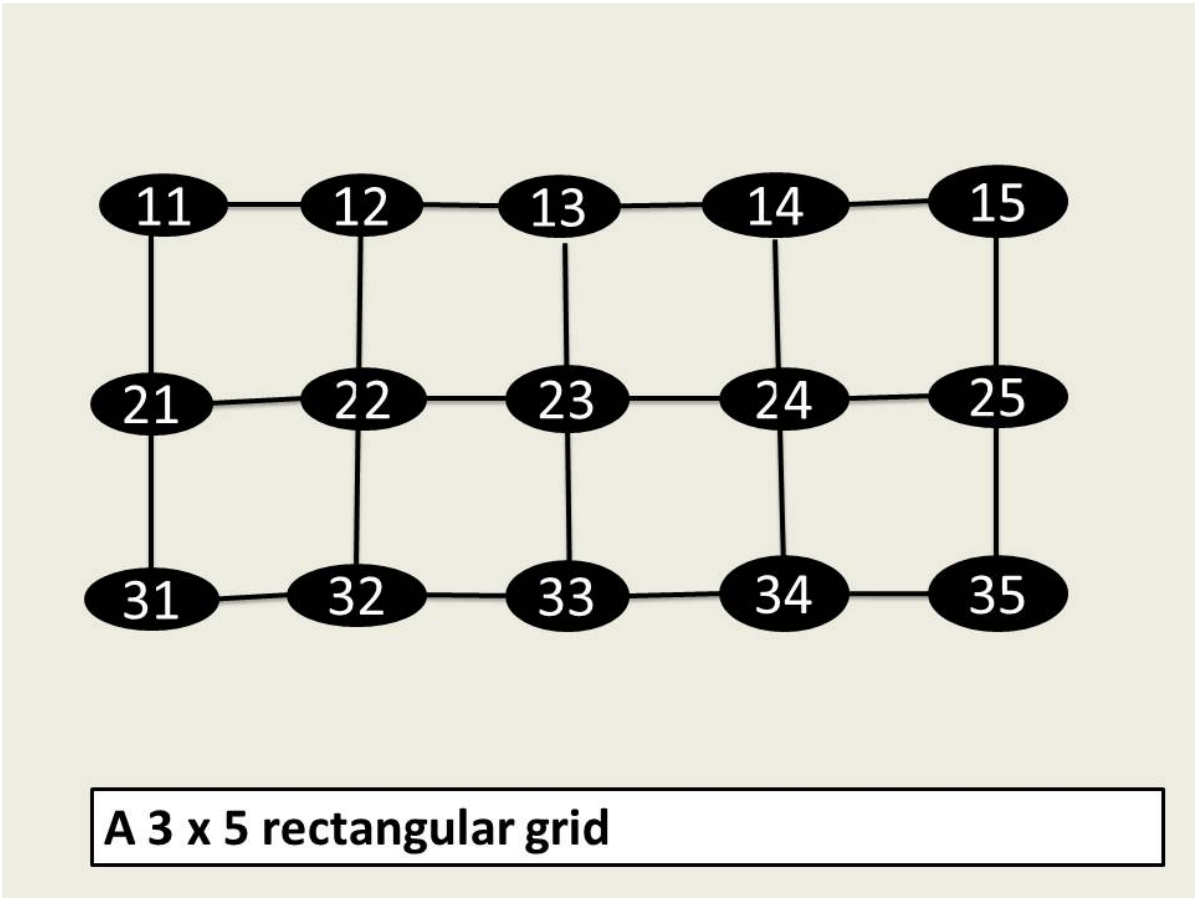

Think of a banana plantation where the disease agent can move only by a distance of at most 1.

Special case 2: Rectangular grids with diagonal edges

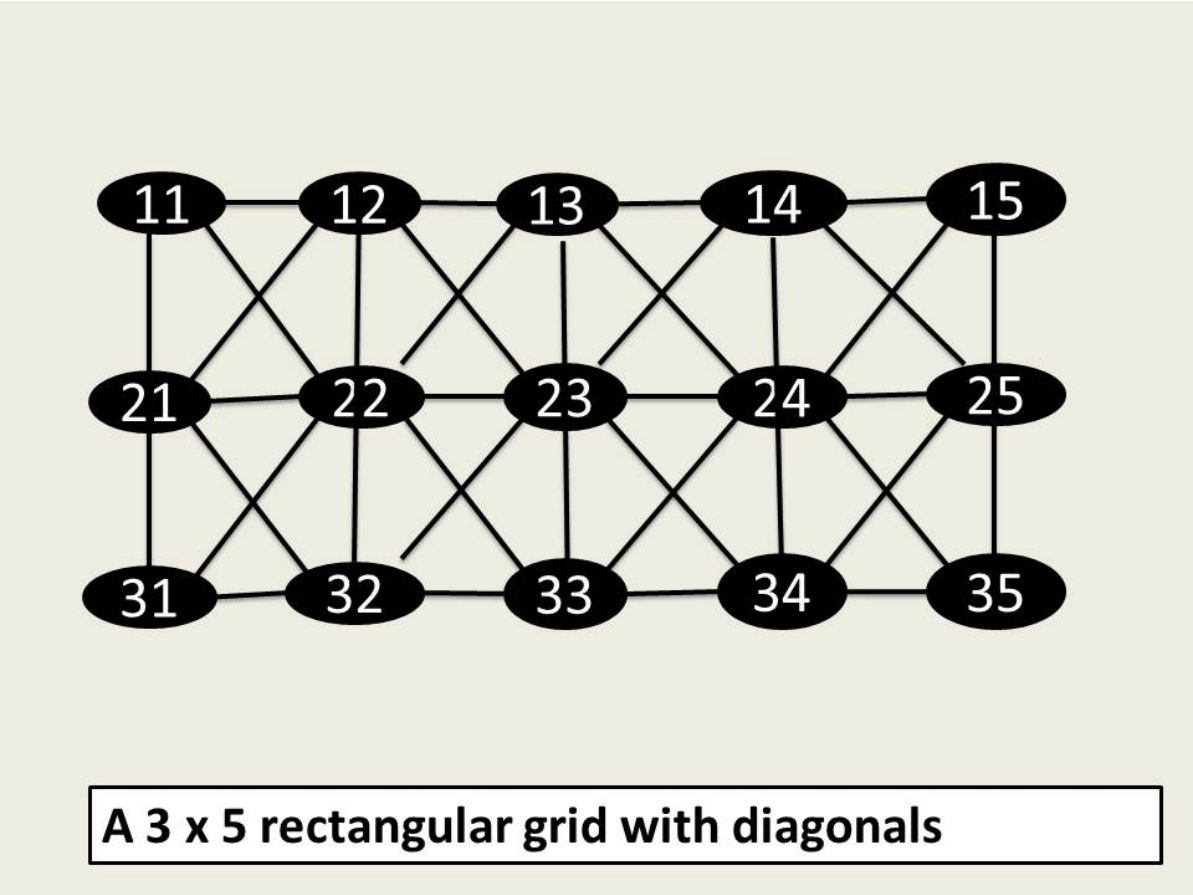

Think of a banana plantation where the disease agent can move only by a distance of at most 1.5.

Exhibit C: Trees

A tree is a connected graph with exactly one path between each pair of nodes. The nodes with degree 1 are called leaves.

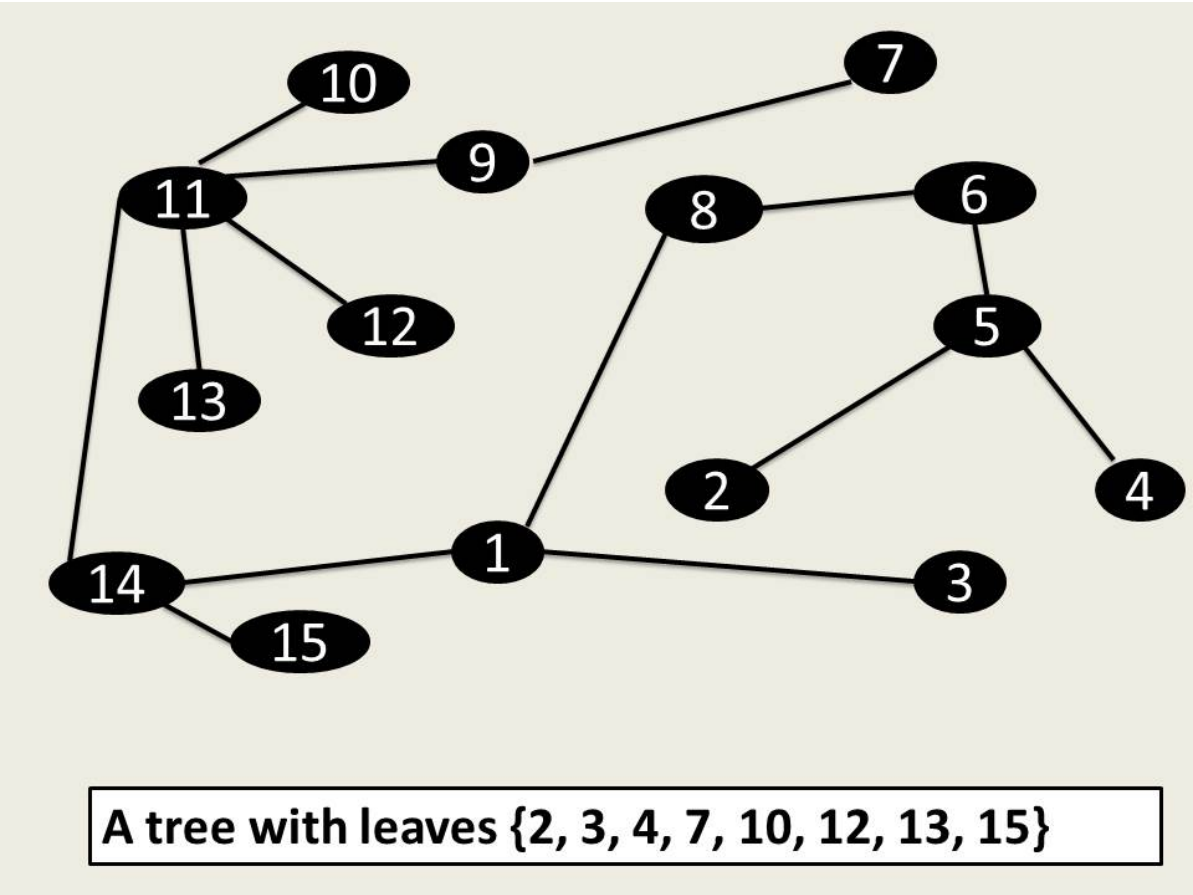

Think of a river system.

Exhibit C: Protecting rivers

Think of a river system. An invasive species, for example, freshwater mussels, may spread along the river. This seems different from disease transmission, but think again:

- The mussels would be the disease agents here.
- The hosts would be branching points of the river system.
- An *SI*-type network model seems appropriate.
- The spread of mussels along river segments could be blocked by physical barriers. This is mathematically equivalent to behavior modification (think "unfriending").
- Alternatively, branching points could be protected by introducing predators in large enclosures. This is mathematically equivalent to immunization.
- Either control measure is expensive.

Given a limited budget, where should we place the predators or barriers?

We will explore various strategies with the help of random trees.