2 EPIDEMICS ON NETWORKS

2.1 EPIDEMICS

Epidemiology is a branch of science that studies the determinants, distribution and dynamics of a propagation process in a population. Given its relation to a whole population, epidemiology is, by design complex and of great interest when shaping public health policies. We will focus on infectious disease epidemiology, *i.e.* on illnesses that can be transmitted from one person to the other. Notable examples include the bubonic plague, smallpox, the Spanish flu, HIV, influenza and, of course, Covid. Even though we will mainly have in mind "illness" propagation, a similar if not identical mathematical framework can be adopted to anythings that propagates through interactions, such as opinions, computer viruses or information.

Let us now describe in deeper detail some fundamental elements of epidemics and their modeling.

2.1.1 EPIDEMIC MODELING

The basic medical observation that we want to model and capture is that when a person is "sick" – for instance it is positive to influenza – and it is in contact with a person who is not – and never was – then it can pass the infection over to the healthy person. The meaning of "contact" depends on the disease we are considering: think to the extreme difference there

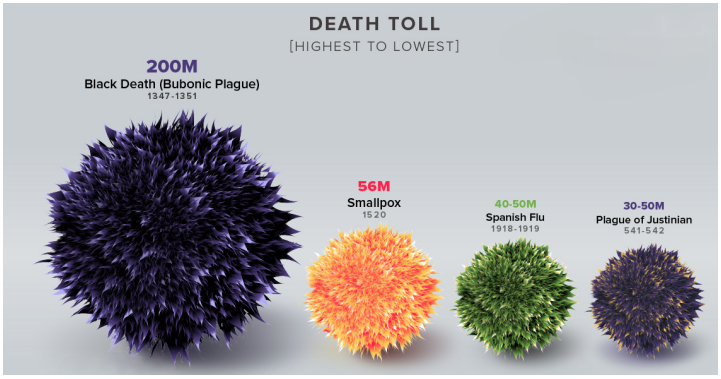

Figure 2.1: **The deadliest pandemics in history**.
Source: visualcapitalist.com/history-of-pandemics-deadliest

is between the HIV and the flu transmissions. Nonetheless, the dynamic processes we have in mind are very similar. We introduce the *Susceptible-Infected-Recovered model* (SIR), a cornerstone of epidemic modeling.

We define three possible states every individual can be in: S, susceptible; I, infected; R recovered. Susceptible people are healthy individuals that can contract the disease. Infectious ones are those who currently carry the disease and can spread it out when they interact with a susceptible individual. The recovered people, instead, are those who used to be infectious and now can no longer be infected. There are two important parameters that take part to this model: β that is the probability *per unit time* to infect a susceptible person and μ, the probability *per unit time* to recover. We can summarize the SIR model with the following equation

The SIR model

$$S + I \xrightarrow{\beta} 2I$$
$$I \xrightarrow{\mu} R$$

The parameters β, μ are disease-dependent and tell us how easily the infection runs across the population. Intuitively, if $\beta \gg \mu$ we are in a situation in which people get infected at a much faster pace than they recover. As a consequence, the epidemic will swiftly spread across the population. On the opposite, if $\mu \ll \beta$ people simply recover very quickly and the epidemic dies out. This concept, that we delineated in intuitive terms, is the so-called epidemic threshold that determines the necessary and sufficient condition for an epidemic spreading to occur and that we now more formally define.

2.1.2 THE EPIDEMIC THRESHOLD

We consider a population in which everybody is susceptible. This condition is a stationary state, because no infection can occur among susceptible people, or, in other words

$$S + S \longrightarrow 2S.$$

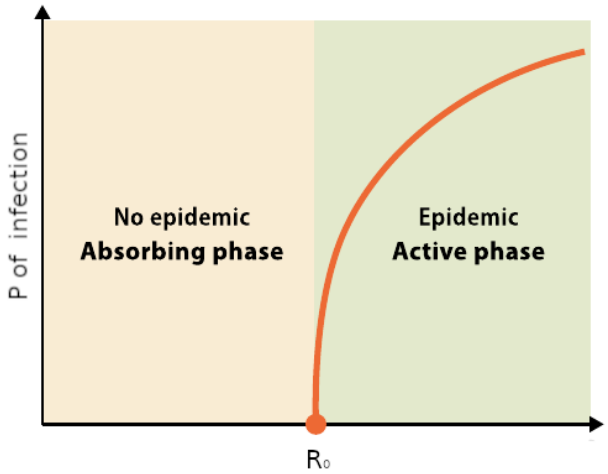

Figure 2.2: **Epidemic threshold phase diagram**. Below the critical value of R_0 we are in an absorbing phase in which there is no epidemic, while for R_0 larger than the critical value, the probability that each node has of being infected is non-null. Source: Satorras, Castellano, Van Mieghem, Vespignani. *Epidemic processes in complex networks.*

We now perturb this state introducing few infectious individuals and ask ourselves whether the system will fall back to an equilibrium having most people being unaffected by the disease, or, if it will spread hitting a large portion of the population. The simple approach to understand this problem lies in the following question

At the beginning of the spreading,
how many people do I infect before recovering?

If the answer is "*more than one*", then we have a cascade in which the propagation grows exponentially fast. If, on the opposite it is "*less than one*" than it dies out because, in most cases, every individual recovers before infecting someone else. We call this the *reproductive number*. As shown in Figure 2.2 it is the control parameter of a phase transition between a disease-free and an epidemic state.

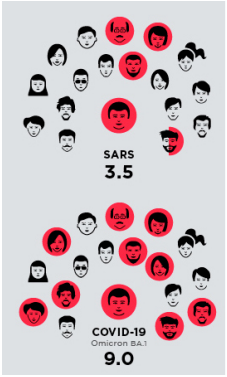

The effect of R_0, some numbers
Source:visualcapitalist.com/history-of-pandemics-deadliest/

Let us stress two important facts. The first one is that R_0 is defined *at the beginning* of the epidemic. Notation-wise, this is why we call it R_0, while R_t is the reproductive number at a given time t. This is important to say because the exponential spreading can occur only on a short time scale and then saturate due to the finite population size. The second fact is that the question refers to a non-better specified "I", implying that the answer is the same for everybody. Of course, if we think of our everyday life, the probability of infection vary a lot across individuals, according to their sociability, the interaction with children, the spaces they occupy *etc.* So, of course, we want an *average* answer and we will study it in the next section.

2.2 EPIDEMIC MODELING ON NETWORKS

As we mentioned, the epidemic spreading propagates through contacts that are well modeled by the edges of a graph. We now want to formally define the epidemic threshold transition for an arbitrary graph. This is of key importance to understand if, given the structure of the network and the disease parameters, the spread will touch a large fraction of individuals or not.

2.2.1 THE STATE EVOLUTION EQUATION

We consider an agent-based model in which every individual $i \in \mathcal{V}$ is associated to a discrete variable $x_i(t) \in \{S, I, R\}$ determining the state i is in. When a susceptible person i is in contact with an infectious one j ($A_{ij} = 1$) for a time dt, then it gets infected with a probability βdt. An infected person recovers with a probability μdt. We can write the probability of being infected at the time-step $t + dt$ as a function of t as follows:

The infected state equation of the SIR on a graph

$$\mathbb{P}(x_i(t+dt) = I) = \mathbb{E}\Bigg[\underbrace{\delta[x_i(t) = I](1-\mu)}_{i \text{ was infected and did not recover}} + \underbrace{\delta[x_i(t) = S]\left(1 - \prod_{j\in\mathcal{V}} \left(1 - \beta dt \cdot \delta[x_j(t) = I]\right)^{A_{ij}}\right)}_{i \text{ was susceptible and got infected}} \Bigg]. \tag{2.1}$$

This equation features two terms: the case in which i was susceptible and got infected and the can in which it was infected and did not recover. Note that the probability of being infected is written as 1 minus the probability of not being infected. We can simplify this equation in the limit for $dt \to 0$, focusing on the term describing the probability of not being infected.

$$\begin{aligned}
&\lim_{dt\to 0} \prod_{j\in\mathcal{V}} \left(1 - \beta dt \cdot \delta[x_j(t) = I]\right)^{A_{ij}} \\
&\overset{(a)}{=} \lim_{dt\to 0} \exp\left\{ \sum_{j\in\mathcal{V}} A_{ij} \log\left(1 - \beta dt \cdot \delta[x_j(t) = I]\right) \right\} \\
&\overset{(b)}{=} \lim_{dt\to 0} \exp\left\{ -\beta dt \sum_{j\in\mathcal{V}} A_{ij} \delta[x_j(t) = I] \right\} \\
&\overset{(c)}{=} \lim_{dt\to 0} 1 - \beta dt \sum_{j\in\mathcal{V}} A_{ij} \delta[x_j(t) = I],
\end{aligned}$$

where in (a) we used the identity $x = e^{\log x}$; in (b) we performed the expansion $\log(1+x) = x + o(x)$ and in (c) the expansion $e^x = 1 + x + o(x)$. Substituting this expression in Equation (2.1) we obtain

$$\begin{aligned}\mathbb{P}(x_i(t+dt) = I) &= \mathbb{E}\left[\delta(x_i(t) = I)\right](1-\mu) + \beta dt \sum_{j\in\mathcal{V}} A_{ij}\mathbb{E}\left[\delta[x_i(t) = S]\delta[x_j(t) = I]\right] \\ &= \mathbb{P}\left(x_i(t) = I\right)(1-\mu) + \beta dt \sum_{j\in\mathcal{V}} A_{ij}\mathbb{P}\left(x_i(t) = S, x_j(t) = I\right).\end{aligned}$$

Taking the first term on the right hand-side to the left and dividing by dt, we obtain the derivative of the probability that reads

$$\partial_t \mathbb{P}(x_i(t) = I) = \beta \sum_{j\in\mathcal{V}} A_{ij}\mathbb{P}\left(x_i(t) = S, x_j(t) = I\right) - \mu\mathbb{P}(x_i(t) = I). \tag{2.2}$$

Following the same passages, we get the evolution equations for all three states. Note that, $\partial_t\mathbb{P}(x_i(t) = S) + \partial_t\mathbb{P}(x_i(t) = I) + \partial_t\mathbb{P}(x_i(t) = R) = 0$, because the probability of being in one of the three states sums up to one.

SIR model on a graph

$$\begin{aligned}\partial_t \mathbb{P}(x_i(t) = S) &= -\beta \sum_{j\in\mathcal{V}} A_{ij}\mathbb{P}\left(x_i(t) = S, x_j(t) = I\right) \\ \partial_t \mathbb{P}(x_i(t) = I) &= \beta \sum_{j\in\mathcal{V}} A_{ij}\mathbb{P}\left(x_i(t) = S, x_j(t) = I\right) - \mu\mathbb{P}(x_i(t) = I) \\ \partial_t \mathbb{P}(x_i(t) = R) &= \mu\mathbb{P}(x_i(t) = I).\end{aligned} \tag{2.3}$$

In order to obtain the reproductive number we must study the stability of this system of equations that, however, is still non-linear and hard to study because it involves the marginal distributions $\mathbb{P}\left(x_i(t) = S, x_j(t) = I\right)$ for which we do not have an explicit expression. To cope with this problem, we adopt the simple naïve mean field approximation.

2.2.2 NAÏVE MEAN FIELD

The *naïve mean field* (NMF) approximation consists in considering all variables as independent *i.e.* in factorizing the marginals as follows

$$\mathbb{P}(x_i(t) = S, x_j(t) = I) = \mathbb{P}(x_i(t) = S)\mathbb{P}(x_j(t) = I). \tag{The NMF approximation}$$

This approximation greatly simplifies the problem. In fact, in (2.3) we have defined the evolution of $3n$ equations concerning the node marginal probabilities, but there are $2|\mathcal{E}|$ equations (where $\mathcal{E}$ is the set of edges) that are unspecified. Factorizing the probabilities with naïve mean field, we simply get rid of these terms. Let us first make some comments about this approximation, its limits and when we expect to be a good method to proceed.

Some notes on the naïve mean field approximation

Given its simplicity NMF is a commonly adopted strategy to first tackle a problem, but is it accurate? In other words, we are asking to what extent we can assume that the event that i is susceptible is independent from the event that its neighbor j is infected. Given the context of the model, the answer seems necessarily to be negative since the contagion is transmitted through the contacts.

The most relevant setting in which NMF should be considered is that of *dense* networks, *i.e.* those in which every node has a large degree. Suppose we have a fully connected network: the fact that the edge A_{ij} exists is simply irrelevant because all edges exist. One can show that indeed, in this setting the NMF approximation becomes asymptotically exact and, in general, the denser the network is the more the NMF approximation is accurate. An example of how to go beyond this approximation is discussed in chapter 3.

With NMF approximation at hand, Equation (2.2) turns into

$$\partial_t \mathbb{P}(x_i(t) = I) = \beta \mathbb{P}(x_i(t) = S) \sum_{j \in \mathcal{V}} A_{ij} \mathbb{P}(x_j(t) = I) - \mu \mathbb{P}(x_i(t) = I).$$

We now linearize this equation around the stationary state $\mathbb{P}(x_i(t) = S) = 1$ to get the reproductive number.

2.2.3 THE REPRODUCTIVE NUMBER WITH NAÏVE MEAN FIELD

As we explained before, we want to see the effect of perturbing the stationary state in which everybody is susceptible by adding a small probability of being infected. For simplicity, we denote $\mathbb{P}(x_i(t) = I) := p_i(t)$ and move to a vector form of the equations. We let $P_i(x_i(t) = S) = 1$ and obtain

The linearization around the diseases-free point under the NMF approximation

$$\partial_t \boldsymbol{p}(t) = (\beta A - \mu I_n) \boldsymbol{p}(t).$$

If we want that $\partial_t p_i(t) < 0$ for all i and all t, we must impose that $\beta \rho(A) - \mu < 0$, where $\rho(A)$ denotes the spectral radius of A. If this condition is satisfied, then we end up in the disease-free region. If on the opposite $\beta \rho(A) - \mu > 0$, the probability of being infected grows at each time step and the virus has a broad diffusion on the network. We thus obtain the following value for the reproductive number

Reproductive number with the NMF approximation

$$R_0 = \frac{\beta \rho(A)}{\mu} \tag{2.4}$$

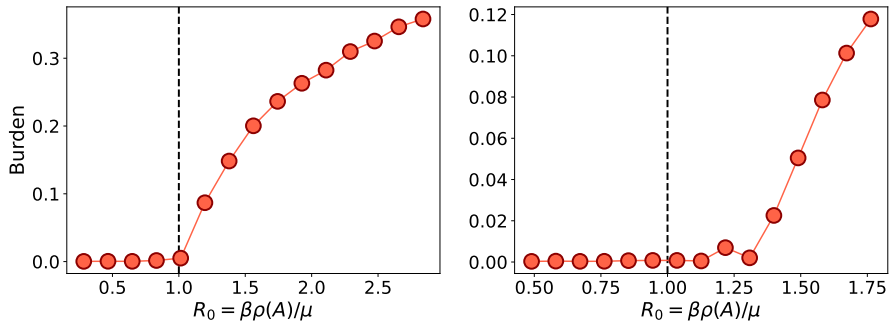

Figure 2.3: **Epidemic threshold: theoretical prediction versus simulated data**. We run a SIR model for different β values on a dense graph (left panel) and on a sparse one (right panel) generated from the random configuration model. We plot the burden (*i.e.* the fraction of non-susceptible individuals) as a function of R_0 as predicted by the NMF approximation.

In Figure 2.3 we compare this prediction with an empirical simulation on a dense (left panel) and on a sparse one (right panel), evidencing the goodness of the approximation only in the former case. Now, as a last step, we give some simple results relating the spectral radius of A with its structure.

2.2.4 GRAPH STRUCTURE AND REPRODUCTIVE NUMBER

Studying the spectral properties of the adjacency matrix for different generative models is a problem of great interest, that however goes beyond the scope of this course. Here we provide two simple examples with intuitive and non-rigorous arguments to characterize the value of $\rho(A)$ and understand the role of density and degree heterogeneity in determining the threshold.

Erdős Renyi random graph

The spectral behavior of the *Erdős-Rényi* (ER) random graph changes dramatically according to whether its expected average degree grows with its size or not. Letting p be the probability of being connected, then we can define two different regimes: the dense one in which $\log(n)/pn = o_n(1)$ and the sparse one in which the opposite is true, *i.e.* $pn/\log(n) = o_n(1)$. In words, if the average degree grows faster than $\log(n)$ we say the network to be *dense*. This is a game-changer because, under this hypothesis, the degree distribution is concentrated, *i.e.* , for all large n, with probability one

$$\max |d_i - np| = o_n(np).$$

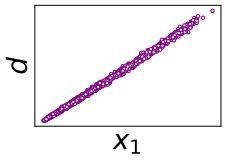

The leading eigenvalue of A, $\boldsymbol{x}_1$ against the degree vector on a random dense graph

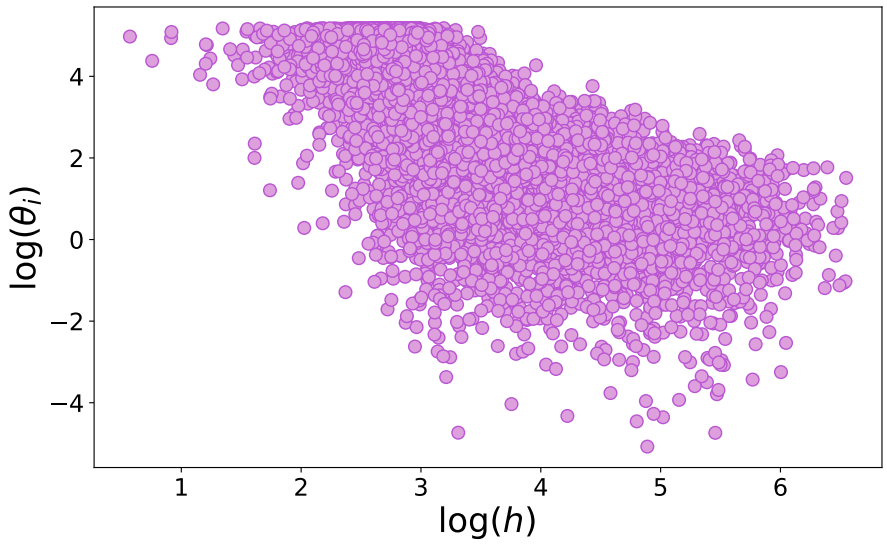

Figure 2.4: **Hitting time as a function of the expected degree**. We consider a random graph generated from the configuration model and denote with θ_i the expected degree of node i. In the plot we show the scatter plot of θ_i against the hitting time h, defined as the number of iterations after which the node got infected in a SIR simulation.

Again, in words, this means that a dense ER graph is quasi regular. In this case, we can heuristically[1] write the following equation

$\rho(A) = \langle d \rangle + o_n(d)$ *on dense ER graphs*

$$(A\mathbf{1}_n)_i = \sum_{j \in \mathcal{V}} A_{ij} = d_i \approx np \approx \langle d \rangle (\mathbf{1}_n)_i.$$

This implies that $\mathbf{1}_n$ is a close approximation of the leading eigenvector[2] and the average degree is an approximation of $\rho(A)$. From this result, we obtain that in denser networks an epidemic spreading runs faster, as one could reasonably expect.

Let us now consider the case in which the graph is generated from a configuration model with an arbitrary degree distribution.

Configuration model

Once again we operate under the assumption of being in a sufficiently dense regime and derive a heuristic expression of the leading eigenvector of A, which we suppose in this case to be $\boldsymbol{d}$, the degree vector.

$\rho(A) = \dfrac{\langle d^2 \rangle}{\langle d \rangle} + o_n(d)$

for dense random graphs with an arbitrary degree distribution

$$(A\boldsymbol{d})_i = \sum_{j \in \mathcal{V}} A_{ij} d_j \approx \sum_{j \in \mathcal{V}} \frac{d_i d_j^2}{2|\mathcal{E}|} = d_i \frac{\langle d^2 \rangle}{\langle d \rangle}.$$

We thus get that $\rho(A) = \langle d^2 \rangle / \langle d \rangle$ implying that a broad degree distribution makes the spreading run even faster on the network. This is because of the

1 For simplicity we derive this result heuristically, but it can be formally proved.
2 Due to Perron-Frobenius theorem.

role played by hubs, *i.e.* nodes with a high degree. Since they have a lot of connections, they are very likely to get infected in the earlier stages of the epidemic and then become super-spreaders. Notably, the value of $\rho(A)$ that we just may diverge for scale-free networks in which the second moment of the degree distribution goes to infinity, making the transition go to zero. Figure 2.4 shows for each node the infection hitting time (*i.e.* the time it takes to get infected) as a function of their degree. The plot evidences a strong negative correlation, confirming the intuition that nodes with a large degree are the first to get infected and then they are responsible of the spreading. Note that in practice the second moment of the degree may only diverge in the asymptotic theoretical limit. All real world networks are finite and so is $\langle d^2 \rangle$. Nonetheless, in real-world settings, it can be overwhelmingly large and drive the epidemic to unfold very fast on the network.

2.2.5 FROM THEORY TO PRACTICE: EPIDEMIC MITIGATION

Let us now discuss some basic facts about epidemic mitigation based on our results. Suppose we have a vaccine and we add a fourth compartment to our model, that of vaccinated people that behaves exactly like the recovered one. From our analytical view-point, we can still deploy the results we obtained, because vaccinated people simply do not take part to the process, since they cannot change their compartment. For this reason it is as if they were not part of the network.

We ask ourselves how vaccination impacts the epidemic spread. To answer this question we add a Boolean variable s_i the equals 0 if i is vaccinated and cannot transmit the disease and is 1 otherwise. The matrix that determines the epidemic threshold is then now

$$W = A \circ (\boldsymbol{s}\boldsymbol{s}^T).$$

Using a non-rigorous argument, we will see how the vaccination determines the epidemic threshold.

Methodological remark

The approach we used to study $\rho(A)$ was non rigorous but leads to the correct result because we assume A to be dense. By vaccinating a large portion of the population, instead W becomes sparse by design and not only the method we adopt but also the result is incorrect. Formally, we can only see what is the effect of vaccination on the spreading, assuming that a small fraction of the population has been vaccinated and W is still dense. This heuristic result, however, gives us some important intuition that we can verify numerically and that can be rigorously proved with other more rigorous methods.

Let $\tilde{d}_i = d_i s_i$. Then

$$\left(W\tilde{\boldsymbol{d}}\right)_i = \sum_{j\in\mathcal{V}} A_{ij} s_i s_j \tilde{d}_j \approx= \sum_{j\in\mathcal{V}} \frac{\tilde{d}_i d_j^2 s_j}{2|\mathcal{E}|} = \tilde{d}_i \frac{\boldsymbol{s}^T \boldsymbol{d}^2}{\mathbf{1}_n^T \boldsymbol{d}},$$

so $\frac{\boldsymbol{s}^T \boldsymbol{d}^2}{\mathbf{1}_n^T \boldsymbol{d}}$ is a close approximation of $\rho(W)$. Now, if s_i is Bernoulli random variable with probability p, we get

$$\rho(W) \approx p \frac{\langle d^2 \rangle}{\langle d \rangle}.$$

So, the vaccination decreases the R_0 and allows one to stay below the epidemic threshold. From a simple observation, however, one sees that this is not the optimal strategy. In fact, in we fix $\boldsymbol{s}^T \mathbf{1}_n$, *i.e.* the number of vaccines, and attempt to minimize $\boldsymbol{s}^T \boldsymbol{d}^2$ one immediately sees that the solution lies in vaccinating the nodes with the highest degree, *i.e.* the *hubs*. This target immunization significantly helps in improving the mitigation effectiveness.

2.3 EXTENSIONS

In the previous sections we only considered the SIR to model an epidemic spreading. While this is one the most relevant models in epidemiology, it must be mentioned that several alternatives exist. The simplest one is the SI in which individuals cannot recover and is equivalent to the SIR for $\mu = 0$. In this case one can see that no epidemic threshold exists and, for how little is the transmission parameter, the epidemic will certainly involve all the population sooner or later. Different is the case of the SIS model in which an infected individual recovers but is once again susceptible. This model accounts for the fact that having experienced an infection does not imply one is immune, in some cases. From a mathematical perspective this slightly changes things: as we commented already, in the SIR, an infected individual cannot have been infected by a susceptible neighbor. On the opposite in the SIS this can happen: a infected individual can have been infected by someone who *now* is susceptible but that recovered. Actually, the difference between these two regimes cannot be understood from the NMF approximation because it is indeed not able to capture this dynamic. If one adopts a more refined approximation strategy, however, it is indeed possible to see that in the two cases the epidemic threshold varies.

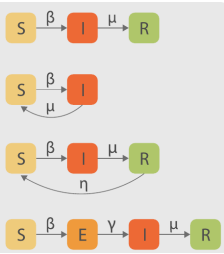

Other models realistically add more compartments to the equations. Some of the most common are the *exposed*, *vaccinated* and *dead* compartments. Exposed people are those who already contracted the virus but that are still not contagious. After some time, they turn infected. The addition of these or other compartments may take into account of more complex medical and behavioral factors. On top of this, the model parameters may add further

depth. We assumed β, μ to be constant and equal for all individuals, while one may assume that they depend, for instance, on age or mask-wearing.

It is worth mentioning that we only talked about simple contagion, *i.e.* the process in which one infected individual passes the disease over to a susceptible one. Thinking however of epidemiology in a broader sense, this is not the only possible alternative. Contagion may occur, for instance, only if one is exposed several times to an infected individual, each one passing a "piece". Only when all "pieces" are passed one becomes infected. Alternatively one can imagine contagion as a process in which it is necessary to have several infectious people interacting *at once* for the disease to be transmitted. We talk in these cases of *complex contagion*.

2.4 REFERENCES

- A.L. Barabasi, *Network Science, Chapter 10*, networksciencebook.com/chapter/10
  This chapter gives a wide and detailed view of epidemics on networks and can be used as a reading to get a bigger picture of the problem that we only treated in a very schematic way. From the analytical viewpoint, the book gives more results on the dynamics and considers also other models but does not use the NMF as it was presented here, but rather presents the degree-based mean field approach that is closely related to NMF.

- Pastor-Satorras, Castellano, Van Mieghem, Vespignani, *Epidemic processes in complex networks*, Reviews of modern physics 87.3 (2015): 925, arxiv.org/pdf/1408.2701.pdf
  This is a long but very important review. The intent of this review is less pedagogical that the Barabasi's book, but it takes a broader look at the problem from a technical perspective and summarizes different results.

- P. Van Mieghem, *Exact Markovian SIR and SIS epidemics on networks and an upper bound for the epidemic threshold*, arxiv.org/pdf/1402.1731.pdf
  This article details some rigorous results based on the NMF approximation with a notation very similar to the one adopted in this chapter.