Lecture Notes in Mathematical Epidemiology

Fred Brauer[1], P. van den Driessche[2], and Jianhong Wu[3]

[1] Department of Mathematics, University of British Columbia
   Vancouver, BC V6T 1Z2, Canada
   `brauer@math.ubc.ca`
[2] Department of Mathematics and Statistics, University of Victoria
   Victoria, B.C. V8W 3P4, Canada
   `pvdd@math.uvic.ca`
[3] Department of Mathematics and Statistics, York University
   Toronto, Ontario M3J 1P3, Canada
   `wujh@mathstat.yorku.ca`

Contents

A light introduction to modelling recurrent epidemics

David J.D. Earn

Department of Mathematics and Statistics
McMaster University, Hamilton ON L8S 4K1, Canada
earn@math.mcmaster.ca

Summary. Epidemics of many infectious diseases occur periodically. Why?

1 Introduction

There are many excellent books that provide broad and deep introductions to the mathematical theory of infectious disease epidemics, ranging from monographs and textbooks [1, 2, 3, 4, 5, 6, 7] to collections of articles from workshops and conferences [8, 9, 10, 11, 12]. My goal in this article is to spark an interest in mathematical epidemiology that might inspire you to dig into the existing literature (starting with the rest of this book[13]) and, perhaps, to engage in some research yourself in this fascinating area of science.

I will discuss some famous epidemics that present challenging theoretical questions, and—without getting bogged down in technical details you can find elsewhere—I will try to convince you that you can fairly easily build and analyze simple models that help us understand the complex patterns evident in these data. Not wishing to give you a false impression of the field, I will also briefly mention some epidemic patterns that do not appear to be explicable in terms of simple models (at least, not in terms of simple models we have thought of!).

Several of the following sections are based in part on an even lighter introduction to the subject of mathematical epidemiology that I wrote for a high school mathematics magazine [14]. Here, I do not limit myself to high school mathematics, but I hope the bulk of the article will be easily accessible to you if you have had an elementary course in ordinary differential equations.

2 Plague

One of the most famous examples of an epidemic of an infectious disease in a human population is the Great Plague of London, which took place in 1665–

1666. We know quite a lot about the progression of the Great Plague because weekly bills of mortality from that time have been retained. A photograph of such a bill is shown in Figure 1. Note that the report indicates that the number of deaths from plague (5533) was more than 37 times the number of births (146) in the week in question, and that wasn't the worst week! (An even worse plague occurred in the 14th century, but no detailed records of that epidemic are available.)

Fig. 1. A bill of mortality for the city of London, England, for the week of 26 September to 3 October 1665. This photograph was taken by Claire Lees at the Guildhall in London, England, with the permission of the librarian.

   Putting together the weekly counts of plague deaths from all the relevant mortality bills, we can obtain the *epidemic curve* for the Great Plague, which I've plotted in the top left panel of Figure 2. The characteristic exponential rise, turnover and decline is precisely the pattern predicted by the classic susceptible-infective-recovered (SIR) model of Kermack and McKendrick [15] that I describe below. While this encourages us to think that mathematical modelling can help us understand epidemics, some detailed features of the epidemic curve are not predicted by the simple SIR model. For example, the model does not explain the jagged features in the plotted curve (and there would be many more small ups and downs if we had a record of daily rather than weekly deaths). However, with some considerable mathematical effort, these "fine details" can be accounted for by replacing the differential equations of Kermack and McKendrick with equations that include stochastic (*i.e.,* random) processes [2]. We can then congratulate ourselves for our modelling success... until we look at more data.


Fig. 2. Epidemic curves for plague in London (left panels) and measles in New York City (right panels). For plague, the curves show the number of deaths reported each week. For measles, the curves show the number of cases reported each month. In the top panels, the small ticks on the time axis occur at monthly intervals.

The bottom left panel of Figure 2 shows weekly mortality from plague in London over a period of 70 years. The Great Plague is the rightmost (and highest) peak in the plot. You can see that on a longer timescale, there was a complex pattern of plague epidemics including extinctions and re-emergences. This cannot be explained by the basic SIR model (even if we reformulate it using stochastic processes). The trouble is likely that we have left out a key biological fact: there is a reservoir of plague in rodents, so it can persist for years, unnoticed by humans, and then re-emerge suddenly and explosively. By including the rodents and aspects of spatial spread in a mathematical model, it has recently been possible to make sense of the pattern of 17th century plague epidemics in London [16]. Nevertheless, some debate continues as to whether all those plagues were really caused by the same pathogenic organism.

3 Measles

A less contentious example is given by epidemics of measles, which are definitely caused by a well-known virus that infects the respiratory tract in humans and is transmitted by airborne particles [3]. Measles gives rise to characteristic red spots that are easily identifiable by physicians who have seen many cases, and parents are very likely to take their children to a doctor when such spots are noticed. Consequently, the majority of measles cases in developed countries end up in the office of a doctor (who, in many countries, is required to report observed measles cases to a central body). The result is that the quality of reported measles case data is unusally good, and it has therefore stimulated a lot of work in mathematical modelling of epidemics.

An epidemic curve for measles in New York City in 1962 is shown in the top right panel of Figure 2. The period shown is 17 months, exactly the same length of time shown for the Great Plague of London in the top left panel. The 1962 measles epidemic in New York took off more slowly and lasted longer than the Great Plague of 1665. Can mathematical models help us understand what might have caused these differences?

4 The SIR model

Most epidemic models are based on dividing the host population (humans in the case of this article) into a small number of compartments, each containing individuals that are identical in terms of their status with respect to the disease in question. In the SIR model, there are three compartments:

- *Susceptible*: individuals who have no immunity to the infectious agent, so might become infected if exposed.
- *Infectious*: individuals who are currently infected and can transmit the infection to susceptible individuals who they contact.

- *Removed*: individuals who are immune to the infection, and consequently do not affect the transmission dynamics in any way when they contact other individuals.

It is traditional to denote the number of individuals in each of these compartments as S, I and R, respectively. The total host population size is $N = S + I + R$.

Having compartmentalized the host population, we now need a set of equations that specify how the sizes of the compartments change over time. Solutions of these equations will yield, in particular, $I(t)$, the size of the infectious compartment at time t. A plot of $I(t)$ should bear a strong resemblance to observed epidemic curves if this is a reasonable model.

The numbers of individuals in each compartment must be integers, of course, but if the host population size N is sufficiently large we can treat S, I and R as continuous variables and express our model for how they change in terms of a system of differential equations,

$$\frac{dS}{dt} = -\beta SI \,, \tag{1a}$$

$$\frac{dI}{dt} = \beta SI - \gamma I \,. \tag{1b}$$

Here, the *transmission rate* (per capita) is β and the *recovery rate* is γ (so the *mean infectious period* is $1/\gamma$). Note that I have not written a differential equation for the number of removed individuals. The appropriate equation is $dR/dt = \gamma I$ (outflow from the I compartment goes into the R compartment) but since R does not appear in equations (1) the equation for dR/dt has no effect on the dynamics of S and I (formalizing the fact that removed individuals cannot affect transmission). This basic SIR model has a long history [15] and is now so standard that you can even find it discussed in some introductory calculus textbooks [17].

If everyone is initially susceptible ($S(0) = N$), then a newly introduced infected individual can be expected to infect other people at the rate βN during the expected infectious period $1/\gamma$. Thus, this first infective individual can be expected to infect $\mathbb{R}_0 = \beta N/\gamma$ individuals. The number $\mathbb{R}_0$ is called the *basic reproduction number* and is unquestionably the most important quantity to consider when analyzing any epidemic model for an infectious disease. In particular, $\mathbb{R}_0$ determines whether an epidemic can occur at all; to see this for the basic SIR model, note in equations (1) that I can never increase unless $\mathbb{R}_0 > 1$. This makes intuitive sense, since if each individual transmits the infection to an average of less than one individual then the number of cases must decrease with time.

So, how do we obtain $I(t)$ from the SIR model?

5 Solving the basic SIR equations

If we take the ratio of equations (1) we obtain

$$\frac{dI}{dS} = -1 + \frac{1}{\mathbb{R}_0 S}\,, \tag{2}$$

which we can integrate immediately to find

$$I = I(0) + S(0) - S + \frac{1}{\mathbb{R}_0} \ln\left[S/S(0)\right]. \tag{3}$$

This is an exact solution, but it gives I as a function of S, not as a function of t. Plots of $I(S)$ for various $\mathbb{R}_0$ show the phase portraits of solutions (Figure 3) but do not give any indication of the time taken to reach any particular points on the curves. While the exact expression for the phase portrait may seem like great progress, it is unfortunately not possible to obtain an exact expression for $I(t)$, even for this extremely simple model.

In their landmark 1927 paper, Kermack and McKendrick [15] found an approximate solution for $I(t)$ in the basic SIR model, but their approximation is valid only at the very beginning of an epidemic (or for all time if $\mathbb{R}_0$ is unrealistically close to unity) so it would not appear to be of much use for understanding measles, which certainly has $\mathbb{R}_0 > 10$.

Computers come to our rescue. Rather than seeking an explicit formula for $I(t)$ we can instead obtain an accurate numerical approximation of the solution. There are many ways to do this [18], but I will briefly mention the simplest approach (Euler's method), which you can implement in a few minutes in any standard programming language, or even a spreadsheet.

Over a sufficiently small time interval Δt, we can make the approximation $dS/dt \simeq \Delta S/\Delta t$, where $\Delta S = S(t+\Delta t) - S(t)$. If we now solve for the number of susceptibles a time Δt in the future, we obtain

$$S(t + \Delta t) = S(t) - \beta S(t)I(t)\Delta t\,. \tag{4a}$$

Similarly, we can approximate the number of infectives at time $t + \Delta t$ as

$$I(t + \Delta t) = I(t) + \beta S(t)I(t)\Delta t - \gamma I(t)\Delta t\,. \tag{4b}$$

Equations (4a) and (4b) together provide a scheme for approximating solutions of the basic SIR model. To implement this scheme on a computer, you need to decide on a suitable small time step Δt. If you want to try this, I suggest taking Δt to be one tenth of a day. I should point out that I am being extremely cavalier in suggesting the above method. Do try this, but be forewarned that you can easily generate garbage using this simple approach if you're not careful. (To avoid potential confusion, include a line in your program that checks that $S(t) \geq 0$ and $I(t) \geq 0$ at all times. Another important check is to repeat your calculations using a much smaller Δt and make sure your results don't change.)

Fig. 3. Phase portraits of solutions of the basic SIR model (1) for a newly invading infectious disease. The curves are labelled by the value of the basic reproduction number $\mathbb{R}_0$ (2, 4, 8 or 16). For each curve, the initial time is at the bottom right corner of the graph ($I(0) \simeq 0$, $S(0) \simeq N$). All solutions end on the S axis ($I \to 0$ as $t \to \infty$). A simple analytical formula for the phase portrait is easily derived (equation 3); from this it is easy to show that $\lim_{t \to \infty} S(t) > 0$ regardless of the value of $\mathbb{R}_0$ (though, as is clear from the phase portraits plotted for $\mathbb{R}_0 = 8$ and 16, nearly everyone is likely to be infected eventually if $\mathbb{R}_0$ is high).

In order for your computer to carry out the calculations specified by equations (4), you need to tell it the parameter values (β and γ, or $\mathbb{R}_0$, N and γ) and initial conditions ($S(0)$ and $I(0)$). For measles, estimates that are independent of the case report data that we're trying to explain indicate that the mean infectious period is $1/\gamma \sim 5$ days and the basic reproduction number is $\mathbb{R}_0 \sim 18$ [3]. The population of New York City in 1960 was $N = 7,781,984$. If we now assume one infectious individual came to New York before the epidemic of 1962 ($I(0) = 1$), and that everyone in the city was susceptible ($S(0) = N$), then we have enough information to let the computer calculate $I(t)$. Doing so yields the epidemic curve shown in the top panel of Figure 4, which does *not* look much like the real data for the 1962 epidemic in New York. So is there something wrong with our model?
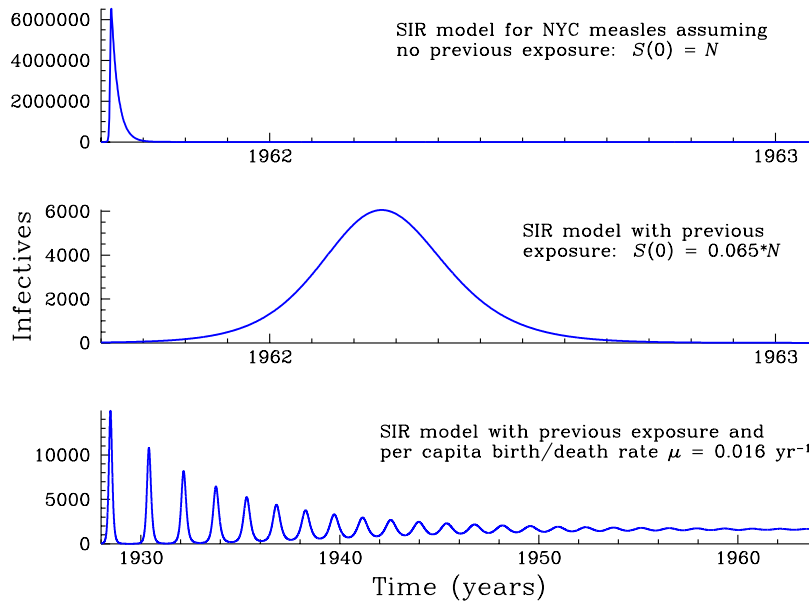
Fig. 4. Epidemic curves for measles in New York City, generated by the basic SIR model. The curves show the number of infectives $I(t)$ at time t. In the top two panels, the small ticks on the time axis occur at monthly intervals. The parameter values and initial conditions are discussed in the main text, except for the initial proportion susceptible used to generate the bottom two panels ($S(0)/N = 0.065$). This initial condition was determined based on the SIR model with vital dynamics, as specified by equations (5). The proportion susceptible at equilibrium is $\hat{S} = 1/\mathbb{R}_0 = 1/18 \simeq 0.056$. At the start of each epidemic cycle that occurs as the system approaches the equilibrium, the proportion susceptible must be higher than $\hat{S}$.

No, but there is something very wrong with our initial conditions. The bottom right panel of Figure 2 shows reported measles cases in New York City for a 36 year period, the end of which includes the 1962 epidemic. Evidently, measles epidemics had been occuring in New York for decades with no sign of extinction of the virus. In late 1961, most of New York's population had already had measles and was already immune, and the epidemic certainly didn't start because one infectious individual came to the city. The assumptions that $I(0) = 1$ and $S(0) = N$ are ridiculous. If, instead, we take $I(0) = 123 * (5/30)$ (the number of reported cases in September 1961 times the infectious period as a proportion of the length of the month) and $S(0) = 0.065 * N$, then we obtain the epidemic curve plotted in the middle panel of Figure 4, which is much more like the observed epidemic curve of Figure 2 (top right panel).

This is progress—we have a model that can explain a single measles epidemic in New York City—but the model cannot explain the recurrent epidemics observed in the bottom right panel of Figure 2. This is not because we still don't have exactly the right parameter values and initial conditions: no parameter values or initial conditions lead to recurrent epidemics in this simple model. So, it would seem, there must be some essential biological mechanism that we have not included in our model. What might that be?

Let's think about why a second epidemic cannot occur in the model we've discussed so far. The characteristic turnover and decline of an epidemic curve occurs because the pathogen is running out of susceptible individuals to infect. To stimulate a second epidemic, there must be a source of susceptible individuals. For measles, that source cannot be previously infected people, because recovered individuals retain lifelong immunity to the virus. So who is it?

Newborns. They typically acquire immunity from their mothers, but this wanes after a few months. A constant stream of births can provide the source we're looking for.

6 SIR with vital dynamics

If we expand the SIR model to include B births per unit time and a natural mortality rate μ (per capita) then our equations become

$$\frac{dS}{dt} = B - \beta SI - \mu S \,, \tag{5a}$$

$$\frac{dI}{dt} = \beta SI - \gamma I - \mu I \,. \tag{5b}$$

The timescale for substantial changes in birth rates (decades) is generally much longer than a measles epidemic (a few months), so we'll assume that the population size is constant (thus $B = \mu N$, so there is really only one new parameter in the above equations and we can take it to be B). As before, we can use Euler's trick to convert the equations above into a scheme that enables a computer to generate approximate solutions. An example is shown in the bottom panel of Figure 4, where I have taken the birth rate to be $B = 126,372$ per year (the number of births in New York City in 1928, the first year for which we have data). The rest of the parameters and initial conditions are as in the middle panel of the figure.

Again we seem to be making progress. We are now getting recurrent epidemics, but the oscillations in the numbers of cases over time damp out, eventually reaching an equilibrium $(\hat{S}, \hat{I})$. Of course, the bottom plot in Figure 4 shows what happens for only one set of possible initial conditions. Perhaps for different initial conditions the oscillations don't damp out? If you try a different set of initial conditions—or lots of different sets of initial conditions—then

I guarantee that you will see the same behaviour. The system will always undergo damped oscillations and converge to $(\hat{S}, \hat{I})$. How can I be so sure, you might ask?

First of all, by setting the derivatives in equations (5) to zero, you can easily calculate (in terms of the model's parameters) expressions for $\hat{S}$ and $\hat{I}$ that are positive (hence meaningful) provided $\mathbb{R}_0 > 1$. Then, by linearizing equations (5) about the equilibrium and computing the eigenvalues of the Jacobian matrix for the system, you will find that the equilibrium is locally stable (the eigenvalues have negative real parts) and approach to the equilibrium is oscillatory (the eigenvalues have non-zero imaginary parts) [3, 7]. But maybe if we are far enough from the equilibrium undamped oscillations are possible?

No, we can prove rigorously that the equilibrium $(\hat{S}, \hat{I})$ is *globally asymptotically stable*, *i.e.,* all initial conditions with $S(0) > 0$ and $I(0) > 0$ yield solutions that converge onto this equilibrium. One way to see this is to scale the variables by population size ($S \to S/N$, $I \to I/N$) and consider the function

$$V(S, I) = S - \hat{S} \log S + I - \hat{I} \log I, \qquad S, I \in (0, 1). \tag{6}$$

With a little work you can show that the time derivative of V along solutions of the model, *i.e.,* $\nabla V \cdot (dS/dt, dI/dt)$ with dS/dt and dI/dt taken from equations (1), is strictly negative for each $S, I \in (0, 1)$. V is therefore a *Lyapunov function* [19] for the basic SIR model. The existence of such a V ensures the global asymptotic stability of the equilibrium $(\hat{S}, \hat{I})$ [19].

Finding a Lyapunov function is generally not straightforward, but functions similar the one given in equation (6) have recently been used to prove global stability of equilibria in many epidemic models [20]. The upshot for our present attempt to understand measles dynamics is that this rigorous argument allows us to rule out the basic SIR model: it cannot explain the real oscillations in measles incidence in New York City from 1928 to 1964, which showed no evidence of damping out. Back to the drawing board?

7 Demographic stochasticity

One thing we have glossed over is the presence of noise. While it is true that for sufficiently large population size N it is reasonable to treat S and I as continuous variables, it is not true that the true discreteness of the number of individuals in each compartment has no observable effect. This was recognized by Bartlett [1] who found that a relatively small amount of noise was sufficient to prevent the oscillations of the basic SIR model from damping out. Whether we recast the SIR model as a stochastic process [2, 21] or simply add a small noise term to the deterministic equations, we can sustain the oscillations that damp out in the bottom panel of Figure 4.

Again, this is progress that has arisen from an important mechanistic insight. But we are left with another puzzle. If you look carefully at the New

York City measles reports in the bottom right panel of Figure 2 you'll see that before about 1945 the epidemics were fairly irregular, whereas after 1945 they followed an almost perfect two-year cycle. Even with oscillations sustained by noise, the SIR model cannot explain why the measles epidemic pattern in New York City changed in this way. Have we missed another important mechanism?

8 Seasonal forcing

So far, we have been assuming implicitly that the transmission rate β (or, equivalently, the basic reproduction number $\mathbb{R}_0$) is simply a constant and, in particular, that it does not change in time. Let's think about that assumption. The transmission rate is really the product of the rate of contact among individuals and the probability that a susceptible individual who is contacted by an infectious individual will become infected. But the contact rate is *not* constant throughout the year. To see that, consider the fact that in the absence of vaccination the average age at which a person is infected with measles is about five years [3], hence most susceptibles are children. Children are in closer contact when school is in session, so the transmission rate must vary seasonally.

A crude approximation of this seasonality is to assume that β varies sinusoidally,

$$\beta(t) = \beta_0(1 + \alpha \cos 2\pi t) \,. \tag{7}$$

Here, β_0 is the mean transmission rate, α is the *amplitude* of seasonal variation and time t is assumed to be measured in years. If, as above, β is assumed to be a periodic function (with period one year) then the SIR model is said to be *seasonally forced*. We can still use Euler's trick to solve the equations approximately, and I encourage you to do that using a computer for various values of the seasonal amplitude α ($0 \leq \alpha \leq 1$).

You might think that seasonal forcing is just a minor tweak of the model. In fact, this forcing has an enormous impact on the epidemic dynamics that the model predicts. If you've ever studied the forced pendulum then you might already have some intuition for this. A pendulum with some friction will exhibit damped oscillations and settle down to an equilibrium. But if you tap the pendulum with a hammer periodically then it will never settle down and it can exhibit quite an exotic range of behaviours including chaotic dynamics [22, 23] (oscillations that look random). Similarly complex dynamics can occur in the seasonally forced SIR model.

Importantly, with seasonal forcing the SIR model displays *un*damped oscillations (and it does this, incidentally, even in the absense of stochasticity). More than that, for different parameter values, the seasonally forced SIR model can produce all the different types of oscillatory measles patterns I have ever seen in real data. So are we done?

14   David J.D. Earn

No. As noted in the previous section, the observed measles epidemics in New York City show very clearly that the dynamical pattern changed over time (bottom right panel of Figure 2) and other significant qualitative changes have been observed in measles case series in other places [24]. How can we explain *changes over time* in the *pattern* of measles epidemics?

9 Slow changes in susceptible recruitment

Once again, the missing ingredient in the model is a changing parameter value. This time it is the birth rate B, which is not really constant. Birth rates fluctuate seasonally, but to such a small extent that this effect is negligible. What turns out to be more important is the much slower changes that occur in the average birth rate over decades. For example, in New York City the birth rate was much lower during the 1930s (the "Great Depression") than after 1945 (the "baby boom") and this difference accounts for the very different patterns of measles epidemics in New York City during these two time periods [24].

A little more analysis of the SIR model is very useful. Intuitively, reducing the birth rate or increasing the proportion of children vaccinated both affect the rate at which new susceptible individuals are recruited into the population. In fact, it is possible to prove that changes in the birth rate have exactly the same effect on disease dynamics as changes of the same relative magnitude in the transmission rate or the proportion of the population that is vaccinated [24]. This equivalence makes it possible to explain historical case report data for a variety of infectious diseases in many different cities.

Interestingly, it turns out that while most aspects of the dynamics of measles can be explained by employing seasonal forcing without noise, both seasonal forcing and stochasticity are essential to explain the dynamics of other childhood diseases [25].

10 Not the whole story

I should emphasize that while the seasonally forced SIR model is adequate to explain observed incidence patterns for childhood diseases, it is definitely not adequate for many other diseases that display recurrent epidemics. An important example is influenza (Figure 5). Influenza viruses evolve in ways that evade the human immune system within a few years, making it possible for each of us to be infected by influenza many times. Influenza models must take into account the simultaneous presence in the population of many different strains that interact immunologically and compete for human hosts. The epidemic pattern shown in Figure 5 bears some similarities to the measles pattern shown in Figure 2, and the effects of seasonal forcing help explain this and other influenza patterns to some extent [26]. But we are far from having

a simple model that can account for both the annual incidence patterns of influenza in humans and the evolution of the virus [27].


Fig. 5. Monthly mortality attributed to pneumonia and influenza (P&I deaths) in the USA in the 20th century. The inset plot shows the period 1910–1930 on a much larger scale, revealing the magnitude of the three peaks that extend beyond the top of the main panel: 1918–1919, 1919–1920 and 1928–1929. Mortality before 1934 is underestimated. It is traditional to combine pneumonia and influenza because many deaths categorized as having pneumonia as the underlying cause are triggered by an influenza infection.

11 Take home message

One thing that you may have picked up from this article is that successful mathematical modelling of biological systems tends to proceed in steps. We begin with the simplest sensible model and try to discover everything we can about it. If the simplest model cannot explain the phenomenon we're trying to understand then we add more biological detail to the model, and it's best to do this in steps because we are then more likely to be able to determine

which biological features have the greatest impact on the behaviour of the model.

In the particular case of mathematical epidemiology, we are lucky that medical and public health personnel have painstakingly conducted surveillance of infectious diseases for centuries. This has created an enormous wealth of valuable data [28] with which to test hypotheses about disease spread using mathematical models, making this a very exciting subject for research.

Acknowledgements

References

[1] M. S. Bartlett. *Stochastic population models in ecology and epidemiology*, volume 4 of *Methuen's Monographs on Applied Probability and Statistics*. Spottiswoode, Ballantyne & Co. Ltd., London, 1960.

[2] N. T. J. Bailey. *The mathematical theory of infectious diseases and its applications*. Hafner Press, New York, second edition, 1975.

[3] R. M. Anderson and R. M. May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, Oxford, 1991.

[4] D. J. Daley and J. Gani. *Epidemic modelling, an introduction*, volume 15 of *Cambridge: Studies in Mathematical Biology*. Cambridge university press, Cambridge, 1999.

[5] H. Andersson and T. Britton. *Stochastic epidemic models and their statistical analysis*, volume 151 of *Lecture notes in statistics*. Springer-Verlag, New York, 2000.

[6] O. Diekmann and J. A. P. Heesterbeek. *Mathematical epidemiology of infectious diseases: model building, analysis and interpretation*. Wiley Series in Mathematical and Computational Biology. John Wiley & Sons, LTD, New York, 2000.

[7] F. Brauer and C. Castillo-Chavez. *Mathematical models in population biology and epidemiology*, volume 40 of *Texts in Applied Mathematics*. Springer-Verlag, New York, 2001.

[8] D. Mollison, editor. *Epidemic models: their structure and relation to data*. Publications of the Newton Institute. Cambridge University Press, Cambridge, 1995.

[9] V. Isham and G. Medley, editors. *Models for infectious human diseases: their structure and relation to data*. Publications of the Newton Institute. Cambridge University Press, Cambridge, 1996.

[10] B. T. Grenfell and A. P. Dobson, editors. *Ecology of infectious diseases in natural populations*. Publications of the Newton Institute. Cambridge University Press, Cambridge, 1995.

[11] C. Castillo-Chavez, with S. Blower, P. van den Driessche, D. Kirschner, and A-A. Yakubu, editors. *Mathematical approaches for emerging and reemerging infectious diseases: an introduction*, volume 125 of *The IMA Volumes in Mathematics and its Applications*. Springer, New York, 2002.

[12] C. Castillo-Chavez, with S. Blower, P. van den Driessche, D. Kirschner, and A-A. Yakubu, editors. *Mathematical approaches for emerging and reemerging infectious diseases: models, methods and theory*, volume 126 of *The IMA Volumes in Mathematics and it Applications*. Springer, New York, 2002.

[13] F. Brauer, P. van den Driessche and J. Wu, editors. *Lecture Notes in Mathematical Epidemiology*.

[14] D. J. D. Earn. Mathematical modelling of recurrent epidemics. *Pi in the Sky*, 8:14–17, 2004.

[15] W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London Series A*, 115:700–721, 1927.

[16] M. J. Keeling and C. A. Gilligan. Metapopulation dynamics of bubonic plague. *Nature*, 407:903–906, 2000.

[17] D. Hughes-Hallett, A. M. Gleason, P. F. Lock, D. E. Flath, S. P. Gordon, D. O. Lomen, D. Lovelock, W. G. McCallum, B. G. Osgood, D. Quinney, A. Pasquale, K. Rhea, J. Tecosky-Feldman, J. B. Thrash, and T. W. Tucker. *Applied calculus*. John Wiley & Sons, Toronto, second edition, 2002.

[18] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, second edition, 1992.

[19] S. Wiggins. *Introduction to applied nonlinear dynamical systems and chaos*, volume 2 of *Texts in Applied Mathematics*. Springer-Verlag, New York, 2 edition, 2003.

[20] A. Korobeinikov and P. K. Maini. A Lyapunov function and global properties for SIR and SEIR epidemiological models with nonlinear incidence. *Mathematical Biosciences and Engineering*, 1(1):57–60, 2004.

[21] D. T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22:403–434, 1976.

[22] J. Gleick. *Chaos*. Abacus, London, 1987.

[23] S. H. Strogatz. *Nonlinear Dynamics and Chaos*. Addison Wesley, New York, 1994.

[24] D. J. D. Earn, P. Rohani, B. M. Bolker, and B. T. Grenfell. A simple model for complex dynamical transitions in epidemics. *Science*, 287(5453):667–670, 2000.

[25] C. T. Bauch and D. J. D. Earn. Transients and attractors in epidemics. *Proceedings of the Royal Society of London Series B-Biological Sciences*, 270(1524):1573–1578, 2003.

[26] J. Dushoff, J. B. Plotkin, S. A. Levin, and D. J. D. Earn. Dynamical resonance can account for seasonality of influenza epidemics. *Proceedings of the National Academy of Sciences of the United States Of America*, 101(48):16915–16916, 2004.

[27] D. J. D. Earn, J. Dushoff, and S. A. Levin. Ecology and evolution of the flu. *Trends in Ecology and Evolution*, 17(7):334–340, 2002.

[28] IIDDA. The International Infectious Disease Data Archive, http://iidda.mcmaster.ca.