Use R!

Ottar N. Bjørnstad

Epidemics

Models and Data using R

Springer

Use R!

Series Editors

Robert Gentleman Kurt Hornik Giovanni Parmigiani

More information about this series at http://www.springer.com/series/6991

Ottar N. Bjørnstad

Epidemics

Models and Data using R

Springer

Ottar N. Bjørnstad
Center for Infectious Disease Dynamics
Pensylvania State University
University Park, PA, USA

ISSN 2197-5736 ISSN 2197-5744 (electronic)
Use R!
ISBN 978-3-319-97486-6 ISBN 978-3-319-97487-3 (eBook)
https://doi.org/10.1007/978-3-319-97487-3

Library of Congress Control Number: 2018953687

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

For Katriona, Esme, and Michael

Preface

Despite an undergraduate degree in Zoology and an MSc on the behavior of voles, I have long been fascinated by theoretical biology and the relationship between models and data, and the feedback between statistical analysis and conceptual developments in the area of infectious disease dynamics, in particular, and ecological dynamics in general. My perpetual frustration has been to read all the wonderfully clever books and journal articles exuding all sorts of nifty maths and stats, but not quite being able to *do* any of it myself when it came to infectious diseases that I care about. This frustration led me to attempt to make myself some worked examples of all this cleverness. Over the years the stack of how-tos has grown, and the following chapters are an attempt at organizing these so they may be useful for others. I have tried to organize the chapters and sections in a reasonably logical way: Chaps. 1–10 are a mix and match of models, data, and statistics pertaining to local disease dynamics; Chaps. 11–13 pertain to spatial and spatiotemporal dynamics; Chap. 14 highlights similarities between the dynamics of infectious disease and parasitoid-host dynamics; finally, Chaps. 15 and 16 overview additional statistical methodology I have found useful in studies of infectious disease dynamics. Some sections are marked as "advanced" for one of two reasons: (1) either the maths or stats is a bit more involved or (2) the topic in focus is a bit more esoteric. Although not marked as such, most of Chap. 10 is advanced in this respect. While less run-of-the-mill, I have thought it important to include these sections, because they cover topics that may be less easy to find code for online.

I have had invaluable help from students, colleagues, and collaborators in my quest. The preconference workshops of "Ecology and Evolution of Infectious Diseases" that I co-taught between 2005 and 2008 enhanced my motivation to annotate many worked examples; bare bones of several of the following sections were written during frantic 24-h stints prior to these workshops. Much of the other material arose from interactions with students and postdocs at Pennsylvania State University's Center for Infectious Disease Dynamics (CIDD). Parts of the epidemics on networks and the R_0 removal estimator is from Matt Ferrari's PhD research, and the age-structured SIR simulator and the SIRWS model are from Jennie Lavine's PhD

work. Working with distributed-delay models has been a collaboration with Bill Nelson and my students Lindsay Beck-Johnson and Megan Greischar. Angie Luis and I cooked up the R code to do transfer functions as part of her PhD research. Much of the code on the catalytic model is from collaborations with Laura Pomeroy and then CIDD postdoctoral fellows Grainne Long and Jess Metcalf. The in-host TSIR was also a collaboration with Jess. The Gillespie code arose from collaborations with postdoctoral fellow Shouli Li and my honor student Reilly Mummah. Reilly also taught me how to write my first `Shiny` app. Away from Penn State, Aaron King and Ben Bolker have at various times been unbelievably patient in teaching me bits of maths I did not understand. Roger Nisbet painstakingly guided me through my first transfer functions during my postdoctoral fellowship at NCEAS. During the same period, Jordi Bascompte introduced me to coupled map lattice models. Finally, Bryan Grenfell showed me wavelets and introduced me to the field of infectious disease dynamics some 20 years ago.

The data used has been kindly shared by Janis Antonovics, Jeremy Burdon, Rebecca Grais, Sylvije Huygen, Jenn Keslow, Sandy Leibhold, Grainne Long, and Mary Poss. The first draft of the text was completed while I was on sabbatical at the University of Western Australia and University of Oslo/the Norwegian Veterinary Institute during 2017. My work leading up to this text has variously been funded by the National Science Foundation, the National Institutes of Health, the US Department of Agriculture, and the Bill and Melinda Gates Foundation.

University Park, PA, USA Ottar N. Bjørnstad
May 2018

Contents

Chapter 1
Introduction

1.1 Preamble

The use of mathematical models to understand infectious disease dynamics has a very rich history in epidemiology. Kermack and McKendrick (1927) is the seminal paper that introduced the equations for the general Susceptible-Infected-Removed model and showed how a set of restrictive assumptions lead to the standard SIR model of ordinary differential equations. During the 1950s and early 1960s stochastic theories of disease dynamics were developed by Bailey (1957) and Bartlett (1960b). Bartlett (1956, 1960a) further pioneered the use of Monte Carlo simulations of epidemics with the aid of "electronic computers" (as opposed to regular human computers), while Muench (1959) proposed the "catalytic" framework for understanding age-incidence patterns.[1] The decades to follow saw broad expansions of theories as well as a surge in real-life application of mathematics to dynamics and control of infectious disease.

There are several excellent textbooks of mathematical epidemiology including Anderson and May (1991) and Keeling and Rohani (2008). The purpose of the current text is not to replicate these efforts but rather use these frameworks as a starting point to discuss practical implementation and analysis. The discussion will be centered around a somewhat haphazard collection of case studies selected to explore various conceptual, mathematical, and statistical issues. The text is designed to be more of a "practicum in infectious disease dynamics."

The dynamics of infectious diseases shows a wide diversity of pattern. Some have locally persistent chains-of-transmission; others persist spatially in "consumer-resource metapopulations." Some infections are prevalent among the young, some among the old, and some are age-invariant. Temporally, some diseases have little

[1] Though, as reviewed by Dietz and Heesterbeek (2002), the original calculations leading to the catalytic model was proposed by Daniel Bernoulli in the late eighteenth century.

O. N. Bjørnstad, *Epidemics*, Use R!, https://doi.org/10.1007/978-3-319-97487-3_1

variation in prevalence, some have predictable seasonal shifts, and others exhibit violent epidemics that may be regular or irregular in their timing. Models and "models-with-data" have proved invaluable for understanding and predicting this diversity, and thence help improve intervention and control. The following chapters are an attempt at providing some notes for a "field guide" for working with data, models, and "models-and-data" to understand epidemics and infectious disease dynamics in space and time.

1.2 In-Host Persistence

Infectious diseases can be classified according to their persistence *within* the host and attack rates with respect to age. Some infections result in life-long colonization of a host because the immune system does not clear them. Such "in-host persistence" may be because the immune system permits it—as for the many symbionts that are beneficial to the host (viz. commmensals and mutualists)—or because detrimental symbionts (viz. pathogens) are able to evade clearance. Examples of "in-host persistent" pathogens are retroviruses such as HIV, latent viruses such as, herpes viruses, and a number of bacteria such as the causative agents of tuberculosis (*Mycobacterium tuberculosis*) and leprosy (*M. leprae*).

"Acute" infections, in contrast, result in transient colonization of the host—that in humans can last for days or months depending on the pathogen—followed by clearance. The clearance is usually immune-mediated, though some viruses like canine distemper virus may run out of target cells and some pathogens may have a programmed life cycle within the host. Some coccidian pathogens within the genus *Eimeria*, for example, go through an exact number of replication cycles in the host (as merozoites) before all pathogen cells are expelled into the environment (as oocysts). The more common example of transience is due to immune-mediated clearance. Examples are plentiful and include acute viruses like measles and influenza, bacteria such as many that causes respiratory disease like bacterial meningitis (e.g., *Neisseria meningitidis*) or whooping cough (*Bordetella pertussis* and *B. parapertussis*), and protozoans such as those that cause malaria (*Plasmodium* spp.).

Among the acute infections we further distinguish between those that leave sterilizing immunity following clearance versus those that leave no or short-lived immunity. This can happen via a number of mechanisms including variable gene expression, rapid evolution, co-circulating strain clouds, or other immune evasive maneuvers. *N. meningitidis* and its congener *N. gonorrhoeae* (which cause gonorrhea), for example, are thought to leave little effective immune memory because of the bacteria's ability to express a very variable arsenal of surface proteins (e.g., Stern

et al. 1984; Tettelin et al. 2000). Many influenza subtypes, in contrast, render effective immune memory short-lived because of rapid evolution; high mutation rates lead to "antigenic drift" and viral recombination during coinfection leads to antigenic "shifts." *Plasmodium falciparum* is thought to be comprised of a diverse set of strains with nonoverlapping "antigenic repertoires" (as well as variable antigen expression) that allows repeat reinfection (e.g., Gupta et al. 1998). A number of common viral afflictions of children have a somewhat more limited strain diversity that may allow several reinfection cycles, but the immune system is ultimately able to cover their antigenic space; Examples include rotavirus (Pitzer et al. 2011) and the enterovirus-complex that cause hand-foot-and-mouse disease (Takahashi et al. 2016). Finally, many pathogens have various "anti-immune devices." Respiratory syncytial virus, for example, uses molecular decoys against neutralizing antibodies (Bukreyev et al. 2008) and *Bordetella pertussis* employs the pertussis toxin to, at least transiently, inhibit recruitment of immune effector cells to sites of infection (Kirimanjeswara et al. 2005).

Many of the remaining "acute, immunizing pathogens"—the ones that result in a transient infection followed by life-long sterilizing immunity—are the poster children of mathematical epidemiology. Notable examples are among the classic vaccine-preventable viruses like measles, rubella, and smallpox. From a biological point of view, the complete failure of immune escape of these pathogens is somewhat mysterious (Kennedy and Read 2017), but the resulting simple dynamical clockwork is a joy to anyone hoping to apply mathematics to understand the living world.

From an epidemiological point of view, it is important to make the *functional*—as opposed to taxonomical—classification of pathogens because it allows us to understand the differences in age-specific attack rates and contrasting disease dynamics. The acute, immunizing infections mainly circulate among the young and therefore comprise the many "childhood" infections because most or all older hosts are immune. From the point of view of the compartmental "SIR-like" formalism (Fig. 1.1), it is thus natural to divide the host population in S, I, and R compartments and assume a unidirectional flow from susceptible children through immune ("removed") adults. In contrast, the prevalence of "in-host persistent" infections will tend to accumulate with age. With respect to the SIR formalism, it is thus natural to consider a model with a unidirectional flow from the S class to a terminal I class. The acute but imperfectly immunizing infections should lead to relatively age-invariant attack rates, and $S \to I \to S$ or $S \to I \to R \to S$ flows depending on the duration of immune protection.

The SIR-like framework predicts how the broad expectation for age-prevalence curves will be modulated by factors such as age-specific pattern of mixing and differential mortality between infected and noninfected individuals. Statistical epidemiology can thus be used to probe empirical patterns to discover subtleties in the dynamics of disease transmission that is hard to observe directly.

1.3 Patterns of Endemicity

We can classify the dynamics of infectious disease according to broad "patterns of endemicity." First, there is the distinction between locally persistent *vs* locally non-persistence pathogens. Local persistence fails when a local chain-of-transmission breaks. This can happen for two very different reasons (Fig. 1.1): (i) The transmission bottleneck is when a pathogen is insufficiently transmissible to sustain a chain of transmission; (ii) at the opposite end of the spectrum is the susceptible bottleneck for acute pathogens that are so transmissible that they burn through susceptibles much faster than they are replenished. In measles, for example, prevaccination cities in the USA smaller than a *critical community size* (CCS) of 250k–500k people did not produce enough children to sustain a local chain-of-transmission (Bartlett 1960a) (Fig. 1.2). Recurrence of such pathogens typically involves spatial dynamics and persistence at the metapopulation scale through spread among asynchronous local host communities (Keeling et al. 2004) or core-satellite dynamics in which a few large cities above the CCS serve as persistent sources for spatial dissemination to communities below the CCS (Grenfell and Harwood 1997; Grenfell et al. 2001).

The 1988 and 2002 epidemics of a related morbilli virus, the phocine distemper virus, in European harbor seals is another illustrations of locally non-persistent infections due to high transmission relative to susceptible recruitment rates (e.g., Swinton 1998). Following introduction into each local population ("haul-out"), explosive local epidemics terminated after 1–4 months due to susceptible depletion. When such epidemics happens so fast that recruitment of susceptibles (through birth, immigration, or loss of immunity) is negligible during the course of the outbreak we call it a "closed epidemics." The closed epidemic is the focus of the standard Susceptible-Infected-Recovered model which we will study in Chap. 2. At the opposite end of the transmissibility spectrum, pathogens may bottleneck because transmission is too ineffective. In particular, if the basic reproductive ratio (R_0, the

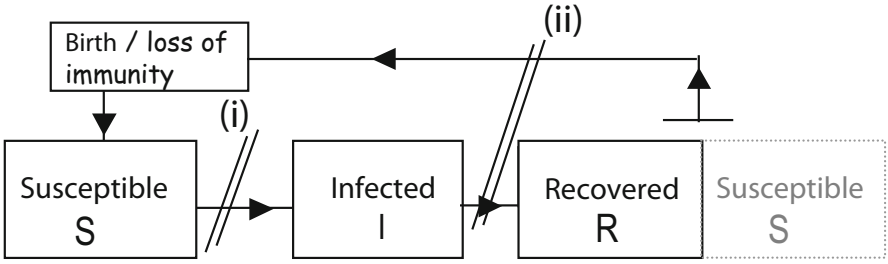

Fig. 1.1 The two bottlenecks for local persistence: (i) the transmission bottleneck for poorly transmitted infections and (ii) the susceptible bottleneck for highly transmissible, acute immunizing (or lethal) pathogens

expected number of secondary cases from a primary case in a completely susceptible population) is smaller than one, we see stuttering ("subcritical") chains of transmission followed by pathogen fade-out. We see this in many zoonoses such as monkey pox and nipah (stage 3 zoonoses in the classification by Lloyd-Smith et al. 2009). Persistent recurrence of these typically involves reservoir host and intermittent zoonotic reintroduction. For example, in their study of Lassa fever in Sierra Leone, Iacono et al. (2015) concluded that about 20% of the human cases were caused by human-to-human transmission (with an average reproductive ratio below one) while the remaining majority was caused by transmission from the multimammate rat (*Mastomys natalensis*) reservoir.

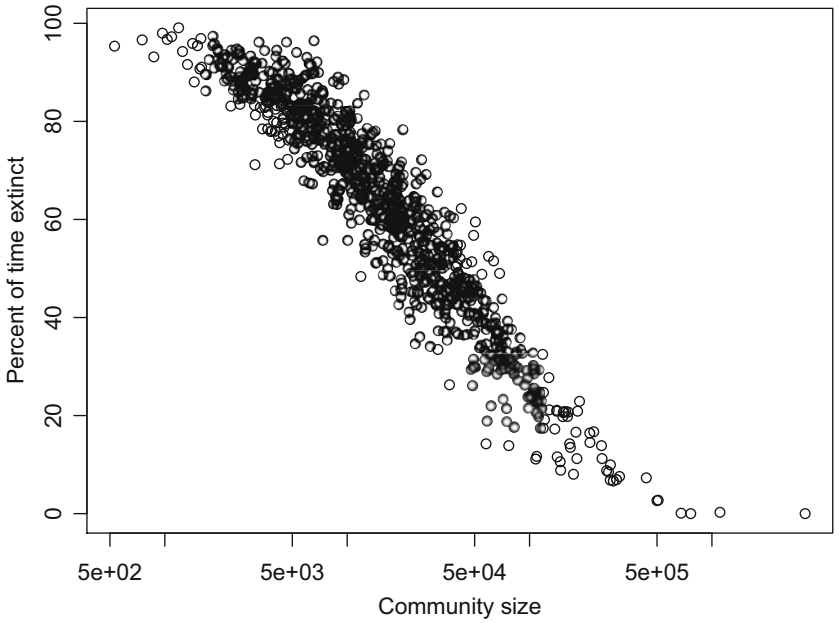

Fig. 1.2 Persistence of measles against population size for 954 cities and villages in prevaccination England and Wales (1944–1964). Communities below 500k exhibited occasional or frequent (depending on size) local extinction of the virus

The locally persistent infections can be classified as: (1) *Stable endemics* that show little variation in incidence through time. Many STDs with SI and SIS-like dynamics like gonorrhea (Fig. 1.3a) and HIV exhibit this pattern. (2) *Seasonal endemics* that show low'ish-level predictable seasonal variation around some mean. Many endemic vector-borne and water-borne infections exhibit this pattern. A classic example is the seasonal two-peaked mortality rate from Cholera in the province

of Dacca, East Bengal (King et al. 2008); The first peak at the beginning of the monsoon season and the second towards the end (Fig. 1.3b). Finally, (3) recurrent epidemics that may be regular or irregular are characterized by violent epidemic fluctuations over time. Many acute, immunizing highly contagious pathogens—measles being the poster-child—follow this pattern (Fig. 1.4).

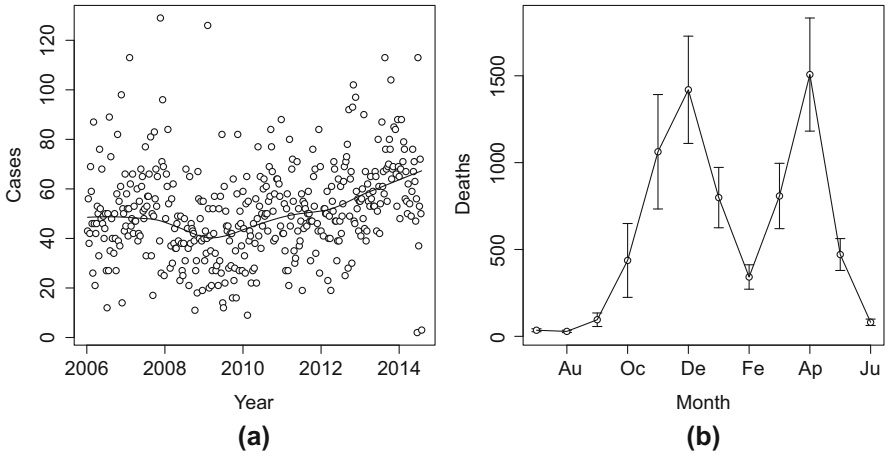

Fig. 1.3 Incidence of (**a**) weekly incidence of gonorrhea in Massachusetts (2006–2015) and (**b**) monthly average (± SE) mortality from cholera in the Dacca district (1891–1940)

1.4 R

To provide a cohesive framework for the practical calculations, all analyses are done in the open-source R-program. The text is written assuming a basic knowledge of this platform. All functions, data, and ShinyApp's discussed in the text are contained in the `epimdr`-package. With the package everything contained herein should be reproducible. The above Figs. 1.2 and 1.4 were for example generated using the following code:

```
#Fig 1.2
data(ccs)
plot(ccs$size, ccs$ext*100, log="x", xlab=
     "Community size", ylab="Percent
     of time extinct")

#Fig 1.3a
plot(magono$time, magono$number, ylab="Cases",
     xlab="Year")
lines(lowess(x=magono$time, y=magono$number, f=.4))
```

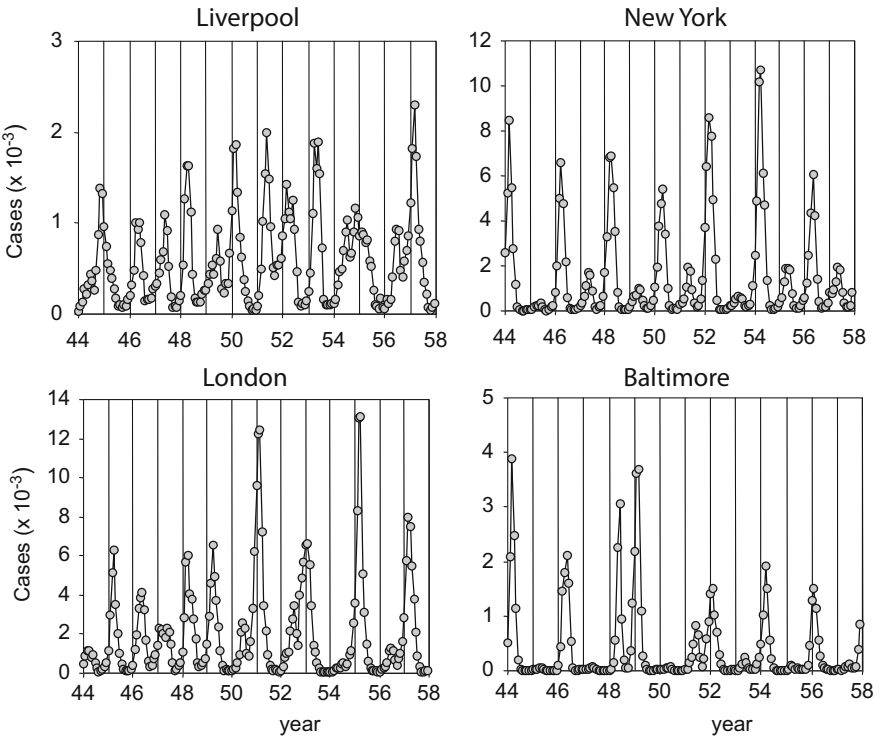

Fig. 1.4 Incidence of measles in various US and UK cities during the pre-vaccination era. The data represent fortnightly incidence (roughly corresponding to the virus' serial interval). The vertical bars mark annual intervals

```
#Fig 1.3b
data(cholera)
ses=sesdv=rep(NA, 12)
ses[c(7:12, 1:6)]=sapply(split(cholera$Dacca,
     cholera$Month), mean, na.rm=TRUE)
sesdv[c(7:12, 1:6)]=sapply(split(cholera$Dacca,
     cholera$Month), sd, na.rm=TRUE)/
     sqrt(length(split(cholera$Dacca, cholera$Month)))
require(plotrix)
plotCI(x=1:12, y=ses, ui=ses+sesdv, li=ses-
     sesdv, xlab="Month", ylab="Deaths")
lines(x=1:12, y=ses)
```

1.5 Other Resources

A 5 min overview of *Patterns of endemicity* can be watched from YouTube: https://www.youtube.com/watch?v=Mf_EZm5amxI. This video is part of the Pennsylvania State University-produced epidemics-MOOC. The entire course is accessible free from https://www.coursera.org/learn/epidemics.

Chapter 2
SIR

2.1 The SIR Model

In 1927, Kermack and McKendrick (1927) published a set of general equations (Breda et al. 2012) to better understand the dynamics of an infectious disease spreading through a susceptible population. Their motivation was

> "One of the most striking features in the study of epidemics is the difficulty of finding a causal factor which appears to be adequate to account for the magnitude of the frequent epidemics of disease which visit almost every population [...] The problem may be summarized as follows: One (or more) infected person is introduced into a community of individuals, more or less susceptible to the disease in question. The disease spreads from the affected to the unaffected by contact infection. Each infected person runs through the course of his sickness, and finally is removed from the number of those who are sick, by recovery or by death. The chances of recovery or death vary from day to day during the course of his illness. The chances that the affected may convey infection to the unaffected are likewise dependent upon the stage of the sickness. As the epidemic spreads, the number of unaffected members of the community becomes reduced [...] In the course of time the epidemic may come to an end. One of the most important problems in epidemiology is to ascertain whether this termination occurs only when no susceptible individuals are left, or whether the interplay of the various factors of infectivity, recovery and mortality, may result in termination, whilst many susceptible individuals are still present in the unaffected population."

Following a general mathematical exposé, they suggested a set of pragmatic assumptions which lead to the standard SIR model of ordinary differential equations

---

This chapter uses the following R-packages: `deSolve`, `rootSolve`, `phaseR`, and `shiny`.
A conceptual understanding of *reproductive ratios* and the *closed epidemic* is useful prior to this discussion. Five minute epidemics-MOOC introductions can be watched from YouTube:
Reproductive number https://www.youtube.com/watch?v=ju26rvzfFg4.
Closed epidemic https://www.youtube.com/watch?v=sSLfrSSmJZM.

O. N. Bjørnstad, *Epidemics*, Use R!, https://doi.org/10.1007/978-3-319-97487-3_2

for the flow of hosts between **S**usceptible, **I**nfectious, and **R**ecovered compartments. In modern notation, their simplest set of equations is (Fig. 2.1):

$$\frac{dS}{dt} = \mu(N - S) - \beta I \frac{S}{N} \tag{2.1}$$

$$\frac{dI}{dt} = \beta I \frac{S}{N} - (\mu + \gamma) I \tag{2.2}$$

$$\frac{dR}{dt} = \gamma I - \mu R. \tag{2.3}$$

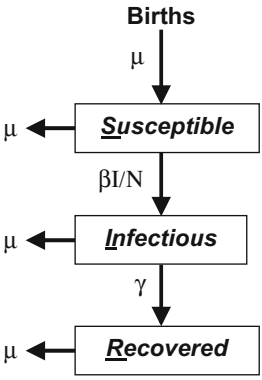

Fig. 2.1 The SIR flow diagram. Flows represent *per capita* flows from the donor compartments

The assumptions of Eqs. (2.1)–(2.3) are:

- The infection circulates in a population of size N, with a per capita "background" death rate, μ, which is balanced by a birth rate μN. From the sum of Eqs. (2.1)–(2.3), $dN/dt = 0$ and $N = S + I + R$ is thus constant.
- The infection causes acute morbidity (not mortality); That is, in this version of the SIR model we assume we can ignore disease-induced mortality. This is reasonable for certain infections like chickenpox, but certainly not for others like rabies, SARS, or ebola.
- Individuals are recruited directly into the susceptible class at birth (so we ignore perinatal maternal immunity).
- Transmission of infection from infectious to susceptible individuals is controlled by a bilinear contact term $\beta I \frac{S}{N}$. This stems from the assumption that the I infectious individuals are independently and randomly mixing with all other individuals, so the fraction S/N of the encounters is with susceptible individuals; β is the contact rate times the probability of transmission given a contact between a susceptible and an infectious individual.

- Chances of recovery or death is assumed not to change during the course of infection.
- Infectiousness is assumed not to change during the course of infection.
- Infected individuals move directly into the the infectious class (as opposed to the SEIR model; see Sect. 3.7) and remains there for an average infectious period of $1/\gamma$ (assuming $\mu << \gamma$).[1]
- The model assumes that recovered individuals are immune from reinfection for life.

The basic reproductive ratio (R_0), defined as the expected number of secondary infections from a single index case in a completely susceptible population, is a very important quantity in epidemiology. Chapter 3 is entirely devoted to this quantity. For this simple SIR model $R_0 = \frac{\beta}{\gamma+\mu}$.

2.2 Numerical Integration of the SIR Model

If there are no (or negligible) births and deaths during the duration of an epidemic ($\mu \simeq 0$), it is commonly referred to as a *closed epidemic*. While it is occasionally possible to derive analytical solutions to systems of ODEs like Eqs. (2.1)–(2.3), we generally have to resort to numerical integration to predict the numbers over time. We use the `deSolve` R-package to numerically integrate the equations. We will numerically integrate a variety of different models. While the models differ, the basic recipe is generally the same: (1) define a R-function for the general system of equations, (2) specify the time points at which we want the integrator to save the state of the system, (3) provide values for the parameters, (4) give initial values for all state variables, and finally (5) invoke the R-function that does the integration. We use the `ode`-function in the `deSolve`-package.

[1] The implicit assumptions that stem from the use of deterministic, ordinary differential equation (ODE) are that the infectious periods (and resident times in all compartments) are exponentially distributed. This is a tractable approximation for exploring overall dynamics, but observed duration of infection periods is often much less variable—the *Eimeria*-gut parasite (a relative of *Plasmodium* that cause malaria) undergoes exactly 8 replication cycles before leaving a host; or much more variable—see superspreader MOOC video: https://www.youtube.com/watch?v=3H1tG4uz9uk. Section 2.7 discusses a practical approach to model dynamics when the exponential assumption is deemed too simplistic.

```
require(deSolve)
```

Step 1: We define the function (often called the gradient-functions) for the equation systems. The `deSolve`-package requires the function to take the following parameters: time,[2] `t`, a vector with the values for the state variables (S, I, R), `y`, and parameter values (β, μ, γ, and N), `parms`:

```
sirmod = function(t, y, parms) {
    # Pull state variables from y vector
    S = y[1]
    I = y[2]
    R = y[3]
    # Pull parameter values from parms vector
    beta = parms["beta"]
    mu = parms["mu"]
    gamma = parms["gamma"]
    N = parms["N"]
    # Define equations
    dS = mu * (N - S) - beta * S * I/N
    dI = beta * S * I/N - (mu + gamma) * I
    dR = gamma * I - mu * R
    res = c(dS, dI, dR)
    # Return list of gradients
    list(res)
}
```

The `ode`-function solves differential equations numerically.

Steps 2–4: Specify the time points at which we want `ode` to record the states of the system (here we use 26 weeks with 10 time-increments per week as specified in the vector `times`), the parameter values (in this case as specified in the vector `parms`), and starting conditions (specified in `start`). In this case we model the *fraction* of individuals in each class, so we set $N = 1$, and consider a disease with an infectious period of 2 weeks ($\gamma = 1/2$), no births or deaths ($\mu = 0$) and a transmission rate of 2 ($\beta = 2$). For our starting conditions we assume that 0.1% of the initial population is infected and the remaining fraction is susceptible.

```
times = seq(0, 26, by = 1/10)
parms = c(mu = 0, N = 1, beta = 2, gamma = 1/2)
start = c(S = 0.999, I = 0.001, R = 0)
```

[2] Though, in the case of the simple SIR model there is no time-dependence in any of the parameters, so this parameter is not called within the gradient function; This will change when we consider seasonality (Chap. 5).

Step 5: Feed `start` values, `times`, the gradient-function and parameter vector to the `ode`-function as suggested by `args(ode)`.[3] For convenience we convert the output to a data frame (`ode` returns a `list`). The `head`-function shows the first 5 rows of `out`, and `round(,3)` rounds the number to three decimals.

```
out=ode(y=start, times=times, func=sirmod, parms=
    parms)
out=as.data.frame(out)
head(round(out, 3))
```

```
  ##   time     S     I R
  ## 1  0.0 0.999 0.001 0
  ## 2  0.1 0.999 0.001 0
  ## 3  0.2 0.999 0.001 0
  ## 4  0.3 0.998 0.002 0
  ## 5  0.4 0.998 0.002 0
  ## 6  0.5 0.998 0.002 0
```

We can plot the result (Fig. 2.2) to see that the model predicts an initial exponential growth of the epidemic that decelerates as susceptibles are depleted, and finally fade-out as susceptible numbers are too low to sustain the chain of transmission.

```
plot(x=out$time, y=out$S, ylab="Fraction", xlab=
    "Time", type="l")
lines(x=out$time, y=out$I, col="red")
lines(x=out$time, y=out$R, col="green")
```

R allows for a lot of customization of graphics—Rseek.org is a useful resource to find solutions to all things R... Fig. 2.2 has some added features such as a right-hand axis for the *effective* reproductive ratio (R_E)—the expected number of new cases per infected individuals in a *not* completely susceptible population—and a legend so that we can confirm that the turnover of the epidemic happens exactly when $R_E = R_0 s = 1$, where s is the fraction of remaining susceptibles. The threshold $R_0 s = 1 \Rightarrow s^* = 1/R_0$ results in the powerful rule of thumb for vaccine induced eradication and herd immunity: If we can—through vaccination —keep the susceptible population below a critical fraction, $p_c = 1 - 1/R_0$, then pathogen spread will dissipate and the pathogen will not be able to reinvade the host population (e.g., Anderson and May 1982; Roberts and Heesterbeek 1993; Ferguson et al. 2003). This rule of thumb appeared to work well for smallpox, the only vaccine-eradicated human disease; Its R_0 was commonly around 5, and most countries saw elimination once vaccine cover exceeded 80% (Anderson and May 1982). The actual code used to produce Fig. 2.2 is:

[3] For further details on usage, do `?function` on the R command-line, i.e., `?ode` in this instance.

```
#Calculate R0
R0=parms["beta"]/(parms["gamma"]+parms["mu"])

#Adjust margins to accommodate a second right axis
par(mar = c(5,5,2,5))
#Plot state variables
plot(x=out$time, y=out$S, ylab="Fraction", xlab="Time",
     type="l")
lines(x=out$time, y=out$I, col="red")
lines(x=out$time, y=out$R, col="green")
#Add vertical line at turnover point
xx=out$time[which.max(out$I)]
lines(c(xx,xx), c(1/R0,max(out$I)), lty=3)

#prepare to superimpose 2nd plot
par(new=TRUE)
#plot effective reproductive ratio  (w/o axes)
plot(x=out$time, y=R0*out$S, type="l", lty=2, lwd=2,
     col="black", axes=FALSE, xlab=NA, ylab=NA,
     ylim=c(-.5, 4.5))
lines(c(xx, 26), c(1,1), lty=3)
#Add right-hand axis for RE
axis(side = 4)
mtext(side = 4, line = 4, expression(R[E]))
#Add legend
legend("right", legend=c("S", "I", "R",
     expression(R[E])), lty=c(1,1,1, 2),
     col=c("black", "red", "green", "black"))
```

2.3 Final Epidemic Size

The closed epidemic model has two equilibria $\{S = 1, I = 0, R = 0\}$ which is unstable when $R_0 > 1$, and the $\{S^*, I^*, R^*\}$-equilibrium which reflects the *final epidemic size*, for which $I^* = 0$ as the epidemic eventually self-extinguish in the absence of susceptible recruitment; S^* is the fraction of susceptibles that escape infection altogether; and R^* is the *final epidemic size*—the fraction of susceptibles that will be infected before the epidemic self-extinguish. For the closed epidemic, there is an exact mathematical solution to the final epidemic size (below). It is nevertheless useful to consider computational ways of finding equilibria in the absence of exact solutions.

The `rootSolve`-package will attempt to find equilibria of systems of differential equations through numerical integration. The function `runsteady` is really just a wrapper function around the `ode`-function that integrates until the system settles on some steady-state (if it exists). It takes the same arguments as `ode`. By varying initial conditions `rootSolve` should find multiple *stable* equilibria if there are more than one stable solution.[4]

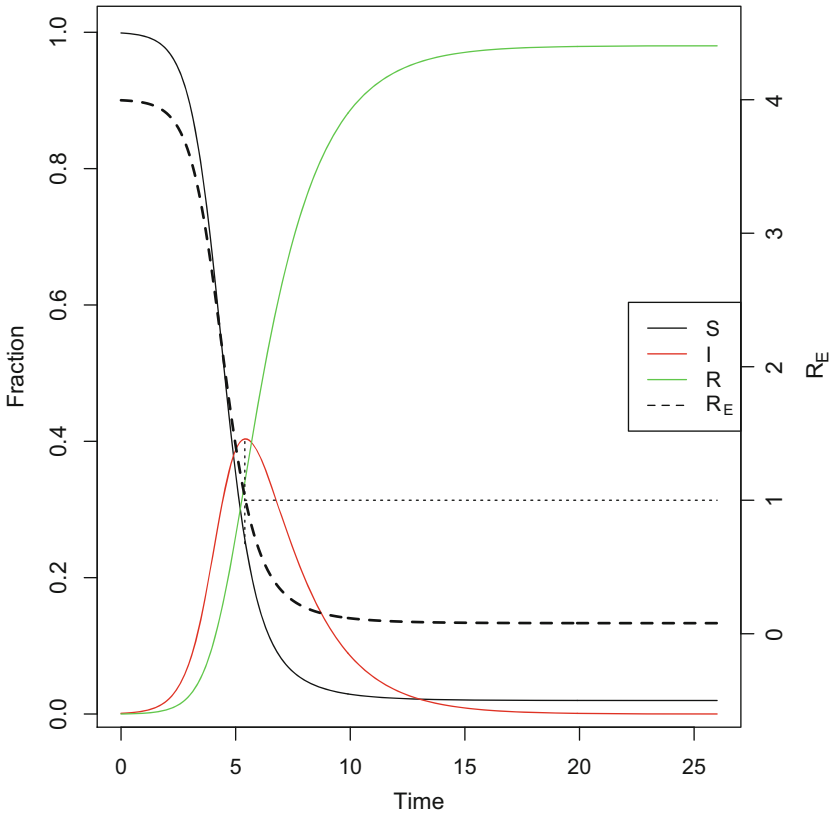

Fig. 2.2 The closed SIR epidemic with left and right axes and effective reproductive ratio, R_E. The epidemic turns over at $R_E = 1$

```
require(rootSolve)
equil=runsteady(y=c(S=1-1E-5, I=1E-5, R=0),
times=c(0,1E5), func=sirmod, parms=parms)
round(equil$y, 3)
```

[4] It will not find unstable equilibria, for these we will need to use other strategies. We will consider finding all equilibria in more depth in Sect. 9.3.

```
##    S    I    R
## 0.02 0.00 0.98
```

So for these parameters, 2% of susceptibles are expected to escape infection altogether and 98%—the final epidemic size—are expected to be infected during the course of the epidemic.

Let us explore numerically how the final epidemic size depends on R_0. Recall that for the specific SIR variant we are working with $R_0 = \beta/(\gamma+\mu)$, and since we are studying the closed epidemic $\mu = 0$. In the above example we assume an infectious period of 2 weeks (i.e., $\gamma = 1/2$), so we may vary β so R_0 goes from 0.1 to 5. For moderate to large R_0 this fraction has been shown to be approximately $1-\exp(-R_0)$ (e.g., Anderson and May 1982). We can check how well this approximation holds (Fig. 2.3).[5]

```
#Candidate values for R0 and beta
R0 = seq(0.1, 5, length=50)
betas= R0 * 1/2
#Vector of NAs to be filled with numbers
f = rep(NA, 50)
#Loop over i from 1, 2, ..., 50
for(i in seq(from=1, to=50, by=1)){
     equil=runsteady(y=c(S=1-1E-5, I=1E-5,
       R=0), times=c(0,1E5), func=sirmod,
       parms=c(mu=0, N=1, beta=betas[i], gamma=1/2))
     f[i]=equil$y["R"]
}
plot(R0, f, type="l", xlab=expression(R[0]))
curve(1-exp(-x), from=1, to=5, add=TRUE, col="red")
```

We see that the approximation is good for $R_0 > 2.5$ but overestimates the final epidemic size for smaller R_0 (and is terrible for $R_0 < 1$).

For the closed epidemic SIR model, there is an exact mathematical solution to the fraction of susceptibles that escapes infection $(1-f)$ given by the implicit equation $f = \exp(-R_0(1-f))$ or equivalently $\exp(-R_0(1-f)) - f = 0$ (Swinton 1998). So we can also find the final size by applying the `uniroot`-function to the equation. The `uniroot`-function finds numerical solutions to equations with one unknown variable (which has to be named `x`).

```
#Define function
fn=function(x, R0){
 exp(-(R0*(1-x))) - x
```

[5] We use a `for`-loop here to calculate the final epidemic size for a range of values of R_0; A loop works by repeating calculations (in this case 50 times), after each repeat the value of the looping variable (in this case `i`) is changed to the next value in the looping vector. So in this example `i` will be 1 first, then 2, then ... until the loop ends after `i`= 50.

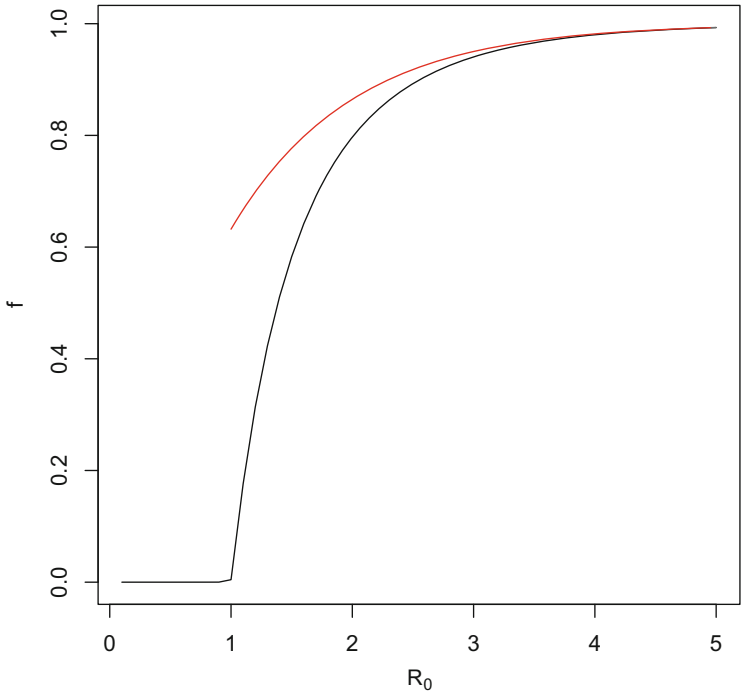

Fig. 2.3 The final epidemic size as a function of R_0. The black line is the solution based on numerically integrating the closed epidemic, and the red line is the approximation $f \simeq 1 - \exp(-R_0)$

```
}
1-uniroot(fn, lower = 0, upper = 1-1E-9,
      tol = 1e-9, R0=2)$root

  ## [1] 0.7968121

#check accuracy of approximation:
exp(-2)-uniroot(fn, lower = 0, upper = 1-1E-9,
      tol = 1e-9, R0=2)$root

  ## [1] -0.06785259
```

So for $R_0 = 2$ the final epidemic size is 79.6% and the approximation is off by around 6.7% points.