

Chapter 23

Modelling Emerging Viral Epidemics for Public Health Protection

Steve Leach and Ian Hall

Abstract

Mathematical models when applied to infectious disease data can provide extremely useful insights into the possible future impacts of potential emerging epidemics and how they might be best controlled or mitigated. Modelling, which is like any other hypothesis-driven approach, aims to develop a better understanding of biological phenomena. However, diseases processes generally, and particularly those related to transmission, will in many cases be imperfectly understood or too complex to systematically describe, so models will necessarily be simplifications of the overall system. It is essential, therefore, that models are designed carefully and used appropriately. Key to this is identifying what specific questions a model might be expected to answer and what data is available to inform the model. A particular type of model might be fine for one particular situation but highly inappropriate for another. It is also important to appreciate and communicate what simplifications and assumptions have had to be made and how this might affect the robustness of the modelling results. It is also particularly important to understand that models frequently make what can be hidden assumptions about underlying processes because of the way they have been constructed and these assumptions also need to be carefully considered and made explicit, particularly for non-expert audiences. This chapter, therefore, provides a brief introduction to some of these aspects of epidemic modelling for those that might be less familiar with them.

Key words: Epidemic modelling, Emerging infectious diseases, Epidemiology, Transmission, Pandemics, Deliberate release, Policy and planning, Preparedness, Prevention and control, Public health interventions, Isolation/quarantine, Vaccination, Real-time modelling

1. Background to Epidemic Modelling

The application of mathematics within the sciences has a long and interesting history and has been used most extensively in the modelling of physical and chemical systems to better understand their underlying processes and test hypotheses. Biological systems have probably seen comparatively less use of mathematical

modelling, due mainly to the systems under consideration being less subject to well-defined “governing laws” (such as often apply at some level in physics and chemistry) and the difficulties of observing biological systems in sufficiently well-controlled and comprehensively understood contexts. There is, nevertheless, a considerable literature associated with the use of mathematical modelling in the context of public health and infectious diseases, with early investigations exemplified by Bernoulli (1) (in relation to smallpox inoculation), Ross (2) (in relation to malaria transmission), and Kermack and McKendrick (3) (in relation to epidemics more generally). In historically more familiar areas, mathematical investigations have often been related to childhood vaccination programmes and combating sexually transmitted infections (4, 5). For the sake of brevity, the extensive background literature related to this subject will not be covered in depth here; but there is a selection of excellent reviews provided in the bibliography (4–10). Similarly, other modelling applications that have also received increasing attention recently include those that aim to provide ongoing advice in the face of such outbreaks; these often being referred to as “real-time” epidemic modelling, “now-casting,” or “forecasting,” depending on context (11–14). Other more recent applications of note, however, have related to providing contingency planning advice ahead of time for potentially high impact outbreaks of emerging infectious diseases (for which we have little current or certain knowledge) arising from acts of bioterrorism or from more natural pandemics, e.g. smallpox, pandemic influenza, and SARS (15–24). Consequently, it is on some of these latter areas of application to contingency planning for emerging infections that this chapter will mainly focus. This chapter is also not intended to be a comprehensive review of the literature nor a technical treatise on how to set-up or use mathematical models. Rather, it is intended to provide a short introduction and pragmatic overview to assist the familiarisation of non-specialist audiences engaged in public health protection. It will thus avoid too much technical description.

2. The Benefits of Using Models

First, what is a mathematical model? Models are generally considered as a representation of a system based on our knowledge of what we understand to be its constituent processes and the relationships that operate between them. Ideally, such representations should be formulated in such a way that it is possible to generalise about the system’s behaviour in every (or most) instance where it is observed. A mathematical model “simply” achieves this representation through one or more mathematical equations.

Necessarily, models, including mathematical ones, will be simplifications of the system under consideration. However, this situation is no different to that which pertains with the types of conceptual models that more generally underlie the development, testing, and potential predictive power of scientific hypotheses. Mathematical models in this context are neither intrinsically no better nor no worse than any other used in science, as long as they are equally supported by data; though they do have some very particular and useful properties. Mathematical models do, nevertheless, often have a vocabulary and notation that is not necessarily easily accessible; and may also come with an air of precision that some might find difficult to accept. It is key then that the (simplifying) assumptions that mathematical models make are made explicit and understandable to more general audiences. This is often best achieved in a collaborative and multidisciplinary environment, including not only the mathematics, but also, for example, the epidemiology, disease and public health expertise. It is also often beneficial to engage with all the constituencies that might come to be dependent on the modelling, including those involved in risk management and risk communication. Indeed, this is well recognised in the broader sphere of risk assessment, and probably most well developed and articulated in the infectious diseases field for food-borne infection risk assessments, as per relevant guidelines of the World Health Organisation (WHO), Food and Agriculture Organisation of the United Nations (FAO), the World Organisation for Animal Health (OIE), and the Codex Alimentarius (25, 26).

3. What Makes a Good Mathematical Model?

Mathematical models do have huge potential for gaining a better understanding of the complex biological and epidemiological systems that underlie emerging infectious disease threats and thus enabling better prospects for their control. Before examining such models in detail, however, it is worth considering what makes a good mathematical model, both in relation to its use and its limitations. Keeling and Rohani (5) identify accuracy, transparency, and flexibility as important aspects that require careful consideration, which can frequently be at odds with each other. Accuracy suggests an ability to quantitatively reproduce observed epidemic data in a consistent fashion (suggesting that a model has predictive power), whereas transparency suggests that a model is well understood in terms of how its various constituent parts interact to generate the resulting epidemic dynamics. The tension here is that more complex (and often less transparent) models which seek to capture the increasingly detailed biology of the

underlying system(s) are generally more likely to better capture their quantitative dynamics (and thus seem more accurate). Put simply, the more parameters/factors that are included in a model, the more likely it is that there will be sufficient parameters that can be (independently) tuned that it is ultimately bound to more accurately reproduce at least one observed epidemic, whether for entirely the right reasons or not. They therefore become increasingly more difficult to understand in terms of how the interactions of the various constituent parts impact on the dynamics and the degree to which the model is able to be generalised. Complex models can also pose other challenges in terms of our ability to parameterise them satisfactorily and in the computational power required to operate them. Contemporary computational capabilities, however, now often make the latter less limiting than does our dearth of knowledge and data concerning diseases and their natural history and transmission, and thus their ability to be satisfactorily parameterised.

Finally, flexibility refers to the relative ease with which models can be adapted to new public health problems. Simpler models can be inherently more flexible than more complex ones, since the latter are more likely to have features that are specific to a particular disease and set of circumstances, several of which may not be relevant to the new problem to be addressed. Possibly, more importantly, they may also lack key features that become crucial. One possibility would be to have one enormously complex model that contains all possibilities in terms of disease natural history (i.e. all the factors that impact on the progression of the disease both within individuals and populations, such as incubation period, viral titres in body fluids, infectivity, etc.), public health interventions, and so forth, with the potential to turn features on and off as required. The problems with this approach would be: first, the unnecessary computational overhead that such a model would always be carrying for any problem that required investigation; second, it is unlikely that every eventuality will have been foreseen at the outset, ultimately requiring the model to be rewritten anyway; third, the potential for over-reliance on “black-box”-type approaches where the underlying model is implicitly trusted but not transparent, and the tendency to be tempted to use a model that is far more complex than is warranted given what little might be known about a particular emerging infectious disease problem. A more reasonable approach is to have a toolbox of models of differing complexity that can be used for the question(s) at hand and a cadre of modelling experts to operate them. Thus, in the spirit of Ockham’s razor (27), a good model is the simplest one (or set) that is suitable to the purposes to which it is to be put, having the right balance of accuracy, transparency, and flexibility, and one that has been constructed with due reference to what is known and preferably measurable.

4. Value of Models

Mathematical representations (models) of systems can take a variety of forms and degrees of complexity, from simply descriptive or explanatory ones concerning a single or a few variables to ones descriptive of more complex multivariable systems. The latter, when suitably validated, have potentially important predictive capabilities and frequently provide insights that are not directly observable or immediately intuitive (23). They also make it possible to undertake experiments (in silico) such as the optimisation of public health interventions that would otherwise be impossible for reasons of, for example, expense, practicality (e.g. no contemporary outbreak/data/evidence) or ethics (e.g. one cannot generally deliberately infect individuals or communities).

5. Introduction to Simple Models and Their Uses in Epidemiology

At one end of the scale of complexity, popular statistical approaches to data analysis and interpretation could be considered as modelling. They provide a more systematic and useful description of a series of observations than simply scanning of a long list of numbers that demonstrates that a parameter in question (e.g. height of mature individuals) varies. If the statistical approach is a parametric one, then there will also be some underlying assumption about the form the data takes. For example, the fitting of a normal, lognormal, or Gamma probability distribution (i.e. ones that have defined parameters such as mean, variance, etc.; hence parametric), might imply that there are underlying processes that produce such a distributional form. With care, for many biological processes, it is often a reasonable assumption that there is an expected distribution to which the data could conceivably conform. In the case of the normal distribution the assumption is that whatever is being measured is subject to multi-factorial (but additive) influences; for example, many genetic loci and environmental influences acting in concert, each with their own “distribution” in relation to their action on what is being measured. Together these interactions produce a joint distribution that is closely approximated by a normal distribution by virtue of the central limit theorem. However, phenomena in nature often produce distributions that are right skewed and with zero probability of values of zero or less. Such systems are generally more closely represented by log normal distributions, where the underlying assumption is that the many factors contribute multiplicatively rather than additively. Further detail on probability distributions

and their applications are outside of the scope of this chapter and can be found elsewhere (28, 29). However, the more complex models that will be discussed later are often built around multiple components that are parameterised as probability distributions. For example, the curves in Fig. 5 and the solid line in Fig. 6 illustrate the types of skewed distribution that are often used to better describe the observed lengths of, for example, the incubation or symptomatic periods of infectious diseases, particularly when these have been sampled from a large enough number of infected individuals in a population to arrive at a sufficiently robust result. Thus, the length of any one individual's incubation period will differ from those of others with some durations (those closer to the average) more frequently observed than others. Clearly, incubation periods of less than zero would be meaningless and use of the lognormal avoids these. Further, the lognormal, and other potentially skewed distributions such as the Gamma, capture well the distribution of the duration of the positive values, probably because incubation periods depend on a wealth of multi-factorial influences (not always well understood) acting multiplicatively, such as the differing pre-existing genetic backgrounds and immunological susceptibilities of individuals to particular viruses, their age and underlying fitness in terms of any pre-existing co-morbidities, and their nutritional and socio-economic conditions; and any variations in the pathogenicity of the particular virus or microorganism, the dose of virus received by an individual and the route of infection. Incubation periods can vary enormously between different infectious agents both in terms of average incubation period and the degree to which this varies between individuals; for example, the mean (and variances) for influenza, SARS coronavirus (SARS CoV), and smallpox have been estimated to be about 1.3 days (0.5 days²), 4.6 days (15.9 days²), and 11.6 days (3.34 days²), respectively. Appropriately parameterising such aspects can be hugely important to modelling the dynamics (e.g. the rate and extent of spread) of particular infections and the potential impact that different public health control options might have. The extent to which they can be accurately reflected in more complex models can be extremely important and will be discussed later.

If the observed data is a series of measurements that varies systematically over time (i.e. a time series) or varies systematically with some other factor that has also been measured (e.g. infectiousness with respect to virus levels in body fluids), then we can also use regression analysis to statistically interpret such relationships. With regression, the model that is quite often imposed is a linear one ($Y = aX + b$), where Y is the observed data and X the (independent) variable; with a and b the slope and intercept parameters, respectively. A simple example concerns crude analysis of the smallpox outbreaks that resulted from importations into

Europe from abroad between 1959 and 1973, when the disease had otherwise been eradicated from that continent. With these importations, and the outbreaks that quite frequently resulted, the greater the delay in their initial notification to the health services, the greater the likely number of cases at the end of an outbreak (30). Whilst this is entirely logical and to be expected, the clear mathematical relationship that was observed between these two variables was striking. Although for illustration here a linear regression model (Fig. 1, solid line) has been fitted, it is worth pointing out that the data on the y -axis has actually been expressed as \log_{10} of the observations. This is because the relationship between the untransformed variables is in fact an exponential one, with the log transform simply linearising this for ease of exploratory data analysis. The reason for this relationship is that the longer an outbreak is allowed to progress unimpeded (with there being delays in public health response) the more time it will have had to grow exponentially. Indeed, epidemics often show early exponential growth (see later). Because this relationship was so striking it was possible to derive a reasonable estimate of the transmissibility of smallpox based on an equivalent regression model (31) (Fig. 1).

There are also analytically tractable response forms other than linear and exponential that might be appropriate (e.g. quadratic, logistic); right through to full non-linear regression. The fitting process that underpins regression relies on assumptions about the probability distribution of the measurement error (e.g. that the

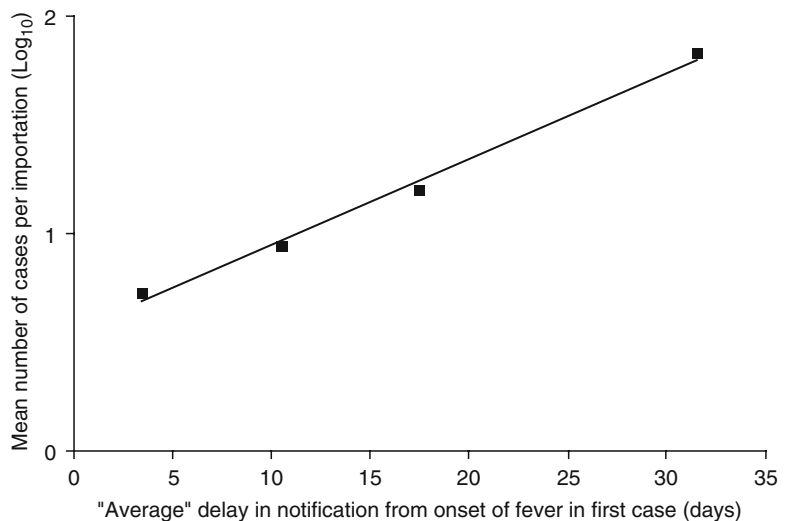


Fig. 1. Fitting a linear regression model to infectious disease data showing the strong relationship between the extent to which smallpox outbreaks expanded as delays in the public health response increased. Regression model: $y=0.0397x+0.5519$, where y is the average size of outbreaks and x is the delay.

deviation of the observed data from the modelled prediction is explained by the normal distribution introduced above). Fuller explanation of regression models is outside the scope of this chapter and more detail can be found elsewhere (32). However, another familiar example of a relationship (model) that is highly non-linear which often has great utility in virology and bacteriology is that between the dose of infectious agent received by an individual (often an experimental animal) and their subsequent probability of infection or death. Usually the probability of becoming infected increases as the dose increases, with the non-linear trend to some extent linearised by expressing the probability of infection or death as a probit transformation (33) plotted against the log of the dose. On the basis of such analyses two key parameters are determined that can be extremely useful when developing more complex models. The first is the infectious dose (ID_{50}) or lethal dose (LD_{50}) that gives a 50% probability that an individual might get infected or die, respectively. The second is the extent to which this changes on the probit scale with every log increase in the dose of organisms administered (i.e. the slope). The infectious dose (50%) varies considerably for different microorganisms and can be important to parameterise appropriately when considering epidemic models. That for *Bacillus anthracis* spores, for example, might be of the order of 10^4 spores, whilst that for *Francisella tularensis* or smallpox virus has been estimated to be of the order of 10 cells/virions or so.

Moving on to more complex infectious disease models, there is an extensive ecological and epidemiological modelling literature that underpins this subject (see reviews in bibliography and references therein). The simplest models of practical utility for emerging infectious diseases are possibly best immediately understood through their schematic descriptions; as the so-called SIR (Fig. 2a), or SIS models (4, 5).

To better understand these, first imagine simply a disease that infects a person for 1 day only and in that time they infected two other people. This means that if we start on day 1 with one case, on day 2 we will have two more; and on day 3 four, day 4 eight, etc., until on day 31 we have more than one billion cases. This process can be represented mathematically by Eq. 1.

$$I_{n+1} = I_n + \beta I_n - \gamma I_n = [1 + \beta - \gamma]^{n+1} I_0 \quad (1)$$

Note that in the simple example given above (γ) *gamma* (the reciprocal of the number of days between infection of a case and their recovery [assuming they are equally infectious throughout] which may be inferred from the generation time or the serial interval data [defined later]) = 1 and (β) *beta* (number of people infected by each case per day) = 2; and I is the number infected, and n denotes the day of the outbreak. Whilst clearly this approach provides only

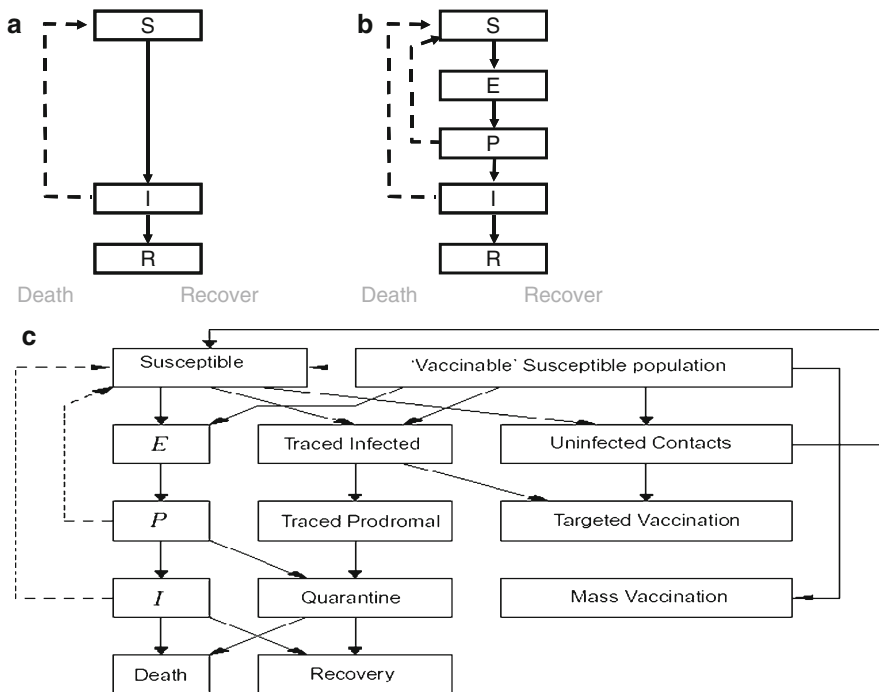


Fig. 2. Three different schematic representations of disease transmission models, from more a simplistic one (a) to a more complex one involving more disease states (b) to one involving a range of public health interventions, in this case for smallpox (c), (see text).

a toy model, this process of exponential growth is often observed (to a first approximation) in real-world epidemic data at the start of outbreaks, as in the smallpox example above. However, common sense tells us that epidemics will not spread that quickly and unchecked through the entire world's population. One reason for this is human contact behaviour. In the scenario above, for example, one of the four cases occurring on day 3 would have had a chance of meeting the same person as one of the other cases (i.e. shared contacts in social groups). Further, in a population of limited size the steadily increasing number of infectious cases would dictate that quite soon some will tend to meet other infected cases (or previously infected and immune individuals) rather than susceptible individuals. Consequently, the number of new contacts that result in new infections at each generation must depend both on a contact rate and a probability that the contacts will be with individuals still susceptible to infection.

This phenomenon can be captured by simply dividing a closed population into two compartments; *S*, meaning that proportion of the population that is (s)usceptible to infection and *I*, that part that is (i)nfected and (i)nfected (note that no distinction is made between these two). This process may be represented by the differential Eq. 2 (note in this instance cases are assumed to become

susceptible again after infection wanes) and solved explicitly as in Eq. 3, taking a logistic form, where t is time and I_0 the number infected at time 0, and $I(t)$ the number infected at time t .

$$\frac{d}{dt}I(t) = \beta[1 - I(t)]I(t) - \gamma I(t) \quad (2)$$

$$I(t) = \frac{(\beta - \gamma)I_0}{\beta I_0 + [\beta - \gamma - \beta I_0] \exp[-(\beta - \gamma)t]} \quad (3)$$

Whilst mathematically tractable such a model is still a fairly limited representation for most diseases, particularly for those diseases where cases become immune to further infection, at least for a time (or die), for example, influenza, smallpox, and SARS CoV. To account for this the schema can be extended by dividing the population into three compartments or classes: with S , and I , as defined above and R , that part that is (r)emoved (i.e. immune or dead). The distinction between immunity and death is generally of no importance to the dynamics of the model unless considered over a time period much longer than the timescale of the disease or a single epidemic. If longer timescales are important to the problem under investigation (as with assessing vaccination programmes for vaccine-preventable childhood diseases such as measles and mumps) then disease mortality has to be factored into the schema, along with deaths from all causes and also new births, the latter providing new susceptibles through the relevant birth rate (4). Factoring such things as mortality and hospitalisation into models can clearly also be important from other perspectives, for example, when estimating the impacts on society and health-care systems. Clearly, the severity of diseases such as smallpox with a case fatality ratio estimated to be in the range of 30% would result in rather different set of consequences compared to a disease such as pandemic influenza with its usually much lower estimates for the case fatality ratio (only up to about 2.4% in the main wave of the 1918/9 pandemic and much lower in the 1957 and 1968 ones).

Such simple SIR models have been used to great effect and can be expressed most simply as a series of differential Eq. 4 that describe the time-dependent transition of proportions of the population through these stages (e.g. from I to R at rate γ in Eq. 4, where β and γ have same definitions as in Eqs. 1 and 2).

$$\begin{aligned} \frac{d}{dt}S(t) &= -\beta S(t)I(t) \\ \frac{d}{dt}I(t) &= \beta S(t)I(t) - \gamma I(t) \\ \frac{d}{dt}R(t) &= \gamma I(t) \end{aligned} \quad (4)$$

Such models are also dependent on that part of the population that is infectious being able to transmit infection (dotted line, Fig. 2a) to that part that is susceptible at some probability or rate (e.g. between I and S with transmission rate β in Eq. 4, β being a composite of the number of contacts made per day and the probability that transmission occurs given that a contact is susceptible). As noted above, this rate will change over time as it depends both on the number of infected individuals and the number of susceptible individuals, as well as the more “intrinsic” transmissibility of the infection. The latter is often described by a “fundamental” parameter of many epidemic models, usually referred to as the basic reproductive number or ratio, designated R_0 , defined in Eq. 5,

$$R_0 = \frac{\beta}{\gamma} 1 - I_0 \tag{5}$$

The simplest understanding of this parameter is the average number of secondary cases caused by each primary case within an entirely susceptible population (and in the absence of public health interventions). The effective reproductive ratio, R_E (Eq. 6), on the other hand, has a similar description except that the level of susceptibility to infection within the population and the effects of public health intervention are taken fully into account. Thus, R_E changes over time as the relative proportions of the infectious and susceptible population change over the course of the epidemic (Fig. 3),




Fig. 3. Illustrative dynamics of a simple SIR model showing the change over time in the number of individuals that remain susceptible and then become infected and ultimately recover.

$$R_E = \frac{\beta S}{\gamma} = SR_0 \quad (6)$$

Even though the SIR Eq. 4 is not amenable to explicit temporal resolution, they are amenable to approximation in certain phases of the epidemic (e.g. the early exponential growth discussed above) and numerical solution. Some of the further mathematical analysis that is possible on this Eq. 4 can provide fundamental insights into aspects of the expected severity and prospects for control of epidemics. The key parameter here is R_0 (Eq. 5), which can be shown to be the parameter that defines the stability of the system (i.e. whether the disease is likely to become a major public health problem or not); since if it is greater than 1, the introduction of cases of disease into a population will likely cause an epidemic, whereas if it is less than one, the introductions will fade-out. Thus, from this parameter one can estimate the proportion of a population (or the population number if this proportion is converted by reference to the population size) that might need to be immunised to control an infection (V , i.e. that proportion that is required to bring R_E below 1 – Eq. 7, Fig. 4), and also, one may derive (4), the likely final size of an uncontrolled epidemic in a closed population (Eq. 8, Fig. 4), where R_∞ is the final attack size.

$$V = \frac{R_0 - 1}{R_0} \quad (7)$$

$$1 - R_\infty = S_0 \exp - R_0 R_\infty . \quad (8)$$

All of this assumes that the various approximations to real life that the model employs (some of which will be covered later) still allow meaningful interpretation of the model output in relation to the real setting in question. Sufficiently often these approximations do not completely compromise the results and allow useful observations to be made. A simple corollary of the relationship in Fig. 4 regarding final attack size is that in the range of R_0 between 1 and 2, one can relatively robustly infer R_0 from final attack size, and vice versa. This range of R_0 between 1 and 2 is relevant to the case of past pandemics of influenza, such that useful comparisons can sometimes be made between these two measures. For R_0 greater than 2, however, the discrimination between final attack sizes for different R_0 becomes much less and often within the bounds of the error in the data that might be available to independently determine final attack size or R_0 . Therefore, for diseases such as smallpox (R_0 in the range 3–6), SARS CoV (R_0 about 3–4) and measles (with one of the highest estimated R_0 's for an infectious disease, variously reported to be in the range 12 to in excess of 20), the inference of R_0 from final attack size would likely be much less clear. It is probably worth mentioning in passing that the estimates of R_0 given above vary, even for a single




Fig. 4. The simplistic dependence of two basic epidemic disease features on the fundamental disease parameter R_0 .

disease, because they will depend on time, place, and context of the study where R_0 was inferred. R_0 will be influenced, for example, depending as it does on contact rates (see above) by factors such as overcrowding and socio-economic conditions. Even for diseases with lower R_0 such as pandemic influenza, assessments are actually made more complicated by the fact that a good proportion of cases won't seek medical attention and others may well become infected and to some degree infectious whilst remaining asymptomatic. These individuals do not therefore get counted among the clinical cases (and so the final attack size may not actually be directly observed), though they might be observed through changes in their immune status if this was to be measured by serological surveys; or, more fundamentally, as a consequence of careful interpretation of the underlying epidemic dynamics (34). It is thought, for example, that maybe only about 60% of individuals infected with influenza actually develop reportable clinical symptoms, whilst by comparison cases of smallpox and SARS CoV are rarely, if ever considered, to remain completely asymptomatic following infection, which makes the epidemiology somewhat easier to interpret in these latter cases.

For the many problems that are not amenable to such explicit treatment one must turn to the numerical solution of these equations. As before, one must define the initial conditions (i.e. what

proportion of individuals are infected, immune or susceptible at the point of introduction of the infectious agent) and we also require a so-called equation of state, $S + I + R = 1$. Typically, for a new or emerging infectious disease initial conditions are given such that no people are immune (R at time $0 = 0$) and only a small proportion are infected; though other situations can be readily investigated by adjusting the proportions for S , I , and R accordingly. Numerical schemes essentially make the continuous differential equations (like Eqs. 2 and 4) discrete with respect to time so that they become difference equations, the precise form dependent on the accuracy of the solution demanded. Difference equations introduce a time step h and the accuracy of the numerical solution to the exact one is inversely dependent on this time step. The choice of numerical scheme is dependent on the form of the equations being approximated and the available computational resources. Essentially, however, the initial conditions will be substituted into the equations (at time 0) and the results calculated for the first time step. These results are then fed back into the equations as the starting conditions for the next time step and so on until the results for sufficient time steps have been calculated to describe the required course of the epidemic. Further discussion of such methods (e.g. Euler, Runge-Kutta, etc.) can be found elsewhere (35).

6. Some Problems with Simple Models

The mathematical formulation of the SIR schema, based on the series of differential equations described above, makes a number of implicit assumptions that need to be appreciated. The first is the often criticised one of homogeneous mixing such that an infected individual has an equal probability of infecting any one susceptible individual in the population as any other. As will be discussed later, this simplification can be addressed by introducing population heterogeneities, such as the probabilities of different age classes mixing with one another or the different geographical limits that might reasonably apply to population mixing over longer distances.

Differential equation sets are also deterministic (“clock-work”), that is, each time the model is run from the same starting conditions and with the same parameters it will produce exactly the same results across each and every time step for the entire “epidemic.” They are also continuous which means that they will allow fractional people to be counted among the cases. Real epidemics, however, are prone to stochasticity, based on individual events that occur probabilistically, which if neglected can present a major issue, particularly at the start and end of epidemics. For example, although the average number of secondary cases caused by each primary case (i.e. R_0) might be observed to be around 2,

say, for pandemic influenza, for any one individual case this might vary from 0 to some rather larger number than 2 depending on circumstances (e.g. the number of contacts the particular case might make with others, the concentration of virus a specific individual sheds and for how long). Thus for diseases with a R_0 closer to 1, say between 1 and 3 for diseases such as pandemic influenza, pneumonic plague or SARS CoV, a one off introduction of a single case of disease into a population would have a much greater chance of causing no further cases and not starting/contributing to an epidemic than for diseases where the R_0 is much larger, such as measles, where there would be a much greater probability of one imported case causing at least one further infection. Diseases with low R_0 therefore have a greater chance of experiencing what is termed stochastic fade-out and this can be extremely important to capture appropriately in models depending on what is being investigated. Similarly, as will be discussed in greater detail later, whilst there might be some concept of an average infectious period (i.e. length of time in I), the duration of such disease states will also vary between cases. Individual-level variability in features such as infectious and/or incubation period has already been discussed previously and where relevant can often be best captured by modelling them with lognormal or Gamma distributions. Differential equation-based models can be formulated within stochastic frameworks to take such individual-level variability into account, along with allowing for the concept of whole, discrete individuals rather than fractional ones to be enabled. The ways in which such formulations can be achieved (5) are beyond the scope of this chapter, but it is worth pointing out that the resulting models will usually have to be run large numbers of times (often 100–1000 s depending on the number and range of uncertainty on parameters that has to be stochastically varied) to generate a whole family of epidemics in order to ensure that a representative selection is collected. These then need to be statistically analysed to better understand the problem being investigated. All of this often increases the computational cost of such models.

7. Introducing Better Descriptions of Disease Natural History: Non-uniform Infectiousness

It may be sufficient for the purposes of a model to simply employ the concept of an average generation time or serial interval; the latter being the observed time between onset of specific symptoms in one case and the onset of the same symptoms in the subsequent cases caused by that case; whilst the former is the time between the infection of a primary case and the infection of each of its secondary cases (36, 37). In the models discussed so far observational data that are related to these intervals are often used as a surrogate for the period of time spent in the I class, whilst at the same time

naively assuming people are equally infectious throughout this period. This relatively unrealistic assumption of uniform infectiousness over time can, however, be solved reasonably well by introducing more infectious compartments into the model schema, each of which can be attributed different infectiousness, I_1 , I_2 , I_3 , and so forth. Thus for influenza it might be appropriate to have an I_1 class to cover the first 24 h following symptom onset that is more infectious with individuals then passing into subsequent I_2 , I_3 , etc. classes of defined duration that are progressively less infectious. Influenza infectiousness is thought to peak very abruptly and then decline somewhat more slowly (19, 38).

8. Introducing Better Descriptions of Disease Stages

The simple SIR model structure (Fig. 2a), whether deterministic or stochastic, can be made more realistic in other ways. For example, as discussed previously, individuals do not necessarily progress from being susceptible to being infectious (and symptomatic) without some intervening latent period. Accepting this potentially alters the observed dynamics of epidemic models in ways that may or may not be important to the specific questions being asked (5, 36, 39). For diseases such as smallpox, for example (Fig. 2b), it might be important to consider five separate disease classes and extend the differential equation set accordingly (34); though there are many ways in which the following aspects of the disease natural history might be reasonably represented with fewer or even greater numbers of classes; for example, to better capture the time varying infectiousness discussed above (5). Thus, there could be a period between infection and the first non-specific symptoms (often referred to as the (E)xposed class), then the period with non-specific symptoms (the (P)rodromal period), followed by the (I)nfectious and (R)ecovered classes. It can be important to capture the natural history in this way for diseases such as smallpox, which have both prodromal and infectious periods, since in the case of smallpox both are infectious (signified by the two dotted lines in the Fig. 2b) but with the latter much more so than the former, according to our knowledge and analysis of previous smallpox outbreaks (30, 40).

9. Some Problems with Data and the Parameterisation of Disease Natural History

The specific timings of events at an individual level can be highly critical, especially the relative infectiousness through the infectious period, which is rarely, if ever, uniform. This can be

important; for example, in relation to modelling public health interventions, which if applied early and before the peak infectiousness of each case will clearly have more impact on the control of the onward spread of an infection than if applied later. In which case it is important to better understand and appropriately capture the disease natural history in models (Figs. 5 and 6). Thus, when modelling pandemic influenza, it is thought that if antiviral drugs (such as the neuraminidase inhibitors) are to have much impact on the onward transmission of disease (through minimising viral replication and viral loads in the secretions, as opposed to simply ameliorating the course of infection and reducing the probability of hospitalisation and death (41) the drugs probably need to be administered within (a challenging) 12–24 h of symptom onset because of the extremely short infectious period of influenza and the rapid rise and fall of viral titres and infectiousness (19, 22). Whilst serial interval times are often more easily




Fig. 5. Timelines of infection for two different viral diseases, given roughly to scale.

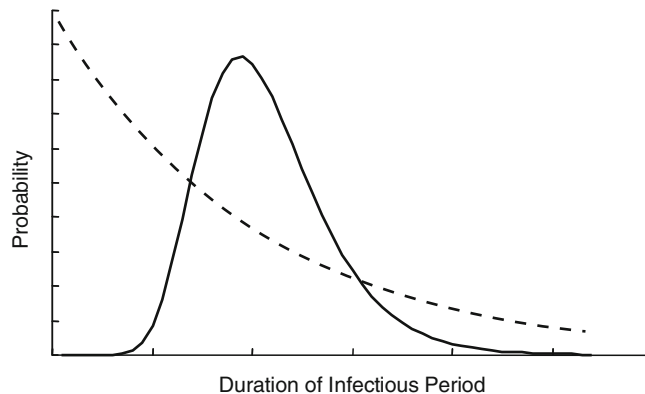


Fig. 6. Two different distributions with the same mean that models sometimes use for representing the duration of the infectious period, one exponential (dashed line) the other lognormal (solid line).

observed than some other intervals in the disease process, they are in themselves convolutions of the other intervals (periods), or parts of them. For example, as discussed previously, there is usually a period between the initial acquiring of infection by a case to the onset of symptoms (incubation period) and to the onset of infectiousness (latent period), which may or may not be coterminous periods and will vary between persons and the disease (Fig. 5). This distinction between incubation and latent period can be extremely important since diseases that become infectious before the onset of symptoms can make them much harder, or impossible, to control through the traditional means of isolating cases and quarantining contacts. Thus, diseases that can be asymptomatic at points during infection or relatively mild overall, or chronic or recurring such as influenza, HIV, or tuberculosis can theoretically be much harder to control by these means than diseases such as SARS CoV or smallpox (18). In this respect, the eradication of SARS CoV from the human population was highly dependent on (and blessed by), amongst other things, the fact that cases were generally not significantly infectious before showing symptoms and had a reasonably long incubation and infectious periods. Incubation/latent periods can often be followed by other defined periods that are relevant to the disease natural history, including: a symptomatic period, during some part of which cases are usually their most infectious (infectious period) – though, infectiousness is likely to vary over time, often rising rapidly to a peak and tailing off more gradually. For some infections there may also be a prodromal period (involving non-specific disease symptoms) between the incubation and symptomatic periods that might also be infectious (Fig. 5).

At its simplest, the parts of these periods that contribute to, for example, the serial interval are the incubation period of a secondary case and some part of the infectious period of the primary case. The latter being the time until a relevant contact has been made between persons such as to permit the transmission of infection, which will also depend on other factors, such as variation in contact rates with the rest of the population. Deconvoluting generation time distributions into their constituent distributions (or attempting the inference of the other distributions) can prove problematic to achieve in a statistically rigorous sense. Details of this are outside of the scope of this overview but such problems and their implications for epidemic models are important given what has been discussed already concerning modelling the prospects of ameliorating or controlling (or not) outbreaks of diseases such as pandemic influenza, SARS CoV, and smallpox. Technical reflections on this subject have been discussed in depth elsewhere (36, 39).

Thus, although it might seem trivial to try to obtain the distributions of the different periods by direct observation, actual

measurement of some of these processes and their distributions can be problematic. For example, quite often the point in time when infection occurs (necessary for the estimation of generation interval and incubation period) is usually not observed in population-based studies for perfectly understandable and pragmatic reasons, except in those rarer situations where one person can reliably be known to have had only one contact with a single case and at a single point in time. Even then the precision of such observations is often limited to being differentiated to the nearest whole day. This will obviously matter more for diseases that have shorter generation/serial interval times: for example, a day can be a relatively long time in the course of influenza compared to smallpox (Fig. 5). Depending on the natural history of the disease, and possibly the prevalence of the disease in the rest of the community, such observations can be more or less feasible. For diseases such as smallpox or SARS CoV that possibly have a better marked clinical course in relation to infectiousness it is easier to define the timing of contacts with cases in relation to disease symptoms, as long as the disease is not so prevalent in the rest of the community that it is difficult to identify infected contacts uniquely with respect to the case that was responsible for them. Smallpox was also a disease for which contact tracing and quarantine was an important part of controlling outbreaks and so observations of the timings of contacts were more routinely made and, happily, sometimes recorded for posterity (30). Useful observational data can be more difficult to obtain for other diseases where the course might be either more rapid (influenza) and/or less well defined (influenza, measles, rubella) with respect to the infectiousness of cases in relation to their symptoms, and especially when the number of infections more widely in the community might also be quite high (such as for influenza) so as to potentially “mask” unique infection events. This problem can sometimes be overcome to some extent; for example, in the case of influenza this has been achieved by rigorous statistical analysis of studies undertaken in defined contexts, such as households where the time between subsequent infections can be more easily inferred (11), or volunteer challenge studies (38) where the time of infection is known.

10. Further Improvement of Descriptions of Disease Natural History: Non-exponential Disease Periods

The differential equation sets described above also hold an implicit assumption that the residence times in the I class, although having the correct mean duration, are exponentially distributed (they exhibit so-called Markovian dynamics because the result at a given point in time depends only on the state at the previous time as no

other history is encapsulated in the model) as opposed to something more realistic (e.g. lognormal or Gamma). As can be seen from Fig. 6, this means that although the average residence time in I is correct from both distributions (i.e. same mean for exponential and lognormal), an appreciable proportion of the residence times for the exponential will be unrealistically short (left hand end of distribution – dashed), and another proportion will be unrealistically long (right hand tail of distribution – dashed).

The consequences of adopting such simplifying assumptions may matter to a greater or lesser extent but is essentially a mathematical convenience to improve the tractability and computational ease of the problem. This assumption can, however, at a computational cost, be revised in a number of ways that allows the utilisation of more reasonable distributions. The different means of achieving this are largely beyond the scope of this chapter (5, 42). However, one simple approach is to adjust the schema described earlier and to arbitrarily break the I class down into more than one compartment (and therefore introducing another equation and term for each class). The use of several sequential equations and classes rather than one, with each class having an implicit exponential distribution for residence time, will overall combine to approximate a Gamma distribution (more like that in Fig. 6, solid line) that will also be closer in form to the distribution observed in the data.

11. Introducing Better Descriptions of Transmission

The transmission process can also itself be implemented mathematically in more than one way, for example, as a probability or rate determined by the mean estimate of the quantity R_0 or, as with the residence times in each of the disease classes mentioned above, as a more realistic distribution based on prior observations. This is sometimes represented as an “offspring distribution” (10), and often given as a distribution that closely follows observed transmission events. This approach usually better represents the variability that is observed in the transmission process. This is because, depending on the mathematical implementation, simple usage of the concept of average transmission, R_0 , and, for example, implementing this as the mean of a Poisson process, can underestimate the potential impacts of low probability but high transmission events (“super-spreaders” or “super-spreading events”), and also the high probability but low transmission events. That is, a more reasonable distribution to use would probably have more dispersion than the Poisson that frequently gets used in mathematical formulations. This is explored in Fig. 7, which relates to observational data on pneumonic plague (31), which is transmissible person to person at relatively low average

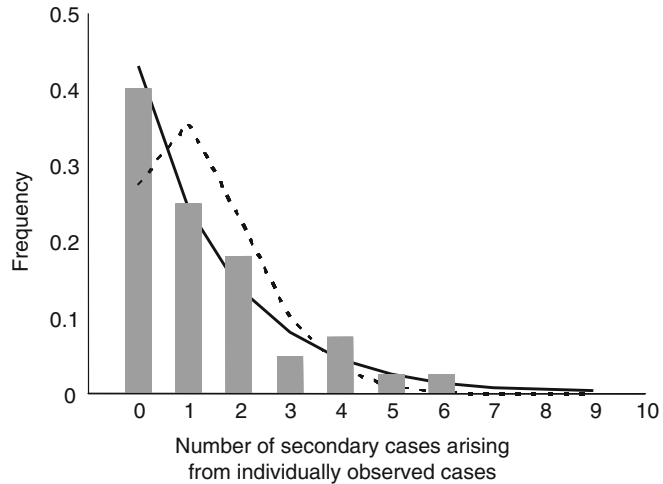


Fig. 7. Offspring distribution for pneumonic plague. Data are the vertical bars; the geometric distribution (*solid line*) is a better fit to the data than the Poisson (*dashed line*).

probability ($R_0 = c \cdot 1.3$). However the data (vertical bars), which are the frequencies (y -axis) with which infected individuals have been observed to infect one, two, three, and so on susceptible individuals (x -axis) can be seen to be better reproduced by a more dispersed geometric distribution given by $f(x) = p(1-p)^x$, where x is the number of secondary cases per primary case, $f(x)$ is the frequency, and $p = 0.43$ (solid line) than a Poisson with an equivalent mean (dotted line). It can be seen from Fig. 7 that there is a greater probability than would be predicted from the Poisson of no transmission occurring from an infected case, and a greater probability of 4 or more secondary cases occurring from a primary case.

The former observation means that if a Poisson was used in a model then there would be a somewhat smaller probability of an epidemic dying out if there were only a very few initial cases, and the latter observation would mean that there would be a smaller probability in the model of generating larger outbreaks purely by chance (31).

12. Introducing More Realistic Descriptions of Public Health Controls

The model structures discussed so far are useful in deriving a better understanding of some aspects of “free-fall” epidemic dynamics, but have not really been discussed so far in relation to assessing the potential impacts of public health interventions, except in other than fairly simple ways. For example, as described already, in relation to calculating the proportion of the whole population that might need to be vaccinated in order to stop transmission

and eradicate a disease. As can be seen simplistically from Fig. 4, the higher the estimate of R_0 for a particular infectious disease the higher the proportion of the population that needs to be vaccinated to create sufficient “herd immunity” to prevent transmission; that is to bring R_E below one. For a pandemic of influenza with an R_0 value of 1.6, for example, this could be as little as around 37% of the population (43), but for smallpox with an R_0 value in the range of 3–6 this might need to be 67–80% or more, respectively (15). As stated previously R_0 can of course depend on local conditions and can vary geographically; often being higher if transmission is promoted by overcrowding and lower socioeconomic conditions. A similar approach could also be simplistically extended to the concept of isolation or quarantine, to estimate what proportion of infected cases and contacts of cases have to be found and completely isolated before they themselves become infectious. This number is similar to the critical vaccination coverage if a simplistic view of case finding, contact tracing and efficient quarantine is assumed (5). Model realism, and therefore complexity, can, however, be extended to investigate public health interventions more directly by allowing other states within the model schema and equations. For example, Fig. 2c shows an extension of the earlier SEPIR structure for smallpox to allow for a whole range of public health interventions, such as mass vaccination of some proportion of the population, case finding/reporting, contact tracing, targeted vaccination, and so forth. The equations exemplified earlier would of course have to be extended to cover these other compartments. All of these processes of course have to be understood and parameterised accordingly; one example being what proportion of the contacts of cases might reasonably be expected to be found and at what point in the course of their disease. This might be based on previously recorded experiences with the disease or with a disease of similar natural history. For diseases such as smallpox, pneumonic plague, SARS CoV, and bacterial meningitis, for example, the finding of potentially infected contacts can be relatively efficient and can often be in excess of 80%, reflecting the relative ease of finding the majority of those who have had sufficient contact with a case to facilitate disease transmission. Those predominantly getting infected tending to be those in (or visiting between) households and those in health-care facilities (see later for a discussion of transmission of pneumonic plague in different settings and contexts). Alternatively, or more likely additionally, robust analysis to parameter uncertainty (sensitivity analysis) would be undertaken. Hence, the parameters related to public health controls, as well as those related to disease natural history, would be systematically varied within ranges considered to be plausible, ensuring appropriate sensitivity and/or scenario analysis was performed [e.g. as in the case of the smallpox studies referred to earlier (20, 23)].

13. Introducing Better Descriptions of Population Mixing vs. Homogeneous Mixing

Returning to the assumption of homogeneous mixing referred to earlier; this is clearly not an entirely reasonable one, and for some purposes may considerably invalidate the use of a model depending on its application. If sufficient is known concerning the contact patterns of groups of individuals, or at least the contexts in which transmission occurs relative to one another, then such features can usefully be incorporated. This can be achieved either through splitting the population into a number of specific groups that share particular defined features (metapopulation models) or indeed into the more computationally demanding concept of individuals, each of which will have some generalised set of (measurable) features that are to some extent different to other individuals, but in combination with all of the other individuals in the model together reflect the characteristics of the population as a whole. These latter models can be implemented as either what are sometimes known as individual-based microsimulation models or network models (5, 19, 20, 22). Metapopulation models introduce heterogeneity into the mixing patterns of the population by identifying specific groups of individuals, within which there is still homogeneous mixing, but where between them there is not. Network and individual-based microsimulation models, necessarily generalise about individuals in some rational way, attributing features to each individual appropriately, and then allowing for heterogeneity of mixing at an individual-based level in respect of those features (5).

Metapopulation approaches are often much more computationally tractable, the extent to which this is true being dependent on the size of the population being simulated and the number of patches into which the population is subdivided. Such models may also have fewer parameters and therefore be easier to more reliably parameterise and understand. Different types of metapopulation approach have been used to good effect. Metapopulations can, for example, be developed on the basis of breaking the population down into different age classes, different economic, social, or functional contexts and/or by geography.

14. Age-Structured Models

Age-structured models at their simplest essentially take the SIR (or more complex compartmental)-type approaches but have a series of parallel schema (and sets of equations) running for each age group in the model (Fig. 8, for a simplistic two age class model), with transmission terms (dashed lines) not just operating




Fig. 8. Simplistic schematic for an age-structured model with two age classes (one *grey*, the other *black*).

within a particular age group, but also between different age groups at potentially different rates.

Parameterising such age-structured models requires some knowledge of the facility with which individuals in each age group infect others in that age group and in each of the others and this will depend to an extent on aspects of the natural history of the disease (e.g. the mode of transmission and the degree of intimacy of contact required; the severity of illness and the degree to which infectious individuals continue to be able to mix). With caveats, such transmission matrices can be estimated by attempting to fit models to infectious disease data where this has been stratified by age group or by reference to data that has been recorded on the relative extent to which different age classes mix with one another and with themselves, and ideally taking into account the intimacy of the contact (e.g. face-to-face conversation of some duration or some level of physical contact). These matrices have been used for some time in relation to developing a better understanding of the dynamics of, for example, childhood vaccine-preventable diseases and informing on optimal vaccination strategies. The derivation and use of such WAIFW (who acquires infection from whom) and similar matrices is beyond the scope of this broad overview but a good introduction can be found in Anderson and May (4), along with more contemporary analyses based on more recent multi-centre European studies (37, 44–46). Essentially mixing among age groups is highly assortative, that is those closer in age tending to mix more frequently with one another than with those in other age groups (but with children also mixing with parental and sometimes grandparental age classes), but with the frequency of contact between children generally being higher than mixing within any other individual age group. This potential disproportionate mixing and potential transmission of disease can have important consequences which are often important to capture in a model. For example, the initial rise in cases during an emerging uncontrolled epidemic might be seen first as a rise in the number of cases in children ahead of the rise in other age groups. Further,

such assortative mixing can have important consequences in relation to potential public health controls such as which age groups to prioritise for vaccination or the value of closing schools in order to try to limit the disproportionate contribution of children to overall disease transmission. For pandemic influenza it is thought such effects are likely to be important, since, for seasonal influenza at least, children do seem to be particularly implicated in transmission (47–49). For smallpox and SARS CoV on the other hand this is much less certain. In the case of smallpox it would seem from the limited historical observations available on populations that had not experienced smallpox for some considerable time (and so no older members of the community were already immune) that the age distribution of cases matched that of the population itself (21, 30). Even when age-dependent assortative mixing is relevant it is important to remember that public health measures such as school closure, for example, would not necessarily reduce the contact rate of that particular age group to zero unless draconian measures were also introduced to prevent them mixing in other contexts out of school. They would be likely to continue to mix to some extent in other contexts such as playing together outside and in other households, and also mix more frequently than before with other age groups such as their parents, household members, and relatives. Such effects are much more difficult to parameterise reliably or to compensate for, but are likely to be important and should not be neglected. Such effects and parameter assignments can in some instances be clarified to an extent by referring back to data from natural experiments and fitting age-stratified rates of influenza-like illness over time during an ongoing epidemic, such as happens, for example, around school holiday periods (47).

15. Socially Structured Models

Depending on what is known about the disease and the purpose of a model, the extent or facility with which disease transmission might occur might also be usefully characterised and subdivided by some social or functional context (e.g. household, workplace, school, hospital). Such contexts have and can be incorporated into a variety of types of model of differing degrees of complexity, including metapopulation and individual-based microsimulation ones (19, 20, 22). For ease, however, they will be considered here more simply in terms of metapopulation models. Taking pneumonic plague again as an example it is clear from historical data that not all contexts and inter-personal relationships were equal in terms of the extent to which transmission was observed to occur (Fig. 9).

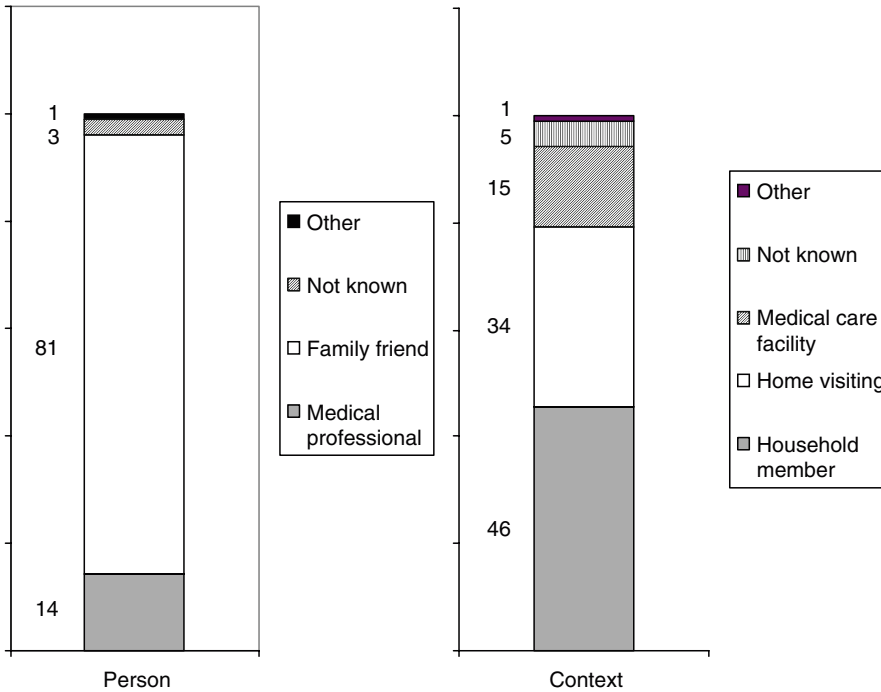


Fig. 9. The contexts in which the transmission of pneumonic plague was observed to occur.

The most frequent context for transmission was within a household (either with another member of, or visitor to, an infected household) or within a medical care facility. Hence, by far the most frequently infected individuals in this case (31) were family relatives and friends followed by health-care workers (together accounting for about 95% of transmission events). This observation in itself probably accounts for the fact that outbreaks of pneumonic plague (with a low overall reproductive potential) were readily brought under control since the infectious contacts of cases were relatively easily identified and quarantined such that R_E was rapidly reduced to below 1. Further, for smallpox, in the latter parts of the eradication era in Europe and other more developed parts of the world, transmission within the hospital context accounted for a significant proportion of all transmission (about 50%) before the disease was correctly identified and subjected to appropriate local infection controls (15, 30). Such contexts as described above clearly have parameters that relate directly to observations and data, and as such can be specifically incorporated into models with each context being represented by a separate metapopulation within the overall model structure, in much the same way as has been described already for age-structured models. In this way the relative frequencies of transmission seen in the data are then replicated by the model in the correct contexts. The contexts of home, workplace, school, etc. have also

been employed in a more sophisticated way within individual-based microsimulation models, where individuals in the model have attributed to them particular home, work, and school locations/interactions (20, 22).

16. Geographically Structured Models

Another degree of complexity that can be introduced into metapopulation models is the concept of separating the overall population into different geographical (or spatially determined) units. This can be done on the basis of relevant administrative areas such as those utilised during the collection of census information. The resulting metapopulation model in principle is not unlike that shown diagrammatically for simpler age-structured models in Fig. 8. In this case, however, each geographically distinct entity might have its own SIR (or more complex) structure, but with the connections (dotted lines in the figure) and probabilities of infection between geographic units parameterised by the extent to which proportions of the populations move between them. These connections may be viewed as largely analogous to the WAIFW matrices described earlier in relation to age-structured models, though with a typically much larger matrix that links each geographic unit with all of the others, and may have dependence on the time of day (to allow for commuting behaviour). Models that reasonably capture space in this way can be crucial when it comes to investigating interventions that have to be given a spatial context, such as vaccinating all the individuals in some geographic region based on there being cases of disease in that region (23). The concept of transmission between geographic entities can also be usefully implemented in other ways. In individual-based microsimulation models, for example, the matrix mentioned above can be converted into some more generalised movement kernel that describes the probability of any one individual moving (and/or causing infection) some distance from their home location by virtue of applying a probability based on such a kernel (20, 22). This probability typically drops off very non-linearly with increasing distance from home.

17. Closing Remarks

It is clear from the discussion above that models of varying degrees of complexity can be constructed to tackle problems related to (re-)emerging infectious disease problems and their control. Before embarking on model development for such issues it is

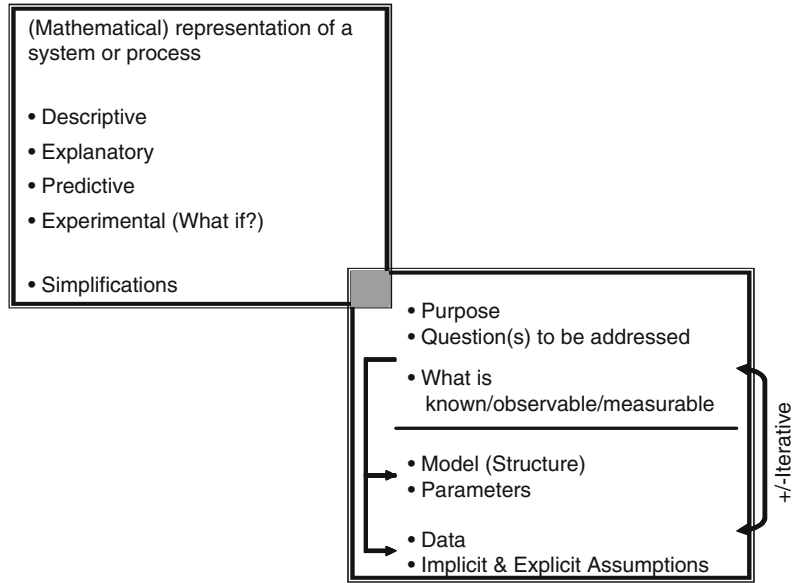


Fig. 10. Inter-relationships between policy and planning and the design of models.

generally useful for those engaged in public health protection to consider some basic practicalities. The first is that there is probably an important initial step before identifying a suitable model structure, either an existing one or one that is to be developed *de novo*, and that is to carefully consider the question or questions that are to be addressed in the light of what might be knowable, observable, or preferably quantifiable features of the disease; that is, the measurable features (parameters and relationships) of the underlying processes that are involved and what it is that the model is intended to determine (Fig. 10).

The former assertion in particular might seem rather facile, but less so when it is realised that in relation to issues of policy, planning, and responses, those requiring the answers to the questions are often not the ones that will be doing the modelling. Questions that might seem at first well specified by one, or even all, parties may fail to take into consideration some contingent factor that was not initially quite so obvious. So to take a very simple example, a question regarding what proportion of a population would need some particular intervention to achieve successful control of an outbreak may at first fail to take into account that it might be advisable to target a particular subset of the population on the basis of, for example, its geographic or demographic features. If the model has not been suitably constructed from the outset then the real question of how to optimise control policies will probably not be able to be addressed without reformulating the model. Similarly, with regards model parameters, it is entirely feasible to develop a model that turns out to require

data on a feature that has never been (reliably) measured or indeed can never be measured (e.g. time of infection for some diseases). Reasonable assumptions about such parameters can sometimes be made, but often a safer recourse is to reformulate (simplify) the model, if possible from the outset, in terms of other parameters that are measurable and for which there are more reliable data. This aspect needs careful consideration and communication among the various stakeholders in the modelling, particularly in relation to the question(s) that need to be addressed. Generally, as long as fit for purpose, the more parsimonious a model, the more readily it will be parameterised and executed, and produce results that are transparent and better able to be understood. As suggested earlier it is probably best to engage iteratively with as comprehensive a stakeholder group as possible that includes all of the disciplines that are relevant from the outset (Fig. 10). It is also important to make clear what are the assumptions and limitations of the models, and employ appropriate sensitivity and scenario analyses to mitigate such problems.

Finally, it is also equally important to set in place real-time data collection and analysis systems so as to be able to recalibrate and rerun models based on real-time data as it arises during an outbreak. Only in this way can an outbreak of an emerging infectious disease really be better understood at the time when the aim is to bring it under control (16).

References

- Bernoulli, D. (1766) An attempt at a new analysis of the mortality caused by smallpox and of the advantages of inoculation to prevent it. *Mem Math Phy Acad R Sci Paris* 1–45.
- Ross, R. (1911) *The Prevention of Malaria*, 2nd Edn. Murray, London.
- Kermack, W.O. and McKendrick, A.G.A. (1927) Contribution to the mathematical theory of epidemics. *Proc R Soc Lond A* 115, 700–721.
- Anderson, R.M. and May, R.M. (1991) *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, Oxford.
- Keeling, M.J. and Rohani, P. (2008) *Modeling Infectious Diseases in Humans and Animals*. Princeton University Press, Princeton.
- Bailey, N.T.J. (1975) *The Mathematical Theory of Infectious Diseases and Its Applications*. Charles Griffin and Company, London.
- Grenfell, B.T. and Dobson, A.P. (1995) *Ecology of Infectious Diseases in Natural Populations*. Cambridge University Press, Cambridge.
- Daley, D.J. and Gani, J. (1999) *Epidemic Modelling*. Cambridge University Press, Cambridge.
- Hethcote, H. (2000) The mathematics of infectious diseases. *SIAM Rev* 42, 599–653.
- Grassly, N.C. and Fraser, C. (2008) Mathematical models of infectious disease transmission. *Nat Rev Microbiol* 6, 477–487.
- Cauchemez, S., Carrat, F., Viboud, C., Valleron, A.J., and Boëlle, P.Y. (2004) A Bayesian MCMC approach to study transmission of influenza: application to household longitudinal data. *Stat Med* 23, 3469–3487.
- ECDC (2007) Meeting Report. Now-Casting and Short-Term Forecasting During Influenza Pandemics: A Focused Developmental ECDC Workshop, Stockholm, 29–30 November 2007. ECDC, Stockholm. At http://ecdc.europa.eu/pdf/080409_meeting_pan.pdf.
- ECDC (2007) Influenza Surveillance in a Pandemic: Paper From ECDC Working Group, August 2007. ECDC, Stockholm. At http://ecdc.europa.eu/en/Health_Topics/Pandemic_Influenza/pdf/070801_Influenza_surveillance.pdf.
- Fraser, C. (2007) Estimating individual and household reproduction numbers in an emerging epidemic. *PLoS ONE* 2, e758.

15. Gani, R. and Leach, S. (2001) Transmission potential of smallpox in contemporary populations. *Nature* 414, 748–751.
16. Ferguson, N.M., Keeling, M.J., Edmunds, W.J., Gani, R., Grenfell, B.T., Anderson, R.M., and Leach, S. (2003) Planning for smallpox outbreaks. *Nature* 425, 681–685.
17. Riley, S., Fraser, C., Donnelly, C.A., Ghani, A.C., Abu-Raddad, L.J., Hedley, A.J., Leung, G.M., Ho, L.M., Lam, T.H., Thach, T.Q., Chau, P., Chan, K.P., Lo, S.V., Leung, P.Y., Tsang, T., Ho, W., Lee, K.H., Lau, E.M., Ferguson, N.M., and Anderson, R.M. (2003) Transmission dynamics of the etiological agent of SARS in Hong Kong: impact of public health interventions. *Science* 300, 1961–1966.
18. Fraser, C., Riley, S., Anderson, R.M., and Ferguson, N.M. (2004) Factors that make an infectious disease outbreak controllable. *Proc Natl Acad Sci USA* 101, 6146–6151.
19. Ferguson, N.M., Cummings, D.A., Cauchemez, S., Fraser, C., Riley, S., Meeyai, A., Iamsrithaworn, S., and Burke, D.S. (2005) Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature* 437, 209–214.
20. Riley, S. and Ferguson, N.M. (2006) Smallpox transmission and control: spatial dynamics in Great Britain. *Proc Natl Acad Sci USA* 103, 12637–12642.
21. Cauchemez, S., Boelle, P.Y., Donnelly, C.A., Ferguson, N.M., Thomas, G., Leung, G.M., Hedley, A.J., Anderson, R.M., and Valleron, A.J. (2006) Real-time estimates in early detection of SARS. *Emerg Infect Dis* 12, 110–113.
22. Ferguson, N.M., Cummings, D.A., Fraser, C., Cajka, J.C., Cooley, P.C., and Burke, D.S. (2006) Strategies for mitigating an influenza pandemic. *Nature* 442, 448–452.
23. Hall, I.M., Egan, J.R., Barrass, I., Gani, R., and Leach, S. (2007) Comparison of smallpox outbreak control strategies using a spatial metapopulation model. *Epidemiol Infect* 135, 1133–1144.
24. Pitzer, V.E., Leung, G.M., and Lipsitch, M. (2007) Estimating variability in the transmission of severe acute respiratory syndrome to household contacts in Hong Kong, China. *Am J Epidemiol* 166, 355–363.
25. FAO/WHO Codex Alimentarius “Official Website of the Codex Alimentarius Commission” (1999) Principles and Guidelines for the Conduct of Microbiological Risk Assessment CAC/GL-30. FAO/WHO Codex Alimentarius, Geneva. At http://www.codexalimentarius.net/download/standards/357/CXG_030e.pdf.
26. FAO/WHO Codex Alimentarius (1999). FAO/WHO Codex Alimentarius, Geneva. Accessed January, 2009, at http://www.codex-alimentarius.net/web/index_en.jsp.
27. Encyclopædia Britannica (2009) “Ockham’s Razor”. Encyclopædia Britannica Online. Encyclopædia Britannica Inc, Chicago, IL. Accessed June 21, 2009, at <http://www.britannica.com/EBchecked/topic/424706/Ockhams-razor>.
28. Stuart, A. and Ord, K. (1994) Kendall’s Advanced Theory of Statistics, Vol. 1: Distribution Theory, 6th Edn. Hodder and Arnold, London.
29. Vose, D. (1996) Risk Analysis, A Quantitative Guide to Monte Carlo Simulation Modelling. Wiley, New York.
30. Fenner, F., Henderson, D.A., Arita, I., Jezek, Z., and Ladnyi, I.D. (1988) Smallpox and Its Eradication. World Health Organization, Geneva.
31. Gani, R. and Leach, S. (2004) Epidemiologic determinants for modeling pneumonic plague outbreaks. *Emerg Infect Dis* 10, 608–614.
32. Stuart, A., Ord, K., and Arnold, S. (1999) Kendall’s Advanced Theory of Statistics, Vol. 2A: Classical Inference and the Linear Model, 6th Edn. Hodder and Arnold, London.
33. Finney, D.J. (1947) Probit Analysis, 1st Edn. Cambridge University Press, Cambridge.
34. Hall, I.M., Gani, R., Hughes, H.E., and Leach, S. (2007) Real-time epidemic forecasting for pandemic influenza. *Epidemiol Infect* 135, 372–385.
35. Press, W.H., Teukolsky, S.A., Vetterling, W.T., and Flannery, B.P. (1992) Numerical Recipes: The Art of Scientific Computing, 2nd Edn. Cambridge University Press, Cambridge.
36. Svensson, A. (2007) A note on generation times in epidemic models. *Math Biosci* 208, 300–311.
37. Wallinga, J., Edmunds, W.J., and Kretzschmar, M. (1999) Perspective: human contact patterns and the spread of airborne infectious diseases. *Trends Microbiol* 7, 372–377.
38. Carrat, F., Vergu, E., Ferguson, N.M., Lemaître, M., Cauchemez, S., Leach, S., and Valleron, A.J. (2008) Time lines of infection and disease in human influenza: a review of volunteer challenge studies. *Am J Epidemiol* 167, 775–785.
39. Wallinga, J. and Lipsitch, M. (2007) How generation intervals shape the relationship between growth rates and reproductive numbers. *Proc Biol Sci* 274, 599–604.
40. Eichner, M. and Dietz, K. (2003) Transmission potential of smallpox: estimates based on

- detailed data from an outbreak. *Am J Epidemiol* 158, 110–117.
41. Gani, R., Hughes, H., Fleming, D., Griffin, T., Medlock, J., and Leach, S. (2005) Potential impact of antiviral drug use during influenza pandemic. *Emerg Infect Dis* 11, 1355–1362.
 42. Ma, J. and Earn, D. (2006) Generality of the final size formula for an epidemic of a newly invading infectious disease. *Bull Math Biol* 68(3), 679–702
 43. Anderson, R.M. (2006) Planning for pandemics of infectious diseases. In: Bugliarello, G. Ed. *The Bridge*, Vol. 36. National Academy of Engineering, Washington, DC, pp. 5–9.
 44. Wallinga, J., Teunis, P., and Kretzschmar, M. (2006) Using data on social contacts to estimate age-specific transmission parameters for respiratory-spread infectious agents. *Am J Epidemiol* 164, 936–944.
 45. Mossong, J., Hens, N., Jit, M., Beutels, P., Auranen, K., Mikolajczyk, R., Massari, M., Salmaso, S., Tomba, G.S., Wallinga, J., Heijne, J., Sadkowska-Todys, M., Rosinska, M., and Edmunds, W.J. (2008) Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS Med* 5(3), e74.
 46. Read, J.M., Eames, K.T., and Edmunds, W.J. (2008) Dynamic social networks and the implications for the spread of infectious disease. *J R Soc Interface* 5, 1001–1007.
 47. Vynnycky, E. and Edmunds, W.J. (2008) Analyses of the 1957 (Asian) influenza pandemic in the United Kingdom and the impact of school closures. *Epidemiol Infect* 136, 166–179.
 48. Cauchemez, S., Valleron, A.J., Boëlle, P.Y., Flahault, A., and Ferguson, N.M. (2008) Estimating the impact of school closure on influenza transmission from Sentinel data. *Nature* 452, 750–754.
 49. Halloran, M.E., Ferguson, N.M., Eubank, S., Longini, I.M. Jr., Cummings, D.A., Lewis, B., Xu, S., Fraser, C., Vullikanti, A., Germann, T.C., Wagener, D., Beckman, R., Kadam, K., Barrett, C., Macken, C.A., Burke, D.S., and Cooley, P. (2008) Modeling targeted layered containment of an influenza pandemic in the United States. *Proc Natl Acad Sci USA* 105, 4639–4644.