

The Mathematics of Vaccination

Matt Keeling*, Mike Tildesley*, Thomas House* and Leon Danon*

Vaccination is one of the major medical advances of the 20th century. Although vaccination was scientifically investigated by Edward Jenner in the late 18th century, it was not until the biochemical advances of the 20th century that cheap, safe vaccines could be produced in large quantities. Vaccination is a powerful tool in the public-health control arsenal, and allows for the mass prevention of infection rather than treating the symptoms of infection. Vaccination saves thousands of lives each year in the UK alone.

In simple terms, vaccines work by forcing the body to mount an immune reaction, which in turn generates immunity to future infections of this particular pathogen (and strain); this immunity is generally assumed to last a lifetime. In the past many vaccines contained small quantities of the infecting organism, but most modern vaccines simply contain sufficient parts of the pathogen to enable an immune response combined with chemicals (adjuvants) that promote this immune reaction. Vaccination is a multi-billion pound industry, with the UK spending over £200 million a year and offering over 10 different regular vaccines.

So where does mathematics fit into this? Various elements of mathematics are used throughout the vaccine development process; but here we focus on the use of mathematics in public health, understanding the impact of vaccination at the population level.

1 Spread of infection and critical vaccination coverage

While there are many mathematical models that seek to capture the spread of infection, the driving force behind all of these models is the transmission of infection from infectious to susceptible individuals. Key to quantifying this transmission is the average number of individuals one infectious person would infect (over their entire period of infectiousness) if everyone in the population

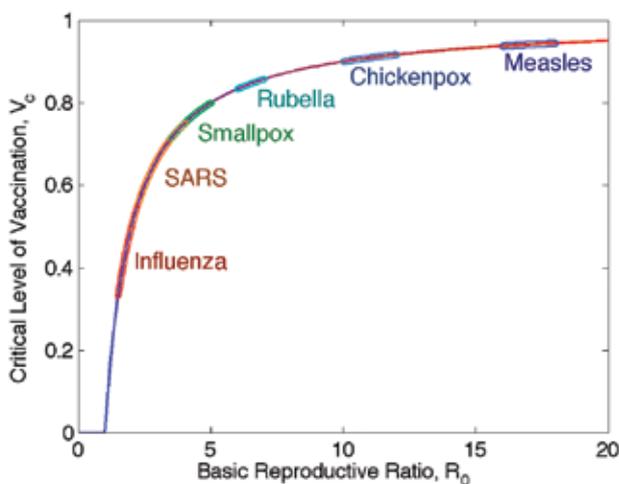


Figure 1: Theoretical impact of perfect vaccination. The graph shows the relationship between the critical level of vaccination needed to eradicate infection and the basic reproductive ratio, R_0 . Superimposed on this curve are parameter estimates for six well-known infectious diseases, ranging from influenza (low R_0) to measles (high R_0).

Box 1 The SIR model with vaccination

In the main text we talk about the Susceptible Infectious Recovered (SIR) differential equation model, which describes the rate of change of the proportion of individuals susceptible to, infected with and recovered from a given infectious disease. To this formulation we add the effects of vaccination

$$\begin{aligned} \frac{dS}{dt} &= B(1 - v_1) - \beta SI - dR - v_2 S \\ \frac{dI}{dt} &= \beta SI - \gamma I - dI \\ \frac{dR}{dt} &= \gamma I - dR \\ \frac{dV}{dt} &= v_1 B + v_2 S - dV \end{aligned}$$

where the terms in grey correspond to the natural birth and death of people in the population, the terms in red correspond to the interaction between susceptible and infected individuals leading to transmission, the terms in blue relate to the recovery of infected individuals, and the terms in green correspond to either the vaccination of infants (a proportion v_1 of all children are assumed to be immunised) or the random vaccination of individuals in the population (at rate v_2 , although this vaccination will only affect those that are currently susceptible).

were susceptible. This quantity is called the *basic reproductive ratio* and is represented by R_0 . Clearly, when R_0 is bigger than one, each case generates more than one additional case before they recover and the epidemic grows; whereas when R_0 is less than one, epidemics are doomed to rapid extinction. Simple mathematics can inform about how vaccination perturbs this pattern.

Two approaches are possible: one constructs simple differential equation models for the number of susceptible, infected, recovered and vaccinated individuals in the population (this is a modification of the well-known SIR model, see Box 1); the second, which we describe now, is based on intuitive reasoning. If an average infected individual causes R_0 secondary cases, then if we can vaccinate (and protect) at least $R_0 - 1$ of these then an epidemic cannot grow. This argument leads us to the critical level of vaccination needed to protect the population

$$V_C = \frac{R_0 - 1}{R_0} = 1 - \frac{1}{R_0}.$$

This same value is derived from the differential equation model and is the threshold that both prevents the spread of a newly invaded infectious disease *and* the threshold that eradicates an existing infectious disease (Figure 1). What is vitally important to recognise is that not all of the population needs to be successfully immunised to control a disease. Each vaccinated individual also confers some protection to the general population, since those they would have infected are now less likely to catch the disease. This is known as *herd immunity*.

*Warwick Mathematics Institute, University of Warwick, Coventry, CV4 7AL.

2 Below the critical vaccination coverage

For many infectious diseases, it is often difficult or impractical to reach the critical level of vaccination needed to eradicate infection. In such situations, it is pertinent to ask, what is the benefit of limited vaccination? Clearly, each individual that is vaccinated is given protection against disease, but moreover it offers protection to others in the population. Using the SIR-type models we observe that the effects of vaccination are linear (Figure 2); if we only achieve half of the critical vaccination level (V_C) then this still reduces the level of infection in the population by half. This implies that (in general) it is always better to do some level of vaccination, even if the critical threshold cannot be achieved.

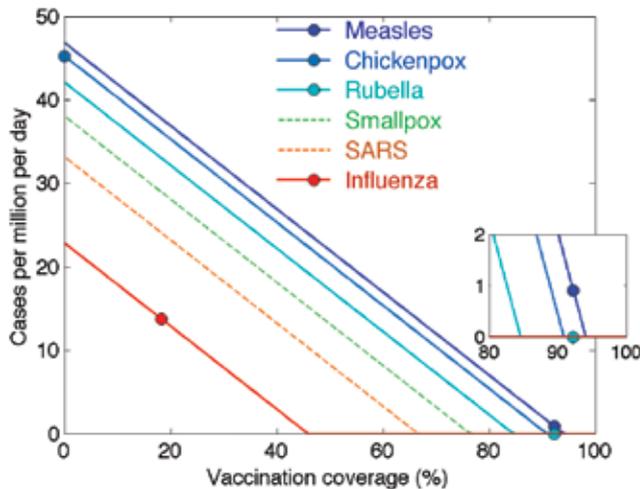


Figure 2: The predicted number of cases of infections declines linearly with vaccination coverage. Eradicated infections (smallpox and SARS) are shown as dashed lines; for all other infections (solid lines) we show the estimated UK vaccination coverage as a dot.¹

3 Beyond the simple models

In practice the simple SIR model given in Box 1 is too naïve to capture the full complexity of human populations and human physiology – although it still provides a tool for qualitative understanding. It is well known that human mixing is assortative, such that individuals predominately mix with, and spread infection to, others of a similar age. In addition, for many diseases the consequences of infection are often more severe in the young and elderly.

Before any new vaccination programme is initiated, detailed mathematical models are required to assess the likely costs and benefits. While such models are based on the SIR paradigm, they need to account for a variety of age-related effects if they are to generate a true picture of the likely impact of any vaccination scheme. Uncertainty in many of the model parameters (including the likely uptake of vaccine and the degree of protection offered by the vaccine) complicates the picture, but the basic methodology is now well developed.

What is surprising, but mathematically verifiable, is that even when complex age-structured mixing patterns are taken into account, the simple vaccination threshold derived above still holds. If we randomly immunise individuals in the population, we achieve disease eradication at $V_C = 1 - 1/R_0$. However, the opportunity to target vaccination raises the potential to control infection more effectively.

4 Vaccines are not perfect

One of the crucial applied factors to account for in any realistic modelling exercise is that vaccines are not perfect. Being injected with a vaccine does not always confer immunity; in such cases the critical threshold applies to the proportion of the population that needs to be immunised (given immunity). For example, if the vaccine fails to generate immunity in a fraction p of those vaccinated, then to achieve herd immunity we need to vaccinate a proportion

$$\frac{R_0 - 1}{R_0(1 - p)}.$$

Hence if p is too big it may be impossible to eradicate infection using that vaccine.

The above description assumes that the vaccine either protects or it does not; an alternative is that the vaccine offers partial protection. Partial protection can arise in many ways: the vaccine can reduce someone's susceptibility to infection (lower the risk of becoming infected), it can reduce subsequent transmission if the person becomes infected, or it may speed up recovery. Such vaccines are described as *leaky*. These can be accommodated within the mathematical framework (Box 1) by allowing vaccinated individuals to become infected and move into a separate infected class (say J) with different parameters. Again a critical vaccination threshold can be established. If R_V is the number of secondary cases produced by an infectious individual when the entire population is vaccinated, then the threshold becomes

$$V_C = \frac{R_0 - 1}{R_0 - R_V}.$$

Clearly, we need R_V to be less than one, or else it is not possible to eradicate the infection. However, we also note that a vaccine works equally effectively if it reduces susceptibility or reduces transmission or speeds up recovery; it is only the impact of R_V that is important. Figure 3 shows that when the vaccine fails to generate any protection in a proportion p of those vaccinated, while the remaining $1 - p$ are fully protected it can be impossible to eradicate the infection using random vaccination. However, when a vaccine is leaky and those vaccinated only have partial protection, then eradication is possible as long as $R_V < 1$.

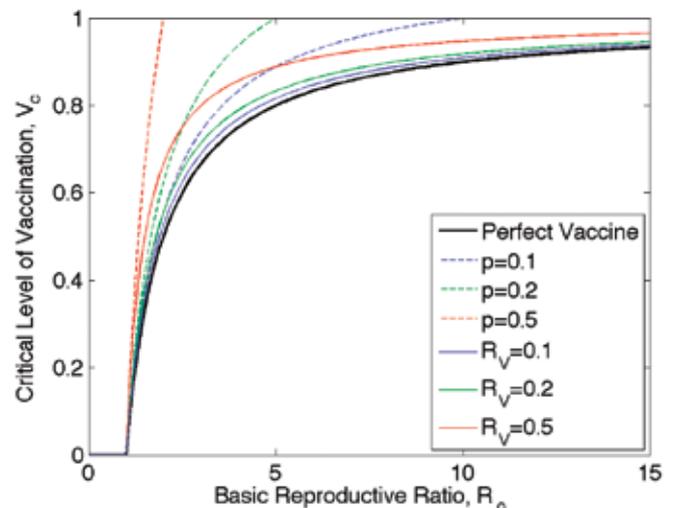


Figure 3: The graph shows how the critical threshold varies as either p (vaccine failure – dashed lines) or R_V (leaky vaccination) varies; in general leaky vaccination has less of an impact than potential vaccine failure.

5 Optimal vaccination

The need to target vaccination, and therefore achieve greater protection than if vaccine was distributed at random, generally arises if vaccines offer short intervals of protection. For the majority of vaccines, protection (possibly after booster vaccines) is considered to be lifelong, so it is always optimal to vaccinate early in life to achieve the greatest duration of protection. However, when the protection offered by the vaccine is short lived, it is generally prudent to target the vaccine at individuals that are most at risk.

An example of this limited duration is the relatively new HPV (human papillomavirus) vaccine that protects against cervical cancer. Given that the vaccine is thought to only offer protection for 5–10 years, it is vital that girls of the appropriate age are targeted; too young and the effects of vaccine are lost before the age of highest risk, too old and some will have contracted HPV before vaccination. Determining the optimal age requires mathematical models that can integrate age-related risks and behaviour with the impact vaccination will have on age-related carriage of infection.

Influenza is another infection where vaccination offers limited protection since the virus is constantly changing and evading the immune response. Again, the question is to determine which ages should be vaccinated to have maximum impact, taking into account that particular groups of the population are likely to suffer health complications following infection and need to be protected – this must be balanced against the costs of the vaccine and the relatively mild effects of influenza on the majority of the population. In 2014, children between 2 and 17 will also be included in the seasonal campaign; this is predicted to have two effects: firstly it will protect children who can often react badly to influenza infection, and secondly it stops spread from children who are a main driver of transmission.

In both these cases, and others, methods from optimisation theory can be used to determine the appropriate targeting of vaccination, with the aim of minimising the costs (in terms of purchasing and administering the vaccine) while maximising the benefits (in terms of lives and treatment costs saved).

6 Optimal vaccination in households

Many people have a strong intuition that they pick up diseases from strangers – for example on public transport or visiting a different department at work. But in fact, it is their closest relatives who often pose the greatest risk and this has important consequences for the deployment of vaccine. The intensity and frequency of interactions between people who share living arrangements causes households to act like local stochastic amplifiers of disease transmission where one initial case can either lead to several more cases within the household or can recover leaving few other household members infected. This leads to a different picture of an epidemic from the homogeneous one based on individuals: the households experience heterogeneous outcomes, and also act as a locus for transmission between generations.

Consideration of a population of households leads to two conclusions for vaccination policy. The first of these is that large households should be preferentially targeted to equalise the number of unvaccinated individuals across different household sizes (Figure 4). This reduction in variability turns out to be a particularly efficient way of allocating vaccine doses if within-household

transmission is significant, and rests on quite general mathematical arguments. While the formal proof strategy is involved, the intuition here is to minimise

$$\sum_{\text{households}} N^2 \times S$$

where N = Number unvaccinated and S = Strength of within-household transmission. While mathematically robust, this strategy is not often implemented for multiple practical reasons.

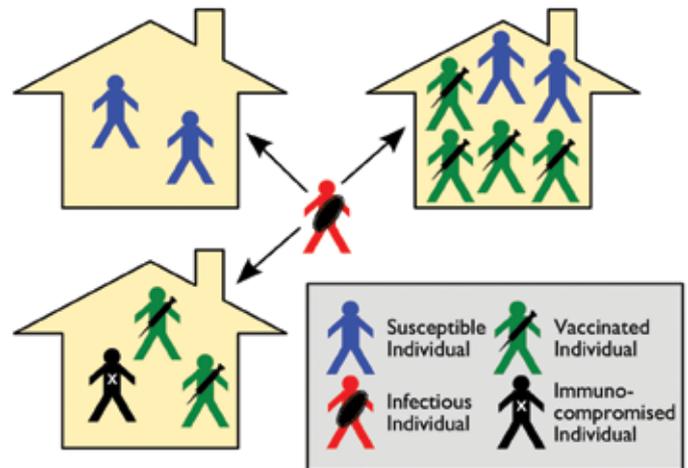


Figure 4: Illustration of optimal vaccination on the basis of household structure.

Another general strategy is to use household structure to protect individuals who have weakened immune systems (called the *immunocompromised*) by generating herd immunity within the household. These individuals are vulnerable to both natural infection and the vaccine, motivating the strategy used during the recent swine flu pandemic of vaccinating all household contacts of immunocompromised individuals (Figure 4).

7 Optimal vaccination in space

Whilst vaccination is used as a preventative measure to reduce the likelihood of disease outbreaks occurring, it also has major benefits if deployed efficiently to reduce disease spread during the course of an epidemic. Such ‘reactive’ vaccination strategies can reduce the need for other intervention strategies that may be less palatable to the general public. In 2001, a major outbreak of foot-and-mouth disease (FMD) caused major devastation to the UK livestock industry and resulted in economic losses totaling around £5 billion. Vaccination was not considered during the outbreak for a number of reasons, not least of which was the fact that the limited vaccination campaign that was logistically feasible at the time would have had little impact upon the outbreak. To control the epidemic the UK government decided to undertake a strategy of reactive culling of ‘at risk’ farms, resulting in the loss of 6.5 million livestock by the end of the epidemic. This policy proved hugely unpopular with the farming community and the general public, and prompted the use of vaccination as part of the UK contingency plan in the event of future FMD outbreaks.

In order for vaccination to be effective during an outbreak, it needs to be targeted in such a way that creates a barrier between the few individuals who are infectious and the susceptible majority. For this reason, ring vaccination around known cases is often considered the best policy to reduce the risk of widespread

transmission. In fact, the World Health Organisation used ring vaccination in the final stages of eradicating smallpox, to contain and eliminate any persisting pockets of infection.

When ring vaccination is deployed, all individuals within a certain distance of detected cases are targeted for vaccination. In the case of FMD, analysis of previous outbreaks indicates that a significant proportion of transmission occurs within only a few kilometres of an infectious farm and this suggests that a policy of ring vaccination has the potential to be highly effective. However, one of the key factors that will determine the success of such a policy is the size of the ring around each reported individual within which vaccination will be carried out. Again, simple ideas concerning the number of secondary cases generated helps to determine a critical ring size (see Figure 5). An effective ring size is one which prevents at least all but one subsequent case; if less than one case escapes from the vaccine ring then the infection will die out, whereas if the vaccination radius is too small the policy will be ineffective as such escapes will occur too frequently and the number of cases will grow.

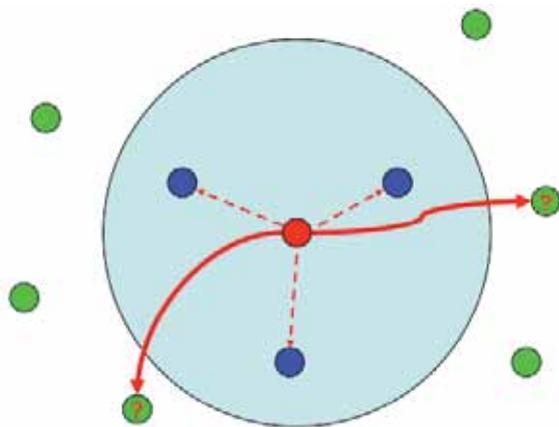


Figure 5: The image shows a caricature of ring vaccination: an infectious individual (red circle) has a ring (pale blue) drawn around it, and those individuals within the ring are vaccinated (blue). The success of this strategy depends on how many of the susceptible individuals (green) outside the ring are infected. If the spread of infection is such that less than one individual outside the ring becomes infected, then the infection cannot spread.

This may lead one to the conclusion that the optimal strategy is to have as large a ring as possible, thus minimising the risk of an escape occurring. Whilst this is a reasonable assumption in theory, in practice there are logistical issues to consider. In the event of large-scale epidemics, it may be impossible to vaccinate all individuals on the day that they are identified. This fact, combined with the time delay from vaccination to immunity and the problem that vaccination may not be 100% effective, means that the size of the ring needs to be carefully optimised in order for it to have maximum benefit. Here, mathematical models and computer simulations play a vital role. If the rings are too large, then by the time all individuals in the ring have been vaccinated, many of them may already have become infected, rendering the strategy ineffective.

The optimal size of the vaccination ring is also dependent upon the spatial spread of the outbreak prior to the introduction of a

vaccination campaign; too many initial cases means that it is impossible to rapidly vaccinate large rings around each. In the case of the 2001 FMD outbreak, had vaccination been considered, resources would have been spread between multiple foci of infection in Cumbria, Devon, the Welsh Borders and Scotland. Whilst ring vaccination could still be effective at combating outbreaks such as this one, the spread of logistical capacity would result in smaller rings being deployed than if the outbreak were contained in one region.

Finally, the introduction of any reactive vaccination strategy must occur rapidly and efficiently – any delay in deployment adds to the risk of the disease spreading to such an extent that ring vaccination may be incapable of controlling the outbreak. In conclusion, reactive ring vaccination has the potential to reduce the impact of future outbreaks of infectious disease. However, there are several considerations that need to be taken into account before such a strategy can be assumed to be successful at reducing spread of disease.

8 The future of vaccination

We have described in this article how vaccination is a key tool in public health (or veterinary health), allowing policymakers to reduce the impact of infection on the population, protect vulnerable groups of individuals and even eradicating certain pathogens globally. Mathematics has played a key role in these processes, allowing policymakers to predict the impact of particular vaccine programmes or to derive more efficient strategies based on mathematical insights. So, do we have all the answers, or does mathematics still have an important role to play in the future? We believe several key challenges remain.

- Current vaccine allocation policy (for seasonal flu as an example) is generally based on age and underlying health problems, could this be better refined? If we account for occupation, household size and location of residence could we improve the efficiency of any vaccine programme?
- We need to have a better understanding of how to optimise the allocation of vaccine during an epidemic; how should policymakers or front-line public-health agencies respond to a changing spatial and demographic pattern of infection?
- Vaccines could potentially have their biggest impact in the developing world. Here, health budgets are limited and diseases often cause high mortality, therefore well-parameterised mathematical models are needed to optimise the allocation of vaccine.
- Finally, models that capture the behaviour of people could be used to understand the uptake of vaccines; what stories in the media or health agencies' press releases have the most impact on the demand for vaccines and how does this impact persist or percolate through the population?

The need for strong mathematics to underpin the use of vaccines is clearly going to continue into the foreseeable future.

Notes

1. Note that in the UK we do not vaccinate against chickenpox; also measles and rubella protection arises from a combined vaccine (the MMR vaccine); currently the uptake of MMR is sufficient to completely control rubella but not measles due to their different R_0 values.