Vaccine 41 (2023) 6630–6636

Contents lists available at ScienceDirect

Vaccine

journal homepage: www.elsevier.com/locate/vaccine

Individual variation in vaccine immune response can produce bimodal distributions of protection

Cameron Zachreson [a,*], Ruarai Tobin [b], Joshua Szanyi [b], Camelia Walker [c], Deborah Cromer [d], Freya M Shearer [f,b], Eamon Conway [e], Gerard Ryan [f,b], Allen Cheng [g], James M McCaw [c,b], Nicholas Geard [a]

[a] School of Computing and Information Systems, The University of Melbourne, Parkville, Victoria, Australia
[b] Centre for Epidemiology and Biostatistics, Melbourne School of Population and Global Health, The University of Melbourne, Parkville, Victoria, Australia
[c] School of Mathematics and Statistics, The University of Melbourne, Parkville, Victoria, Australia
[d] Kirby Institute, University of New South Wales, Sydney, New South Wales, Australia
[e] The Walter and Eliza Hall Institute, Melbourne, Victoria, Australia
[f] Telethon Kids Institute, Nedlands, Western Australia, Australia
[g] Monash Infectious Diseases, Monash Health and School of Clinical Sciences, Monash University, Clayton, Victoria, Australia

ARTICLE INFO

Dataset link: https://github.com/cjzachreson/bimodal_efficacy

Keywords:
Immunogenicity
Correlates of protection
Leaky vaccines
Waning immunity
Vaccine efficacy
Models of vaccine efficacy

ABSTRACT

The ability for vaccines to protect against infectious diseases varies among individuals, but computational models employed to inform policy typically do not account for this variation. Here we examine this issue: we implement a model of vaccine efficacy developed in the context of SARS-CoV-2 in order to evaluate the general implications of modelling correlates of protection on the individual level. Due to high levels of variation in immune response, the distributions of individual-level protection emerging from this model tend to be highly dispersed, and are often bimodal. We describe the specification of the model, provide an intuitive parameterisation, and comment on its general robustness. We show that the model can be viewed as an intermediate between the typical approaches that consider the mode of vaccine action to be either "all-or-nothing" or "leaky". Our view based on this analysis is that individual variation in correlates of protection is an important consideration that may be crucial to designing and implementing models for estimating population-level impacts of vaccination programs.

1. Introduction

Mathematical and computational models of disease transmission are often used to estimate the population-level impact of vaccines and guide the design of immunisation programs. In such models, the effect of a vaccine on protection against infection or other outcomes is typically modelled in one of two ways. The "all-or-nothing" model assumes that a proportion of vaccinated individuals are completely protected from infection, while the remainder are not protected and maintain the susceptibility of an immune-naive individual (see, e.g., [1–3]). The "leaky" model assumes that each vaccinated individual has a reduced risk of becoming infected each time they are exposed (see, e.g., [4–6]).

Estimates of the likely impact of immunisation programs can differ substantially depending on which model of vaccine action is used. In simulations, the magnitude of this difference increases as an outbreak grows, especially for highly contagious pathogens. This difference arises because, in the leaky model, the vaccinated individuals are partially protected and may become infected due to repeated exposures [7,8]. Both the all-or-nothing and the leaky models are abstractions that treat vaccine efficacy as a phenomenological variable that is estimated based on outcome data from phase III vaccine trials and population efficacy studies [9]. Neither model explicitly accounts for the underlying mechanisms of vaccine-induced protection.

While the cohort-level efficacy estimates drawn from phase III trials are useful population-level measures, they do not directly describe the variation in risk over a population. From this perspective, the all-or-nothing model represents a logical extreme by assuming maximal variation over a population, while the leaky model assumes the opposite extreme of minimal variation [10]. Outcome-based models with intermediate levels of variation have been proposed [10]. In these models, risk is estimated based on observed population-stratified vaccine effects. Such approaches highlight the complexity and limitations encountered when constraining models of population heterogeneity using outcome data only.

* Corresponding author.
E-mail address: cameron.zachreson@unimelb.edu.au (C. Zachreson).

https://doi.org/10.1016/j.vaccine.2023.09.025
Received 14 February 2023; Received in revised form 14 September 2023; Accepted 14 September 2023
Available online 2 October 2023

In general, models of vaccine effects based on clinical outcomes do not provide sufficiently flexible methods for accommodating population heterogeneity in vaccine-derived protection. Such heterogeneity can be accounted for by identifying correlates of protection (CoPs, reliable measurements of an individual's vaccine-induced immune response), which have been described as the "holy grail" of vaccine research [11]. Measurements of CoPs from immunogenicity studies of vaccines typically show high levels of variation in the immune response of individuals. For example, the concentrations of neutralising antibodies produced by a vaccine challenge can often vary over orders of magnitude [12–15]. As a result, models of vaccine effects that incorporate the role of CoPs in determining population heterogeneity typically demonstrate high variability in protection against infection or disease outcomes, either explicitly (e.g., through best-fit beta distributions [16,17]) or implicitly through the use of logistic regression methods [11,18–23].

Here, we assess the implications of modelling vaccine efficacy as being mechanistically linked to underlying immune correlates of protection. With SARS-CoV-2 as a well-developed case study, we highlight the hypothesis implicit in such models: that vaccine effects are broadly dispersed over a population due to high levels of variation in CoPs between individuals in that population. The formulation we assess produces a time-varying logit-normal distribution of protection. Through stochastic SIR simulations, we show that this model of protection can act as a flexible intermediate of the leaky and all-or-nothing models, that emerges naturally as a consequence of computing individual-level vaccine-derived protection as a function of an underlying immune correlate. We compare realistic distributions of vaccine-induced or exposure-induced protection to the extremes defined by the all-or-nothing and leaky models.

The model of protection we analyse was developed in the context of the SARS-CoV-2 pandemic after phase II and III trials had been conducted for seven different vaccines including mRNA vaccines, viral vector vaccines, whole virus vaccines and a protein based vaccine. By relating phase III trial outcomes to the levels of neutralising antibodies measured in corresponding phase II trials, these studies demonstrated that neutralising antibodies could serve as a cohort-level correlate of protection against symptomatic infection [24,25]. Notably, declines in neutralising antibody concentrations were found to correspond to waning of levels vaccine efficacy [26].

The model can be calibrated based on empirical evidence relating CoP levels to outcomes, providing nuanced estimates of the impact of vaccines. Furthermore, the model is adaptable to observations of time-varying CoP levels, and can readily incorporate additional layers of heterogeneity such as altered neutralisation levels corresponding to newly emerged viral variants. It therefore provides a flexible approach to modelling vaccine effects that enables the principled investigation of intermediate levels of heterogeneity in vaccine-derived protection.

2. Model description

We employ the model used by Khoury et al. [24] (adapted from CoP studies using logistic regression methods [18–20,22]) in which the level of protection against symptomatic infection varies as a function of a single CoP value (in this case, the concentration of neutralising antibodies, relative to convalescent serum). Note that our focus in this work differs from the focus of Khoury et al. in which they present a model of the average level of protection in a population for a given distribution of neutralising antibody titres in the population. In this work, we interrogate the underlying individual-level model that gives rise to the population model presented in Khoury et al. [24].

This model developed by Khoury et al. assumes that when an immune-naive individual is exposed to antigens, their immune system responds by producing neutralising antibodies which then peak and wane with time. We simplify the dynamic immune response after vaccination by considering only the peak neutralisation level and subsequent decay, while neglecting the relatively short period of time over which the antibody concentration initially increases. The initial peak immune response is modelled as a normally-distributed random variable for each individual i:

$$N_o(i) \sim \mathcal{N}(\mu_n, \sigma_n), \tag{1}$$

in which $N_o(i)$ is the log-transformed neutralising antibody concentration for individual i at the time of peak protection following exposure, μ_n is the population mean, and σ_n is the population standard deviation. This assumption corresponds to the observation that neutralising antibody concentrations are approximately log-normally distributed in studies of immunogenicity following pathogen exposure, vaccination, or combinations thereof (see, for example [12,13,15]).

If an individual is exposed to the same or a similar infectious pathogen following production of neutralising antibodies, their susceptibility to infection will be lower than that of an immunonaive individual. The *efficacy* against infection is defined as:

$$E_{\text{infection}} = 1 - \frac{p_1}{p_o}, \tag{2}$$

where p_o is the probability of infection, given exposure, for an immunonaive individual and p_1 is the probability for an individual with some existing protection derived from vaccination or prior exposure. Recent and past work [18,19,24] has provided support for the use of a logistic mapping between the measured levels of (log-transformed) neutralising antibodies, and the level of protection an individual receives, such that:

$$E_{\text{infection}}(\Delta t, i) = \frac{1}{1 + \exp[-k(N(\Delta t, i) - n_c)]}, \tag{3}$$

where $N(\Delta t, i)$ is the log-transformed neutralising antibody concentration for each individual i at time Δt since vaccination and n_c is the threshold for 50% protection. If we make the simplifying assumption that neutralising antibody concentrations decay exponentially at the same rate λ for each individual, then for a group of individuals who reached peak neutralisation at time t_o, observed at a later time t, then $\Delta t = t - t_o$ and:

$$N(\Delta t, i) \sim \mathcal{N}(\mu_n(\Delta t), \sigma_n), \tag{4}$$

where

$$\mu_n(\Delta t) = \mu_n - \lambda \Delta t, \tag{5}$$

is the time-dependent mean of the log-transformed distribution of neutralising antibody concentrations in the group, after a period Δt since reaching the concentration peak. The above equations produce a time-varying logit-normal distribution of protection over the population:

$$E_{\text{infection}}(\Delta t, i) \sim \text{logit-normal}(\mu^*, \sigma^*), \tag{6}$$

with parameters

$$\sigma^* = k\sigma_n, \tag{7}$$

and

$$\mu^* = k(\mu_n(\Delta t) - n_c). \tag{8}$$

The logit-normal distribution is bimodal when the density of the underlying normal distribution is split through the logistic mapping. Therefore, even if the initial distribution of protection is strongly clustered at high values, the gradual decay of the antibody concentration for each individual will result in periods of time over which the population distribution of protection is bimodal if the CoP distribution is sufficiently broad (Fig. 1). See the Supporting Information for a detailed derivation of the extended parameterisation in Eqs. (7) and (8).

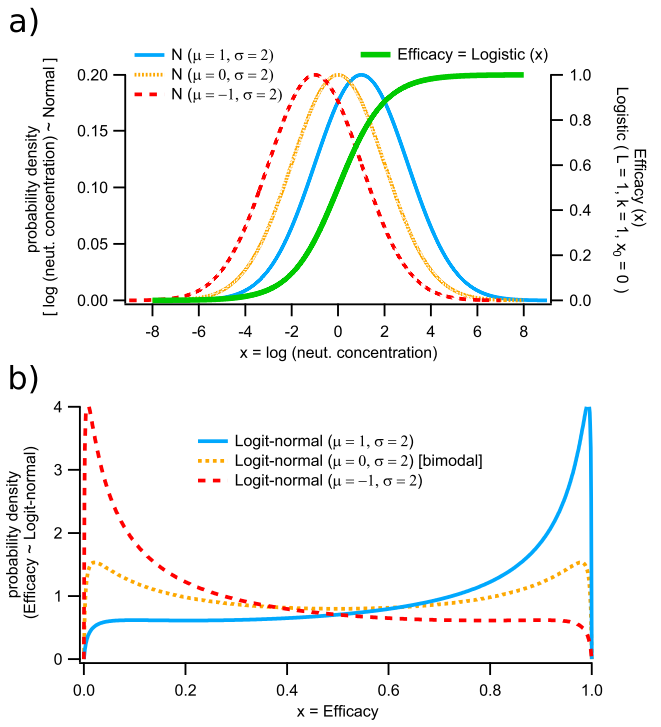

Fig. 1. Examples of how shifting the mean of the underlying normally-distributed control parameter (a) alters the shape of the resulting logit-normal distribution (b) produced through a standard logistic mapping of the control variable. Bimodality occurs when the density of the underlying normal distribution has substantial portions on either side of the logistic mapping function. See the Supporting Information for additional details about the methods used to produce this figure.

3. Results

3.1. Case study: SARS-CoV-2

In a case study, we investigate whether realistic parameterisations of the protection model described above produce bimodal population distributions of vaccine efficacy.

With the model specified by:

- Eq. (1) (distribution of initial immune response),
- Eq. (5) (decay of initial response),
- Eq. (3) (efficacy as a function of antibody concentration) and,
- Eq. (6) (the resulting efficacy distribution),

the key parameters are the population standard deviation of the peak immune response (σ_n), the population mean of the peak immune response (μ_n), the steepness of the logistic function (k), and its inflection point (n_c). The parameters σ_n and μ_n are derived from neutralisation studies, which do not require measurement of infection or disease outcomes. These two parameters can vary as a function of the type of stimulus producing protection (e.g., a vaccine or previous infection), or as functions of the viral variant against which neutralisation is tested. On the other hand, the parameters k and n_c relate to the position and shape of the function relating neutralisation level to relative outcome risk, and their estimation requires clinical trials or population efficacy studies. They may therefore be subject to variation based on population structure and the distribution of underlying risk factors. We posit that the clear conceptual separation between the parameter sets $\{\mu_n, \sigma_n\}$ and $\{k, n_c\}$ is intuitively useful for establishing sources of variation in real-world studies applying this model.

In this case study, we use parameters estimated by Khoury et al. for a convalescent (previously infected) cohort in the context of the ancestral variant of SARS-CoV-2 [24]. We use the parameters corresponding to convalescent serum for two reasons: first, we do not aim to comment on the effectiveness of any particular vaccine (especially now that these efficacy levels are largely irrelevant to the current situation with SARS-CoV-2 due to the emergence of variants). Second, Khoury et al. used convalescent serum as a baseline measure relative to which they computed CoPs for each of the vaccines studied in their work. Convalescent neutralising antibody titres provided an intermediate level of protection, with some vaccines producing stronger responses and others weaker. In this context, the following parameters define the time-varying distribution of immune-derived protection against symptomatic infection:

- $\mu_n = 0$, $\sigma_n \approx 1.07$, $n_c \approx -1.6$, $k \approx 1.3$, and $\lambda \approx 0.0064$,

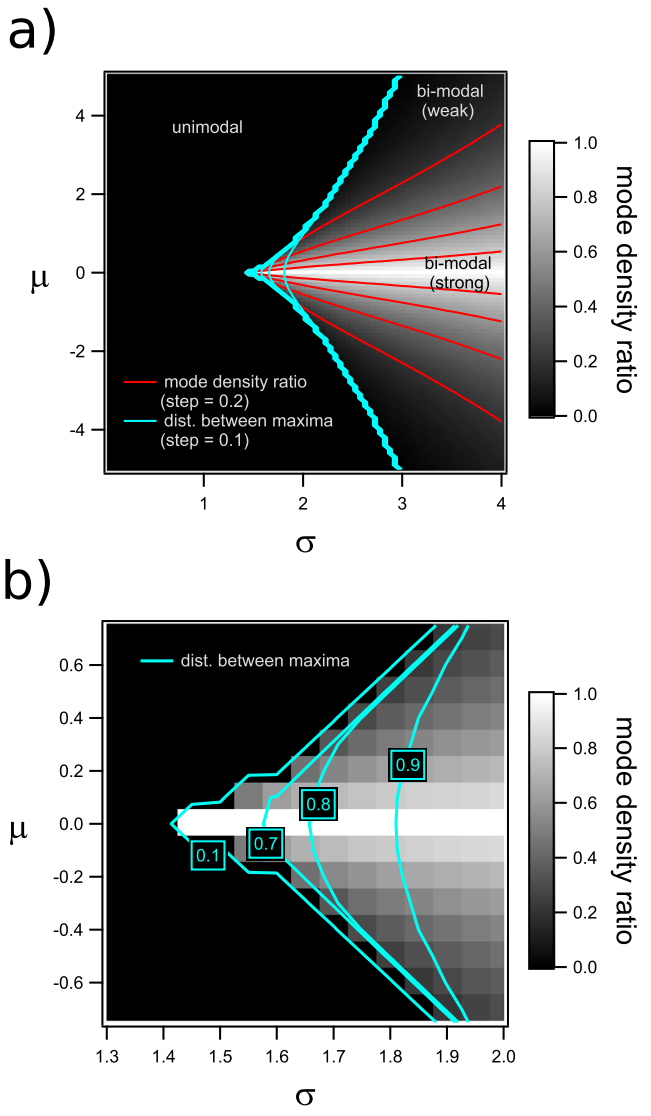

Fig. 2. Numerical characterisation of the standard logit-normal distribution's polarisation as a function of the mean (μ) and standard deviation (σ) of the normally-distributed control variable. Subfigure (a) shows the entire parameter space investigated in this work, while subfigure (b) zooms in to demonstrate the parameter space over which the modes separate. In subfigure (a), shades and red contours correspond to the "mode density ratio" the fraction of the logit-normal probability density lying on either side of the minimum separating the modes (if there are two). The blue contours in (a) and (b) correspond to the distance between maxima (if there is more than one), and demonstrate that the modes rapidly separate as σ increases beyond the threshold for bimodality. See the Supporting Information for additional details about the methods used to produce this figure. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

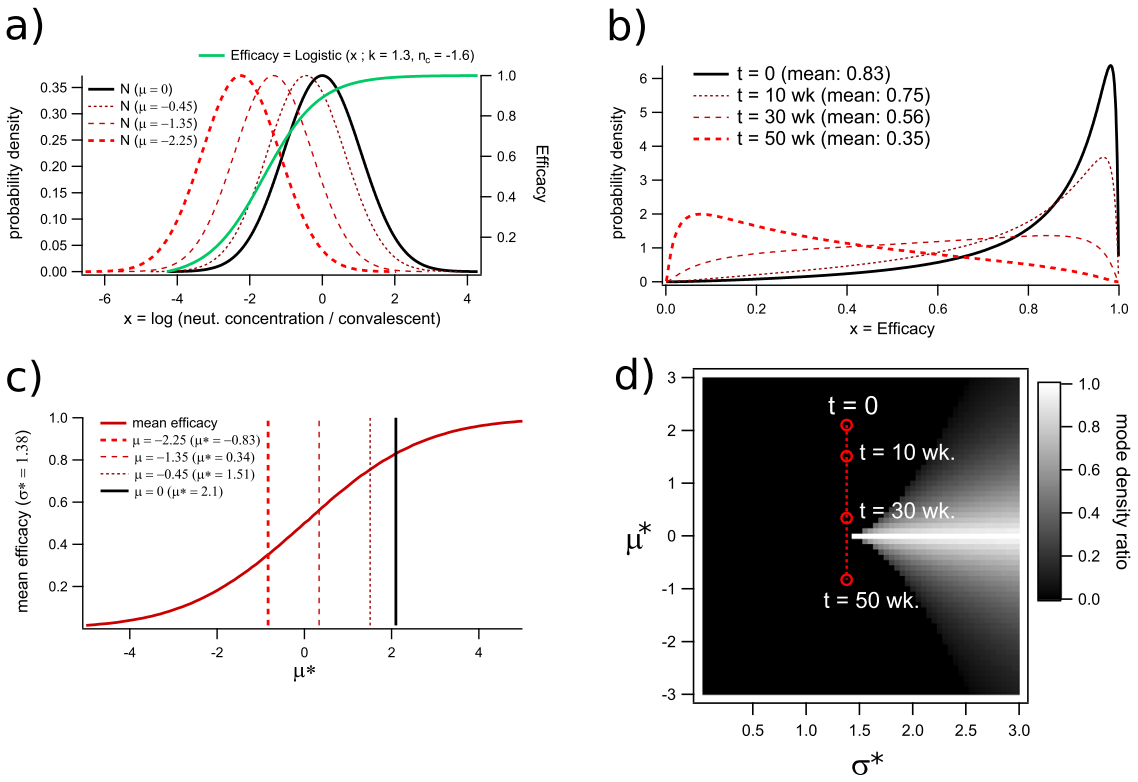

Fig. 3. A case study of COVID-19 infection-derived protection, where the parameters correspond to those computed by Khoury et al. for convalescent serum in the context of the ancestral strains of SARS-CoV-2. The curves in (a) show the underlying distribution of log-transformed neutralising antibody concentrations from peak protection (black curve) and after three periods of waning (dashed curves). The green trace in (a) is the logistic mapping that produces the logit-normal efficacy distributions shown in (b). The mean efficacy as a function of μ^* is shown in (c), with vertical lines indicating the μ^* values for each waning period illustrated in (a) and (b). The distributions in (b) are not bimodal but do demonstrate high levels of dispersion and correspond to a part of the model parameter space that is close to the bimodal phase, indicated by the dashed line and red circles in (d). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

which, at the time of peak protection ($t = t_o$) gives $\mu^* = 2.1$ and $\sigma^* = 1.38$, (according to Eqs. (7) and (8)) corresponding to a high level of protection without bimodality (Fig. 2). See the Supporting Information for further details about parameter selection. The distributions of neutralisation calculated for 0, 10, 30, and 50 weeks after peak neutralisation are illustrated in Fig. 3(a), while the corresponding distributions of protection are shown in 3(b). Because σ^* does not change with time, the population mean efficacy as a function of μ^* is shown in Fig. 3(c). Though the distributions shown in 3(b) are not bimodal, they are widely dispersed, and close to the bimodal part of the model's parameter space (3d). Small increases in either k or σ_n would produce bimodal distributions of protection.

The absence of a bimodal distribution in Fig. 3(b) is indicative of a vaccine that is relatively leaky, suggesting that the results of COVID-19 vaccine trials could have been affected by the presence of repeat exposure, given background infection rates and study duration [27–29]. We anticipate that trials conducted in environments with higher background infection risk or which were conducted for longer duration could conceivably have produced lower estimates of vaccine efficacy than studies conducted for shorter periods, or with lower background prevalence. We note that if accounting for leaky vaccine effects substantially altered estimates of vaccine efficacy, this could potentially alter the best-fit parameter values of the logit-normal model. Such a scenario implies a recursive analysis in which the ascertainment of leakiness (as measured by the inferred value of σ^*) would influence the magnitude of any bias produced by trial exposure conditions. Such a bias would require adjusting the corresponding estimates of efficacy, which could lead to an updated best-fit parameter set with different σ^*.

3.2. Implications for generic epidemic dynamics (SIR)

Here, we depart from our case study of SARS-CoV-2 and use a generic stochastic SIR epidemic simulator to illustrate qualitatively some of the general implications of using a logit-normal distribution of individual vaccine efficacy.

To illustrate some of the implications of using a logit-normal distribution of vaccine efficacy, we depart from our SARS-CoV-2 case study and consider a generic stochastic SIR model. Our goal here is not to examine a particular scenario, but rather to demonstrate how the results of using a logit-normal distribution of protection differ from those produced by more traditional modelling approaches. Note that we use the logit-normal model to describe the population distribution of protection against infection and do not model symptom onset explicitly. We neglect the effects of waning (i.e., we assume that waning occurs on a timescale much slower than that of an outbreak). We also assume implicitly by using the SIR framework that prior exposure produces very high levels of immune-derived protection, which makes re-infection impossible over the duration of the outbreak.

We implemented a discrete-time individual-based disease transmission model setting $R_0 = 2$, and used this transmission simulation to compare three contrasting models for vaccine-derived protection against infection. The leaky model assumes that the force of infection applied to each susceptible individual is reduced by a factor equal to $(1 - \langle E_{\text{infection}} \rangle)$. The all-or-nothing model assumes that each member of the population is protected completely and permanently from infection with probability $p = \langle E_{\text{infection}} \rangle$. Finally, the logit-normal model uses a similar approach to the leaky one, but draws each individual's efficacy value from a logit-normal distribution with parameters (μ, σ), such that $\langle E_{\text{infection}} \rangle$ is the distribution mean (Figure S1). In all scenarios,

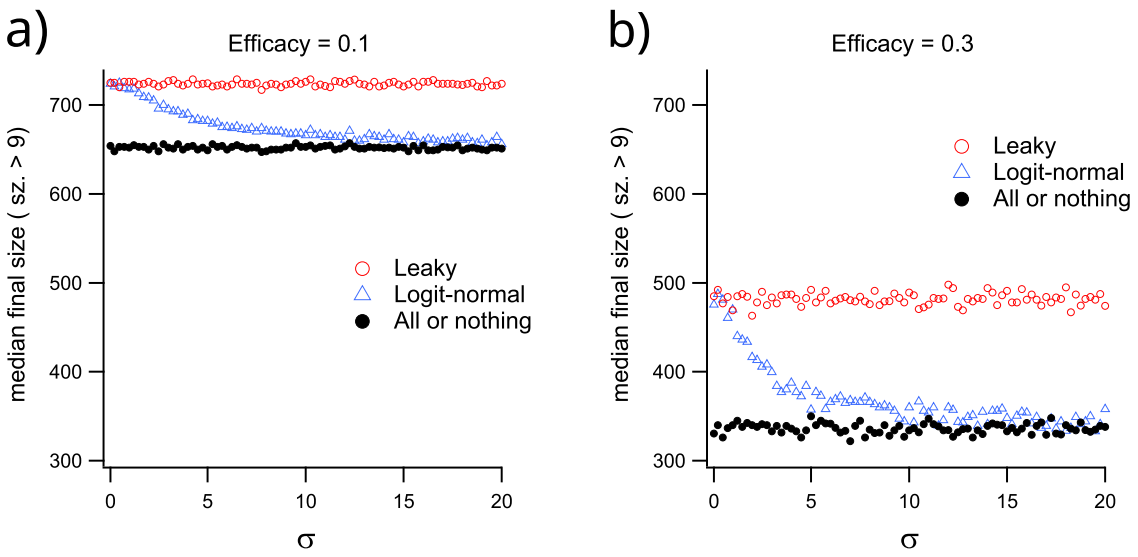

Fig. 4. Comparison of SIR epidemic simulations using three different types of vaccine efficacy (leaky, all-or-nothing, and logit-normal). Plots a and b show the median final size (cumulative number of infections) over 1000 simulations, for efficacy values of 0.1 and 0.3, respectively. Here, we computed the median final size over only those simulation instances for which outbreak size exceeded 9 cumulative cases. Each scenario (point) uses the average efficacy computed from the logit-normal distribution as a function of μ and σ, and μ is varied to hold efficacy constant as σ increases. Values of μ corresponding to each value of efficacy and σ are shown in Figure S2a and the probability of an outbreak with 10 or more cases is shown in S2b. See the Supporting Information for additional details about the methods used to estimate μ for each value of efficacy and σ, and for additional details about how SIR simulations were implemented.

the entire simulated population is vaccinated, and each simulation is initialised with a single infected individual. Comparing the results of repeated stochastic SIR simulations shows that the logit-normal model produces outcomes that lie between those produced by the leaky and the all-or-nothing models. The plots of epidemic final size for different values of mean efficacy in Fig. 4 demonstrate that when σ is near 0, the final epidemic sizes generated by the logit-normal implementation are equivalent to those of the leaky approximation. However, as σ increases, the final size output of the logit-normal model approaches that of the all-or-nothing model, as the density of the distribution of protection accumulates near the extremes.

If neutralising antibody levels are broadly distributed in a vaccinated population (i.e., $\sigma > 1$), this result supports the assertion, made in previous studies, that assuming mean efficacy levels apply uniformly to each member of the population produces overly conservative (pessimistic) predictions of epidemic size [7].

4. Discussion

The logit-normal model described here for immune-derived protection against SARS-CoV-2 was introduced when billions of people globally were being vaccinated in one of the most rapid and extensive vaccination campaigns in human history. While mRNA vaccines appeared promising in their initial capacity to limit transmission of the ancestral variants of the virus, the emergence of variant strains and waning of vaccine-derived protection have led to a endemicity of the SARS-CoV-2 virus, dominated by Omicron variants as of mid-2023 [30–32].

In this context, the role of vaccines in public health efforts to mitigate the disease burden of SARS-CoV-2 has shifted from preventing transmission to mitigation of severe clinical outcomes [33]. However, emerging evidence suggests that waning of vaccine-derived protection against severe disease may occur, though protection against severe disease does appear to be more robust and long-lasting than protection against infection [34]. In order to estimate the future effectiveness of SARS-CoV-2 vaccines in mitigating public health burden, models of transmission must account for the waning of immunity. Models that include quantifiable immune mechanisms will be helpful in this regard because they have the flexibility to connect vaccine effectiveness to immune correlates of protection [35]. Traditional mathematical models of vaccine effects that treat protection as a phenomenological variable guided only by efficacy estimates from phase III trials or observational studies are less well equipped to adapt to changing conditions such as emerging variants and waning immunity. This is because phase III vaccine trials are challenging for endemic pathogens due to highly variable exposure histories [36,37], and present ethical concerns when vaccines have been established to provide benefits. In addition, traditional leaky and all-or-nothing models are inflexible in their implicit treatment of population heterogeneity of immune-derived protection. Using such models will lead either to pessimistic (leaky) or optimistic (all-or-nothing) estimates of the effectiveness of vaccination programs.

While models based on correlates of protection are more flexible and suitable for adaptation to long-term endemic processes, there are many challenges to implementation of these models, both in terms of establishing correlates of protection, and in terms of fitting the individual-level models that map known immune correlates to risk reduction [17,22,30,35]. We showed here that the model adapted by Khoury et al. to SARS-CoV-2 vaccination provides a flexible and intuitive framework for modelling vaccine efficacy as a function of underlying immune correlates of protection on the individual level.

The logit-normal distribution of efficacy is produced by mapping a normally-distributed CoP through a logistic response function. It can be viewed as an intermediate between the leaky and all-or-nothing models of vaccine-induced protection that are commonly implemented. As a heuristic, our results indicate that the all-or-nothing approximation could be appropriate when the distribution of immune correlates is very broad in the population and the effect of those correlates on protection follows a steeply non-linear mapping over a relatively narrow threshold range. If, on the other hand, immune correlates are narrowly distributed over the vaccinated population, or the mapping from CoP to effectiveness has a more gradual slope over a broader range, the leaky model may be a more accurate approximation. In this sense, the standard deviation of the underlying CoP distribution σ corresponds inversely to the protection model's "leakiness" because low values of σ correspond to logit-normal distributions that approach the extreme represented by the leaky model, while high values of σ will cause the model to approach the all-or-nothing extreme.

We have focused here on immune-derived protection against initial infection. Extending this approach to consider protection against severe

clinical outcomes following infection is more challenging. The study by Khoury et al. that correlates neutralising antibody titres with vaccine effectiveness primarily concerns protection against symptomatic SARS-CoV-2 infection. Estimates of the neutralisation level associated with protection against severe infection have been reported, but with the caveat that this implicitly assumes neutralising antibodies play a primary role in the protective immune response against critical or fatal COVID-19 [24,38]. Other immunological mechanisms such as T cell responses may also play an important role in preventing progression to severe disease [39]. As such, population-level effectiveness against severe disease may follow a more complex pattern of temporal decay than the simple exponential model we used in this study [40]. However, it is our conjecture that increasing the complexity of the underlying correlates of protection is likely to introduce additional sources of individual variation. Because bimodality in the distribution of individual-level efficacy is a result of wide variation in the individual correlate of protection, we hypothesise that the general model described here applies qualitatively to downstream clinical outcomes as well, consistent with recent findings [38].

If the same CoPs do apply to multiple stages of a clinical cascade (i.e., infection → severe illness → death) then the dispersion of the resulting efficacy distributions will alter the cohort-level CoP distribution at each subsequent stage due to "selection by natural infection" (see, e.g., Gomes et al. [33]). If we consider infection as a probabilistic sampling process, with weights determined by each individual's immune protection (or other exposure-related risk factors), then we arrive at a scenario in which individuals with lower CoP values are naturally selected at each stage of clinical progression. The larger the dispersion of the distribution of protection applied to each stage of a cascade, the more influence this sampling process will have on the distribution of CoP levels determining protection against subsequent events. In other words, if people with lower levels of neutralising antibodies are more likely to become infected, then the CoP levels in the infected cohort will be lower on average than those in the general population. This natural sampling process will have three general effects on the underlying CoP distribution for subsampled cohorts at different stages of a clinical cascade, relative to the general population. The first is that the CoP mean will decrease. The second is that the CoP variance will decrease. The third is that the CoP distribution will be skewed towards lower values, departing from the normal (Gaussian) distribution function. This third effect means that the quantitative robustness of the logit-normal model presented here could be low with respect to downstream outcomes conditional on infection and any subsequent processes that are influenced by the same CoP that governs relative infection risk. A similar effect would occur if the decay parameter λ were heterogeneous in the population. Because λ is bounded from below (it cannot be negative), its distribution cannot be symmetric around the mean. The more skewed the distribution of decay rates, the more the variance in the CoP distribution would vary as a function of time. This would also skew the CoP distribution and shift the final distribution of efficacy from the logit-normal form. However, we propose that a linear combination of logit-normal models, representing different population strata, could provide a useful and robust approximation. Such an expansion would be able to account for variability in the distribution of CoPs (μ_n), their decay over time (λ), and their relationship to protection (n_c and k) for different subpopulations.

We note that an individual's level of protection could correlate with other characteristics relevant to disease transmission or clinical outcome. We expect that polarised distributions of protection would magnify the effects of any correlations between immune response and behavioural or biological factors that influence such correlations. For example, if those who tend to have weaker immune response (low CoP) also tend to have fewer contacts, as could potentially occur with older cohorts, then the effectiveness of the vaccine on the scale of the whole population would be higher than it would be without such correlations. On the other hand, if those at highest risk of severe disease (given infection) were also more likely to be exposed due to behavioural factors, then the public health burden produced by an epidemic would be higher than expected.

A key limitation of the model as implemented here is that vaccines act solely by reducing susceptibility to infection. Vaccines may also reduce the transmissibility of breakthrough infections. In the leaky model of vaccination, including efficacy against onward transmission reduces the impact of repeat exposures by reducing the risk of transmission associated with breakthrough cases. This additional indirect effect complicates the comparison of the leaky and all-or-nothing models, and can change which of them provides the more pessimistic prediction of overall effectiveness [1].

5. Conclusion

In this work, we detailed a general model of immune-derived protection from infection or disease outcomes. We based our formulation on a model that was adapted and applied by Khoury et al. [24] in the context of vaccine-induced protection against symptomatic COVID-19 disease. The model produces a logit-normal distribution of individual protection, which can be bimodal when the variance of the underlying correlate of protection is large. Using simulations, we demonstrated that at extreme values of the dispersion parameter σ, this model can exhibit the properties of the leaky or all-or-nothing models of vaccine efficacy that are typically used in simulations of infectious diseases. It can therefore be viewed, from a modelling perspective, as a generalisation of these approaches that can be calibrated based on empirical evidence relating the results of immunogenicity studies to the outcomes of phase III trials or vaccine effectiveness studies. We note that even if reliable immunogenicity data is not available, the logit-normal model allows the systematic relaxation of the assumptions implicit in the leaky and all-or-nothing approaches, making it a useful tool for sensitivity analysis. The flexibility of this model enables it to be readily adapted to emerging observations of waning immune correlates of protection, or reduced neutralisation levels observed for new viral variants.

CRediT authorship contribution statement

Cameron Zachreson: Coordinated the research, Performed numerical simulations, Composed the original manuscript text, and the figures, Developed the model, Associated numeric implementations. **Ruarai Tobin:** Developed the model, Associated numeric implementations, Provided the derivation of the extended logit-normal model. **Joshua Szanyi:** Conceptualised the work, Contributed to manuscript composition, framing, and context. **Camelia Walker:** Developed the model, Associated numeric implementations. **Deborah Cromer:** Conceptualise the work, Contributed to manuscript composition, framing, and context, Ensured the COVID-19 case study was an accurate reflection of the model produced by Khoury et al. (2021), Reviewed the formulation of the model. **Freya M Shearer:** Conceptualise the work, Contributed to manuscript composition, framing, and context. **Eamon Conway:** Developed the model, Associated numeric implementations. **Gerard Ryan:** Conceptualise the work, Contributed to manuscript composition, framing, and context, Ensured the COVID-19 case study was an accurate reflection of the model produced by Khoury et al. (2021), Reviewed the formulation of the model. **Allen Cheng:** Conceptualise the work, Contributed to manuscript composition, framing, and context. **James M McCaw:** Conceptualise the work, Contributed to manuscript composition, framing, and context. **Nicholas Geard:** Conceptualise the work, Contributed to manuscript composition, framing, and context.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Cameron Zachreson reports financial support was provided by Australian Government Department of Health and Aged Care. Nicholas Geard reports financial support was provided by Australian Government Department of Health and Aged Care. Allen Cheng reports financial support was provided by Australian Government Department of Health and Aged Care. Joshua Szanyi reports a relationship with Moderna Inc that includes: funding grants. Co-author Allen Cheng is a member of Australian government advisory committees; the views in this paper may not necessarily reflect those of the Australian government.

Data availability

Scripts reproducing our results are available on GitHub: https://github.com/cjzachreson/bimodal_efficacy.

Acknowledgement

We would like to acknowledge useful discussions with David Khoury during the development of this work.

Funding

This work was funded in part by Australian Research Council grant number DP210101920. DC is funded by an NHMRC Investigator Grant, grant number GNT1173528. AC receives funding from the Australian Department of Health and the National Health and Medical Research Council.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.vaccine.2023.09.025.

References

[1] Shim E, Galvani AP. Distinguishing vaccine efficacy and effectiveness. Vaccine 2012;30(47):6700–5.
[2] Abeysuriya R, Hellard M, Scott N. Long-term COVID-19 control requires high vaccination and intermittent control measures. Burnet Institute Public Health Report, 2021.
[3] Patel MD, Rosenstrom E, Ivy JS, Mayorga ME, Keskinocak P, Boyce RM, et al. Association of simulated COVID-19 vaccination and nonpharmaceutical interventions with infections, hospitalizations, and mortality. JAMA Netw Open 2021;4(6). e2110782-2.
[4] Zachreson C, Chang SL, Cliff OM, Prokopenko M. How will mass-vaccination change COVID-19 lockdown requirements in Australia? Lancet Reg Health-West Pac 2021;14:100224.
[5] Kerr CC, Stuart RM, Mistry D, Abeysuriya RG, Rosenfeld K, Hart GR, et al. Covasim: an agent-based model of COVID-19 dynamics and interventions. PLoS Comput Biol 2021;17(7):e1009149.
[6] Germann TC, Kadau K, Longini Jr IM, Macken CA. Mitigation strategies for pandemic influenza in the United States. Proc Natl Acad Sci 2006;103(15):5935–40.
[7] Ragonnet R, Trauer JM, Denholm JT, Geard NL, Hellard M, McBryde ES. Vaccination programs for endemic infections: modelling real versus apparent impacts of vaccine and infection characteristics. Sci Rep 2015;5(1):1–11.
[8] Roberts D, Jamrozik E, Heriot G, Selgelid M, Miller JC. Quantifying the impact of individual and collective compliance with infection control measures for ethical public health policy. 2021, medRxiv.
[9] Andrews N, Tessier E, Stowe J, Gower C, Kirsebom F, Simmons R, et al. Duration of protection against mild and severe disease by Covid-19 vaccines. N Engl J Med 2022;386(4):340–50.
[10] Halloran ME, Haber M, Longini Jr IM. Interpretation and estimation of vaccine efficacy under heterogeneity. Am J Epidemiol 1992;136(3):328–43.
[11] Halloran ME, Longini IM, Struchiner CJ, Longini IM. Design and analysis of vaccine studies, vol. 18. Springer; 2010.
[12] Jackson LA, Anderson EJ, Rouphael NG, Roberts PC, Makhene M, Coler RN, et al. An mRNA vaccine against SARS-CoV-2—preliminary report. N Engl J Med 2020.
[13] Keech C, Albert G, Cho I, Robertson A, Reed P, Neal S, et al. Phase 1–2 trial of a SARS-CoV-2 recombinant spike protein nanoparticle vaccine. N Engl J Med 2020;383(24):2320–32.
[14] Logunov DY, Dolzhikova IV, Shcheblyakov DV, Tukhvatulin AI, Zubkova OV, Dzharullaeva AS, et al. Safety and efficacy of an rAd26 and rAd5 vector-based heterologous prime-boost COVID-19 vaccine: an interim analysis of a randomised controlled phase 3 trial in Russia. Lancet 2021;397(10275):671–81.
[15] Hobson D, Curry R, Beare A, Ward-Gardner A. The role of serum haemagglutination-inhibiting antibody in protection against challenge infection with influenza A2 and B viruses. Epidemiol Infect 1972;70(4):767–77.
[16] Brunet RC, Struchiner CJ, Halloran ME. On the distribution of vaccine protection under heterogeneous response. Math Biosci 1993;116(1):111–25.
[17] Langwig KE, Wargo AR, Jones DR, Viss JR, Rutan BJ, Egan NA, et al. Vaccine effects on heterogeneity in susceptibility and implications for population health management. mbio 2017;8(6):e00796–17.
[18] Dunning AJ. A model for immunological correlates of protection. Stat Med 2006;25(9):1485–97.
[19] Coudeville L, Bailleux F, Riche B, Megas F, Andre P, Ecochard R. Relationship between haemagglutination-inhibiting antibody titres and clinical protection against influenza: development and application of a bayesian random-effects model. BMC Med Res Methodol 2010;10(1):1–11.
[20] Black S, Nicolay U, Vesikari T, Knuf M, Del Giudice G, Della Cioppa G, et al. Hemagglutination inhibition antibody titers as a correlate of protection for inactivated influenza vaccines in children. Pediatr Infect Dis J 2011;30(12):1081–5.
[21] Storsaeter J, Hallander HO, Gustafsson L, Olin P. Levels of anti-pertussis antibodies related to protection after household exposure to bordetella pertussis. Vaccine 1998;16(20):1907–16.
[22] Qin L, Gilbert PB, Corey L, McElrath MJ, Self SG. A framework for assessing immunological correlates of protection in vaccine trials. J Infect Dis 2007;196(9):1304–12.
[23] Okada Y, Kayano T, Anzai A, Zhang T, Nishiura H. Protection against SARS-CoV-2 BA. 4 and BA. 5 subvariants via vaccination and natural infection: A modeling study. Math Biosci Eng 2023;20(2):2530–43.
[24] Khoury DS, Cromer D, Reynaldi A, Schlub TE, Wheatley AK, Juno JA, et al. Neutralizing antibody levels are highly predictive of immune protection from symptomatic SARS-CoV-2 infection. Nat Med 2021;27(7):1205–11.
[25] Cromer D, Steain M, Reynaldi A, Schlub TE, Wheatley AK, Juno JA, et al. Neutralising antibody titres as predictors of protection against SARS-CoV-2 variants and the impact of boosting: a meta-analysis. Lancet Microbe 2022;3(1):e52–61.
[26] Dolgin E, et al. COVID vaccine immunity is waning-how much does that matter. Nature 2021;597(7878):606–7.
[27] Smith P, Rodrigues L, Fine P. Assessment of the protective efficacy of vaccines against common diseases using case-control and cohort studies. Int J Epidemiol 1984;13(1):87–93.
[28] Williams LR, Ferguson NM, Donnelly CA, Grassly NC. Measuring vaccine efficacy against infection and disease in clinical trials: Sources and magnitude of bias in coronavirus disease 2019 (COVID-19) vaccine efficacy estimates. Clin Infect Dis 2022;75(1):e764–73.
[29] Lipsitch M, Goldstein E, Ray GT, Fireman B. Depletion-of-susceptibles bias in influenza vaccine waning studies: how to ensure robust results. Epidemiol Infect 2019;147:e306.
[30] Levin EG, Lustig Y, Cohen C, Fluss R, Indenbaum V, Amit S, et al. Waning immune humoral response to BNT162b2 Covid-19 vaccine over 6 months. N Engl J Med 2021;385(24):e84.
[31] Antia R, Halloran ME. Transition to endemicity: Understanding COVID-19. Immunity 2021;54(10):2172–6.
[32] Koelle K, Martin MA, Antia R, Lopman B, Dean NE. The changing epidemiology of SARS-CoV-2. Science 2022;375(6585):1116–21.
[33] Gomes MGM, Ferreira MU, Corder RM, King JG, Souto-Maior C, Penha-Gonçalves C, et al. Individual variation in susceptibility or exposure to SARS-CoV-2 lowers the herd immunity threshold. J Theoret Biol 2022;540:111063.
[34] Cromer D, Juno JA, Khoury D, Reynaldi A, Wheatley AK, Kent SJ, et al. Prospects for durable immune control of SARS-CoV-2 and prevention of reinfection. Nat Rev Immunol 2021;21(6):395–404.
[35] Lipsitch M, Krammer F, Regev-Yochay G, Lustig Y, Balicer RD. SARS-CoV-2 breakthrough infections in vaccinated individuals: measurement, causes and impact. Nat Rev Immunol 2022;22(1):57–65.
[36] Ray GT, Lewis N, Klein NP, Daley MF, Wang SV, Kulldorff M, et al. Intraseason waning of influenza vaccine effectiveness. Clin Infect Dis 2019;68(10):1623–30.
[37] Lipsitch M. Challenges of vaccine effectiveness and waning studies. Oxford University Press US; 2019.
[38] Cromer Deborah, Steain Megan, Reynaldi Arnold, Schlub Timothy E, Khan Shanchita R, Sasson Sarah C, Kent Stephen J, Khoury David S, Davenport Miles P. Predicting vaccine effectiveness against severe covid-19 over time and against variants: a meta-analysis. Nature Commun 2023;14(1):1633.
[39] Moss P. The T cell immune response against SARS-CoV-2. Nature Immunol 2022;23(2):186–93.
[40] Hogan AB, Doohan P, Wu SL, Mesa DO, Toor J, Watson OJ, et al. Estimating long-term vaccine effectiveness against SARS-CoV-2 variants: a model-based approach. 2023, medRxiv.