Explaining rapid reinfections in multiple-wave influenza outbreaks: Tristan da Cunha 1971 epidemic as a case study
Electronic supplementary material

Anton Camacho[1,*], Sébastien Ballesteros[1,2], Andrea L. Graham[3], Fabrice Carrat[4,5,6], Oliver Ratmann[7,8], Bernard Cazelles[1,9]

1 Laboratoire Eco-Evolution Mathématique, UMR 7625, CNRS-UPMC-ENS-AgroParisTech, 75230 Paris Cedex 05, France
2 Universidade de Lisboa, Centro de Matemática e Aplicações Fundamentais, 1649-003 Lisboa, Portugal
3 Princeton University, Department of Ecology and Evolutionary Biology, Princeton, New Jersey, USA
4 UPMC - Paris 6, UMR-S 707, Paris, F-75012, France
5 Inserm, UMR-S 707, Paris, F-75012, France
6 Assistance Publique Hôpitaux de Paris, Hôpital Saint Antoine, Paris, F-75012, France
7 Duke University, Biology Department, Durham, NC 27708, USA
8 Imperial College London, Dep Infectious Disease Epidemiology, Norfolk Place, London W2 1PG, UK
9 UMMISCO UMI 209 IRD-UPMC, F-93142 Bondy, France

[*] *Corresponding author*:
E-mail: camacho@biologie.ens.fr

Contents

1 Text S1: Model details

1.1 Gamma distributions for the E, I and R states

Stochastic epidemiological models, such as the ones presented in this study, are commonly based on the continuous-time and discrete state-space Markov chains theory. Within this formalism, the duration of each epidemiological state (exposed, infectious, removed) is naturally exponentially distributed. However, the exponential distribution is usually criticized on the grounds of biological realism since it is restrictive in the shape of distribution it can represent. Although this limitation is commonly ignored in analytical studies of epidemiological models, our statistical analysis of a real dataset call for a more realistic alternative. The gamma distribution is a good candidate to improve realism since its flexibility range from the simple exponential distribution to the normal distribution at the expense of only one additional shape parameter k. From a biological point of view, it allows for realistic bell-shaped distribution of biological durations whenever $k > 1$. Within the discrete state-space formalism, the gamma distribution with mean m and shape k is modelled by k consecutive exponential distributions with mean m/k and is called the Erlang distribution due to the integer value of k. Such Erlang distribution has a mean equal to m and a variance equal to m/k.

In all six models, the residence times into the exposed (E), infective (I) and temporary removed (R) states follow Erlang distributions with shapes k_E, k_I and k_R and means $1/\epsilon$,$1/\nu$ and $1/\gamma$ respectively. For simplicity these distributions are not explicitly represented in the main paper and we refer to the Table S1 for a comprehensive interpretation of the Fig 2.

However, unlike other parameters, statistical inference of these shape parameters is combinatorial and thus computationally intractable (see Text 2 for more details). To address this issue, we performed several tests to explore the sensitivity of the maximized log-likelihood as well as the maximum likelihood parameter estimates regarding these shape parameters. We found weak sensitivity of these quantities to the values of k_E and k_I but notable effect of k_R. Indeed, the distribution of the time elapsed in the removed state (R), between the first and second attack, is expected to play a crucial role in the shape and the timing of the second epidemic wave.

Thus, we chose to fix $k_E = k_I = 2$, which is the minimal value to have a realistic bell-shaped distribution and the maximal value to perform analytical calculation (Text 5), and to infer k_R for each model (Table S2). Our choice is even more justified because, unlike the E and I states, the R state has different biological interpretations depending on the model considered. Indeed, while in the temporary removed state, hosts do not participate to the transmission process i.e. they cannot transmit or be reinfected. For the single-strain hypotheses (*AoN*, *PPI* and *Win* models) it mostly corresponds to the protective effect of the specific cellular immune response (see Immunological discussion in the main paper) whereas for the InH model it shapes the intra-host reinfection dynamics. In contrast, for the two-strain hypotheses (*2Vi* and *Mut* models) it should represent the time spent in confinement recovering because specific immunity to one strain is not expected to be fully protective against the other, antigenically differing strain. Although these two-strain models cannot be retained as best explanation of our dataset, we note good agreement between the estimated duration of the R state and the original report by Mantle and Tyrrell: "In mild illness the patient was not confined to the house or to bed for more than a day." [1].

Table S1: Interpretation of the Erlang distribution (with shape k and mean $1/\lambda$) as represented in Figure 2 of the main paper.

Representation (main paper)	Interpretation
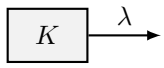	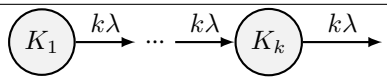
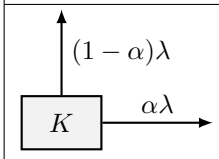	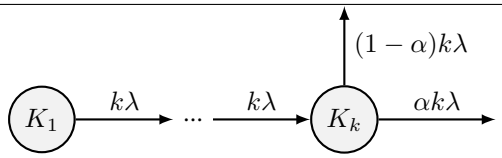

Table S2: Maximum likelihood estimates for the shape k_R and the mean $1/\gamma$ of the Erlang distribution of the duration in the R state

Model	*Win*	*AoN*	*2Vi*	*Mut*	*InH*	*PPI*
k_R	5	4	2	6	9	9
$1/\gamma$ (days)	13.61	11.57	0.64	1.07	5.72	15.90

1.2 Identifiability of the *2Vi* and *Mut* models

Both *2Vi* and *Mut* hypotheses share the same 2-strain history-based model. The general formulation of this model allows different epidemiological parameter values for each strain: β^i, ϵ^i, ν^i, γ^i as well as a cross-immunity parameter σ between the two strains. However, since the original paper by [1] reported that people infected by one of these two viruses formed "a fairly homogeneous group, epidemiologically speaking" and in order to reduce the parameter space of our models, we forced both strains to have the same latent ($\epsilon^i = \epsilon^j$), infectious ($\nu^i = \nu^j$) and temporary resistance ($\gamma^i = \gamma^j$) periods. On another hand, there is a structural non-identifiability issue between β^1, β^2 and σ as these parameters multiply each other [2]. Accordingly, fixing one of these parameters has no effect on the value of the maximized likelihood whereas it decreases the AIC of the model [3]. For the *Mut* model, we chose to fix $\beta^1 = \beta^2$ (and thus $R_0^1 = R_0^2$ since $\nu^1 = \nu^2$ and $R_0^i = \beta^i/\nu^i$) and to estimate the cross-immunity parameter σ. This choice is justified since i) the mutated strain is expected to have a R_0 similar to that of the initial strain and ii) we are mainly interested in the antigenic change conferred by the mutation event and measured by the percentage of antigenic escape σ. In contrast, the *2Vi* model assumes that two different viral agents were introduced at the beginning of the epidemic. We can thus precisely estimate the transmissibility of each virus since there is no biological reason to consider some degree of cross-immunity between the two viruses ($\sigma = 1$).

1.3 Alternative modelling of the cross-immunity in the *Mut* model

A commonly used framework to model the dynamics of co-circulating influenza-strains conferring partial cross-immunity is the history-based formalism introduced by Andreasen *et al.* [4]. Within this formalism, hosts are classified according to their infection history and cross-immunity acts by reducing the susceptibility (RS) to further strains. In our main paper, the mutation hypothesis was modelled using a two-strain history-based formalism and hosts already infected by strain i

have their susceptibility to strain j reduced by a factor $\sigma \in [0, 1]$ (the higher is σ the lower is the cross-immunity i.e. the more distant are the two strains, see Fig. S1A). However, an alternative approach to the RS modelling is the polarized immunity (PI) introduced by Gog *et al.* [5]. Within this framework, each host already infected by a strain i is either entirely susceptible to strain j or completely immune to it. In the case of our simple 2-strain hypothesis, the PI can simply be modelled as in Fig. S1B: following recovery from strain i, hosts remain fully susceptible to the other strain j with probability σ (entering the L_i class) otherwise they become completely protected (entering the L_{12} class). Because these two cross-immunity hypotheses have been previously used for influenza (see references [5, 6] for a complete discussion) we performed a maximum likelihood analysis of our dataset with both models (Fig. S1). We found that the RS hypothesis (Fig. S1A, $\mathcal{L}(\theta_{\text{ML}}) = -115.24$) has a higher log-likelihood than the PI hypothesis (Fig. S1B, $\mathcal{L}(\theta_{\text{ML}}) = -117.99$) which has furthermore the worst $\Delta\text{AIC}_c (= 13.91)$ among the competing models. Although we focused, for simplicity, on the RS hypothesis in the main paper we contend that this likelihood discrepancy between the two cross-immunity hypotheses is somewhat surprising and deserve to be discussed. In particular, the maximum likelihood parameter estimates under these two models notably differ for the basic reproduction number R_0 and the level of cross-protection σ which are crucial epidemiological quantities (Table S3). Although neither of these two hypotheses can be retained to explain the dataset, these results come echoing the main paper and stress the sensitivity of parameter inference and model selection regarding the underlying biological assumptions.

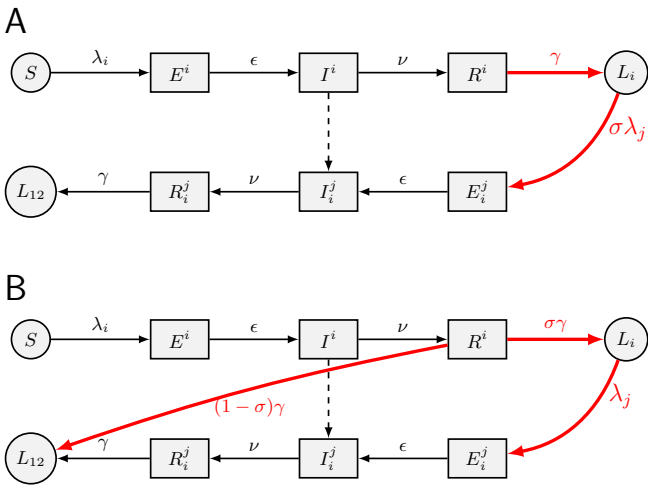

Figure S1: The two *Mut* models. (A) The reduced susceptibility hypothesis of Andreasen *et al.* [4]. (B) The polarized immunity hypothesis of Gog et al. [5]

Table S3: Comparison of the maximum likelihood parameter estimates for the *Mut* model with either polarized immunity (PI) or reduced susceptibility (RS)

Parameter	Description	PI hyp.	RS hyp.
$R_0 = \beta/\nu$	basic reproduction number	20.41	9.16
$1/\epsilon$	mean latent period (days)	2.69	2.38
$1/\nu$	mean infective period (days)	1.94	1.01
$1/\gamma$	mean temporary removed period (days)	16.84	1.07
$1-\sigma$	cross-immunity between the two strains	0.41	0.82
ρ	reporting rate for observation	0.72	0.69
T_{mut}	emergence time of the new variant	8/27	8/24
$\mathcal{L}(\theta_{\text{ML}})$	maximized log-likelihood	-117.99	-115.24
AIC_c	corrected Akaike information criterion	260.56	255.06
ΔAIC_c	rescaled AIC_c	13.91	8.41

1.4 Note on the window-of-reinfection duration

In contrast to the residence times in the E, I and R states, the duration of the reinfection window W in the *Win* hypothesis was modelled using an exponential distribution. We contend that this choice is not only motivated by parsimony but also by profound modelling arguments: *only* hosts with a delayed humoral immune response are expected to experience this window-of-reinfection. The advantage of the exponential distribution relies on its nonzero density for the null duration: this implicitly account for host heterogeneity as *some* hosts with rapid humoral response directly enter the L state. As shown in Fig S2 the Erlang distribution has not such property and as expected, the use of this distribution decreased the likelihood of the *Win* model (data not shown).

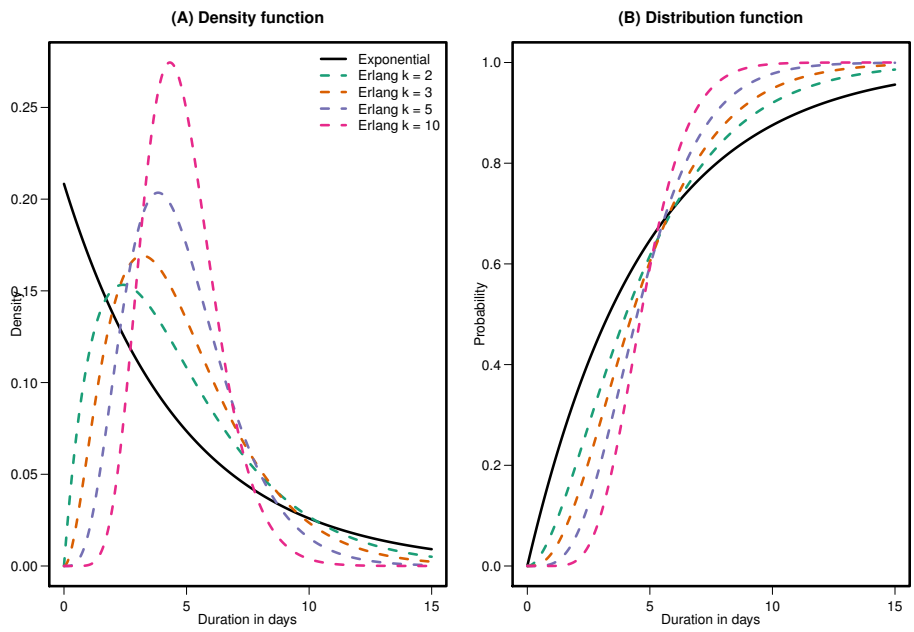

Figure S2: Differences in the density (A) and distribution (B) functions between the exponential (black solid line) and the Erlang distribution (coloured dashed lines). All distributions have the same mean $m = 4.8$ days and we tested different shape parameter $k \in \{2, 3, 5, 10\}$ for the Erlang distribution.

1.5 Transition rates

The Markov Chain events and their transition rates used to stochastically simulate the six models are presented in Tables S4 and Tables S5. The column $Event$ describes the transition, the column $Change$ indicates which compartments are modified and how, the column $Rate$ shows the rate associated with the corresponding transition and the column $Comments$ specifies the distribution for waiting times as well as information on parameters. For simplicity Erlang distributions for waiting times in classes E, I and R are not explicitly represented. In the main paper, Erlang distributions with shape equal k are modelled by k exponential distributions, each exponential rate being equal to k-times the rate reported, as explained in Table S1. All models were simulated using Gillespie's exact algorithm [7] and were implemented in C using the standalone math library from R [8] for random number generation.

Table S4: Stochastic events for the 2 Virus (*2Vi*) and the Mutation (*Mut*) models : $(i, j) \in \{1, 2\}^2$ and $i \neq j$. Upper index corresponds to the infective strain, lower index corresponds to the immunized strains. Coinfection is not allowed.

Event	Change	Rate	Comments
Infection	$(S, E^i) \to (S-1, E^i+1)$	$\beta_i S(I^i + I^i_j)/N$	*Mut*: $\beta_i = \beta$
Infectivity	$(E^i, I^i) \to (E^i - 1, I^i + 1)$ $(E^i_j, I^i_j) \to (E^i_j - 1, I^i_j + 1)$	ϵE^i ϵE^i_j	Erl($k = 2$)
Removing	$(I^i, R^i) \to (I^i - 1, R^i + 1)$ $(I^i_j, R^i_j) \to (I^i_j - 1, R^i_j + 1)$	νI^i νI^i_j	Erl($k = 2$)
Long-term immunity against reinfection by the same virus/variant	$(R^i, L_i) \to (R^i - 1, L_i + 1)$ $(R^i_j, L_{ij}) \to (R^i_j - 1, L_{ij} + 1)$	γR^i γR^i_j	*2Vi*: Erl($k = 2$) *Mut*: Erl($k = 6$)
Reinfection by the other virus/variant	$(L_i, E^j_i) \to (L_i - 1, E^j_i + 1)$	$\sigma \beta_i L_i (I^j + I^j_i)/N$	*2Vi*: $\sigma = 1$ *Mut*: $\beta_i = \beta$
Emergence of the new variant	$(I^1, I^2_1 = 0) \to (I^1 - 1, I^2_1 = 1)$	Punctual change at time $t = T_{mut}$, conditioned by: $I^1 > 0$	Only for *Mut*

Table S5: Stochastic events for the All or Nothing (*AoN*), Partially Protective Immunity (*PPI*), Intra-Host reinfection (*InH*) and Window of reinfection (*Win*) models

Event	Change	Rate	Comments
Infection	$(S, E) \to (S-1, E+1)$	$\beta SI/N$	–
Infectivity	$(E, I) \to (E-1, I+1)$	ϵE	Erl($k=2$)
Removing	$(I, R) \to (I-1, R+1)$	νI	Erl($k=2$)
Return to the susceptible state	$(S, R) \to (S+1, R-1)$	$(1-\alpha)\gamma R$	Only for *AoN* Erl($k=4$)
Reinfection window	$(R, W) \to (R-1, W+1)$	γR	Only for *Win* Erl($k=5$)
Reinfection	*PPI*: $(L, E) \to (L-1, E+1)$ *InH*: $(R, I) \to (R-1, I+1)$ *Win*: $(W, E) \to (W-1, E+1)$	*PPI*: $\sigma\beta LI/N$ *InH*: $(1-\alpha)\gamma R$ *Win*: $\beta WI/N$	*InH*: Erl($k=9$)
Long-term protective immunity	*AoN*, *InH*, *PPI*: $(R, L) \to (R-1, L+1)$ *Win*: $(W, L) \to (W-1, L+1)$	*AoN*, *InH*: $\alpha\gamma R$ *PPI*: γR *Win*: τW	*AoN*: Erl($k=4$) *InH*: Erl($k=9$) *PPI*: Erl($k=9$)

1.6 Choice of the observation process

In the main text, we have assumed that the observation process followed a Poisson distribution. It is straightforward to test whether this assumption is suitable for the epidemic under study. We used the negative binomial distribution (NB) as an alternative to the Poisson distribution. The NB is especially useful when the observations are overdispersed with respect to a Poisson distribution, for which the mean is equal to the variance. Since the NB has one more parameter than the Poisson, the second parameter can be used to adjust the variance independently of the mean. We thus performed log-likelihood profiles on this extra parameter (one for each of the six models) to assess whether this parameter was significantly different from 0 (in this case the NB tends to the Poisson distribution). As shown in Figure S3, we find that the maximum likelihood for the overdispersion parameter is 0 for all six models, which justifies the use of the Poisson distribution.

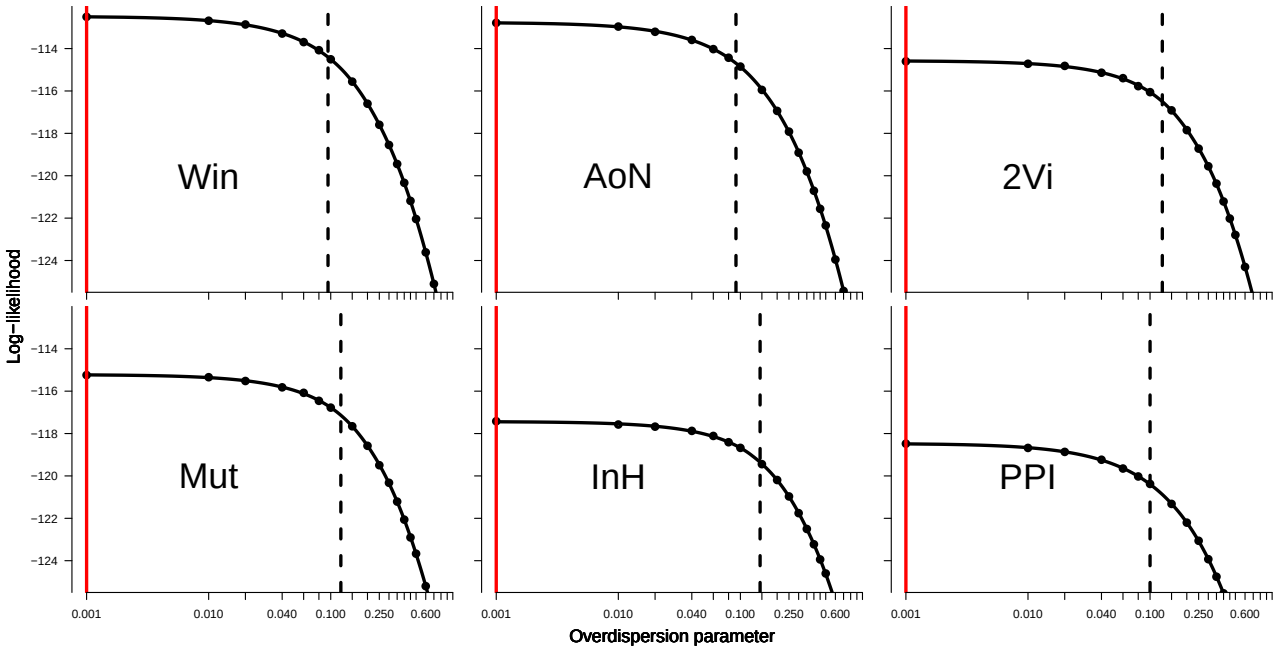

Figure S3: Log-likelihood profiles of the overdispersion parameter assuming a negative binomial distribution (NB) of the observation process. The maximum likelihood estimate for this parameter is indicated by a red line and the 95% confidence intervals by a black dashed line (note the log-scale on the x-axis). For all six models, the maximum likelihood estimate is approached as the overdispersion parameter tends to 0 (note that the NB is not defined when the overdispersion parameter is 0) so that the NB tends to the Poisson distribution.

2 Text S2: Maximum likelihood via Iterated Filtering (MIF)

For a time series $y_{1:T}$ of T successive observations and a model H_i with parameter vector θ, the likelihood is given by the identity $\mathcal{L}(\theta_{ML}|H_i) := P(y_{1:T}|\theta, H_i) = \prod_{t=1}^{T} P(y_t|y_{1:t-1}, \theta, H_i)$, i.e. by the product of the conditional likelihoods at each observation time t. The challenge is to find the maximum of the likelihood as a function of θ.

For our nonlinear stochastic dynamical models (also known as state-space models) the likelihood function is analytically intractable but maximization can be achieved through a numerical method – Maximum likelihood via Iterated Filtering (MIF) – proposed by [9] (see also [10]). In particular, MIF allows to estimate the conditional likelihood at successive observation times and we used such decomposition in Figure 3 of the main paper.

We refer to the references [9, 10] for a complete description of MIF. In the following we only precise our parameterization of the MIF algorithm as well as the different procedures we used to explore the likelihood surface, to assess convergence and to compute confidence intervals.

For all procedures and all models we used the same initial variance multiplier $c^2 = 25$ and the same discount factor $\alpha = 0.95$ [9].

We defined lower and upper bounds for each parameter, based on biological realism, and we used the arctan function in order to transform the bounded parameter space into an unbounded one, as required by the MIF algorithm. The variance for the mutation process of each parameter in the unbounded space is equal to $1/4n$ [9], where n is the number of data points in the time series (n=59). We used a Gaussian perturbation kernel for the mutation process.

During the first 5 iterations of each MIF we used the simple average updating rule proposed by [9] to improve algorithm stability.

For each model we proceeded in 2 steps to reach the maximum likelihood (ML) parameter estimate:

Step 1 Ten crude MIFs with different initial parameter values and wide parameter bounds (see Table S11), 10^4 particles and 150 iterations. For all models, identical parameters have the same lower and upper bounds. This step allows to localize the global maximum of the likelihood surface.

Step 2 A refined MIF with parameter bounds centred on the ML parameter estimates from **Step 1**, 10^5 particles and 200 iterations. This step allow to precisely reach the global maximum of the likelihood surface.

For each step, diagnostic was realized using a convergence plot [9]. Maximum likelihood convergence of **Step 2** was assessed by log-likelihood profiles that allow to separate the global maximum from local maxima.

Since MIF is based on a sequential Monte Carlo (SMC) algorithm, both ML estimates for parameters and maximized log-likelihood ($l(\theta_{ML}|H_i)$) are subject to Monte Carlo variability. In order to handle such inherent variability, the values presented in Figure 4 (see main paper) were obtained as the mean of 10 replicates of the **Step 2** procedure for each model, allowing us to compute the Monte Carlo standard deviation of these estimates (σ^{MC}). For all models, both ML estimates for parameters (see Table S11) and $l(\theta_{ML}|H_i)$ (see Table S6) have σ^{MC} on the order of 10^{-2} or smaller which indicates that the number of particles (10^5) is large enough to neglect the Monte Carlo variability in the model comparison.

Approximate 95% confidence intervals (CI) for all parameters were computed using the procedure detailed in [11]. For a given parameter p and a given model H_i, we computed the log-likelihood profile $\lambda_{H_i}(p)$ by fixing p to several values and maximizing the log-likelihood over all parameters excluding p. The profile was then smoothed using a non-parametric regression. Maximum likelihood convergence was confirmed by comparing the estimate of p in **Step 2** with $\hat{p} = argmax(\lambda_{H_i}(p))$. Lower and upper bounds for an approximate 95% confidence interval are given by $\{p : 2[\lambda_{H_i}(\hat{p}) - \lambda_{H_i}(p)] < \chi^2_{0.95}(1)\}$ where $\chi^2_{0.95}(1)$ is the 0.95 quantile of a χ^2 random variable with one degree of freedom.

For the *Mut* model, the time of emergence of the second strain (parameter T_{mut}) cannot be inferred throughout the time series. We thus performed a log-likelihood profile on T_{mut} and we fixed its value to the argmax of the profile for precise inference of the remaining parameters. Guided by the lowest level of influenza incidence observed between waves in the data, T_{mut} was considered to have occurred sometime between August 21 and September 8.

Similarly, the values for the shape parameters k_E, k_I and k_R cannot be inferred throughout the time series since changing the shape of the Erlang distributions modify the number of intermediate compartments. Indeed, if the shape of a class C changes from k_C^t to k_C^{t+1} between times t and $t+1$ there is no way to divide or rearrange the population from the k_C^t compartments into the new k_C^{t+1} compartments. Accordingly, the only way to infer the shape k is to perform a log-likelihood profile over all its possible values. However, because we must simultaneously infer 3 shape parameters (for the E, I and R classes), the rigorous way to do this is to perform a 3-dimensional log-likelihood profile of complexity $k_E^{max} \times k_I^{max} \times k_R^{max}$, where k_C^{max} is the upper bound of the 3D profile in the direction of the shape of class C. For example, fixing $k_{\cdot}^{max} = 10$ results in 6000 simulations of the MIF algorithm for all six models. After several sensitivity tests (see Text 1.1), we choose to fix $k_E = k_I = 2$ and to perform a likelihood profile over $k_R \in [1, 20]$.

The MIF algorithm was implemented in C using the standalone math library from R [8] for random number generation and probability law distributions (code available upon request). A numerical implementation of MIF is also available via the software package **pomp** [12] which operates in the free, open-source, R computing environment [8].

Table S6: Likelihood of all six models. $l(\theta^i_{ML}|H_i)$ is the maximized log-likelihood, estimated by sequential Monte Carlo and σ^{MC} is the Monte Carlo standard deviation of this estimate. k is the number of inferred parameters plus initial conditions.

Model	$k = \|\theta\|$	$l(\theta_{ML}\|H_i)$	σ^{MC}
Win	9	-112.49	0.03
AoN	9	-112.78	0.02
2Vi	10	-114.61	0.03
Mut	10	-115.24	0.04
InH	9	-117.42	0.04
PPI	9	-118.44	0.03

3 Text S3: Detailed analysis of the model selection based on Akaike Information Criterion

The Akaike Information Criterion (AIC) is a measure of the amount of information about the observed data that is contained in a single, fitted mathematical model [13]. It is defined by

$$\mathrm{AIC}^i = -2l\big(\theta_{\mathrm{ML}} \,\big|\, H_i\big) + 2k,$$

where θ_{ML} is the maximum likelihood point estimate of the unknown model parameters θ of model H_i, and k describes the dimensionality of H_i. Intuitively, the above equation conveys that the AIC penalizes against model complexity as measured by k. More complex models H_j attain a higher value AIC^j only if the difference in the numbers k_i, k_j is compensated by a better log-likelihood fit to the data.

AIC is a statistical quantity based on the observed data points; it estimates a well-defined information-theoretic criterion of model fit that implicitly penalizes model complexity. Under certain regularity conditions, AIC is an unbiased estimator of this quantity when k is set to the number of parameters in the model [13]. When the number of data points T is low relative to k ($T/k < 40$), a second-order correction is commonly used,

$$\mathrm{AIC}^i_c = -2l\big(\theta_{\mathrm{ML}} \,\big|\, H_i\big) + 2k + \frac{2k(k+1)}{T-k-1},$$

to reduce/ eliminate the bias of AIC to the information-theoretic criterion it seeks to estimate [14]. Here k is set to the number of model parameters and T is the number of time series points. In many practical situations, the potential bias and the variance of different AIC statistics is not known. Nevertheless, AIC has proven to be a very efficient measure to distinguish between models [14].

To compare different models H_i in a given set of candidates $\mathcal{H} = \{H_1, \ldots, H_N\}$, it is common to consider the differences

$$\Delta_i = \mathrm{AIC}^i_c - \min_j \mathrm{AIC}^j_c$$

or the so-called Akaike weights

$$\omega_i = \frac{\exp(-\Delta_i/2)}{\sum_{h\in\mathcal{H}} \exp(-\Delta_h/2)},$$

from which "evidence ratios"

$$\omega_1/\omega_i$$

between the best model H_1 and each competing model H_i can be computed. Table S7 lists the respective values of each of these quantities for the models considered in this study. These values allow us to rank our models: the Win model is closely followed by the AoN model, and all other models are trailing substantially behind.

The differences Δ_i can be interpreted as an estimate of the information loss experienced when using the fitted model H_i rather than the best model, H_1. Based on this interpretation, the following rules of thumb have been proposed [14]: models H_i with $\Delta_i \leq 2$ have substantial support, those with $4 \leq \Delta_i \leq 7$ have considerably less support, and models with $\Delta_i > 10$ have essentially no support under the data considered. Under the same interpretation, the evidence ratios can be regarded as betting odds, and support the fact that relatively small differences in Δ_i may convey significant differences in model performance [14]. Regarding our study, Table S7 suggests that the selected best model is not convincingly best: the evidence ratio for the Win model versus the AoN model is only 1.337. However, given the data considered in this study, we must conclude that it is extremely unlikely that one of the remaining models provides an adequate explanation of the Tristan da Cunha time series (evidence ratio > 34).

Table S7: Detailed analysis of the model selection based on Akaike Information Criterion (AIC)

Model	AIC_c^i	Δ_i	Akaike weight ω_i	Evidence ratio ω_1/ω_i
Win	246.65	0	0.555	1
AoN	247.23	0.58	0.415	1.337
2Vi	253.80	7.15	0.016	34.688
Mut	255.06	8.41	0.008	69.375
InH	256.51	9.86	0.004	138.750
PPI	258.55	11.90	0.002	277.500

4 Text S4: Detailed evaluation of the predictions of the six epidemiological models

The Akaike Information criterion is very efficient in discriminating between the fitted models on the reported Tristan da Cunha time series (see Table S7). Moreover, the difference between the Win and AoN model does not appear to be substantial.

This begs the question if the differences are due to different model behaviour during the characteristic second wave on Tristan da Cunha, which is the focus of this study. Alternatively, do these differences accumulate across the entire time series? Decomposing the maximized log-likelihood into conditional log-likelihoods (Figure 3A-E, lower panels), we showed that the differences in the log likelihoods (and thus AIC) accumulate during the second epidemic wave and

the extinction period after the second wave.

To investigate if the differences in AIC correspond to visible, practically significant differences in reproducing the second wave, we performed the following predictive checks. We simulated 10^5 time series under each of our six ML-fitted models. Figure 3 in the main text qualitatively compares the mean behaviour of the best model (Win) with each competitor. Since the six models have very different extinction probabilities during the low-prevalence inter-wave period (Figure 3F), it is natural to separate our simulations between those that have gone extinct before the second wave and the surviving. The information on extinction is summarized in Figure 3F while Figure 3 A-E (upper panels) shows the mean as well as the 95% confidence envelope of the dynamics conditioning on the occurrence of a second epidemic wave.

More precisely, for each model, we removed from the 10^5 simulated incidence time series all those with at least one 0 between day 28 and 35. Indeed, this time interval corresponds to the peak of the second epidemic wave. Finally, we used the remaining time series, which form an empirical distribution of the observed incidence, to compute the mean and the 2.5–97.5 percentile of the daily incidence. Although these 95% confidence envelopes are wide and contain all the data points (for our two best models Win and AoN), qualitative differences between the six models remain apparent during the second epidemic wave. Here, we supplement a suite of quantitative analyses, based on the same predictive simulations, to compare the ability of our six models in reproducing the observed incidence time series.

First, we investigated the independency of the mean standardized predicted residuals

$$\hat{\epsilon_t} = [y_t - \sum_{i=1}^{N} x_{t,i}/N]/[\sum_{i=1}^{N} x_{t,i}^2/N - (\sum_{i=1}^{N} x_{t,i}/N)^2]^{1/2}$$

where y_t is the observed incidence (data) at time t, $x_{t,i}$ is the prediction at time t for realization i and N is the number of simulated time series with a second epidemic wave. Figure S4 shows that the residuals $\hat{\epsilon}_t$ of all models, except *Win* and *AoN*, present significative autocorrelation (p-value < 5% for rank autocorrelation test via the function AutocorTest from the R [8] package FinTS).

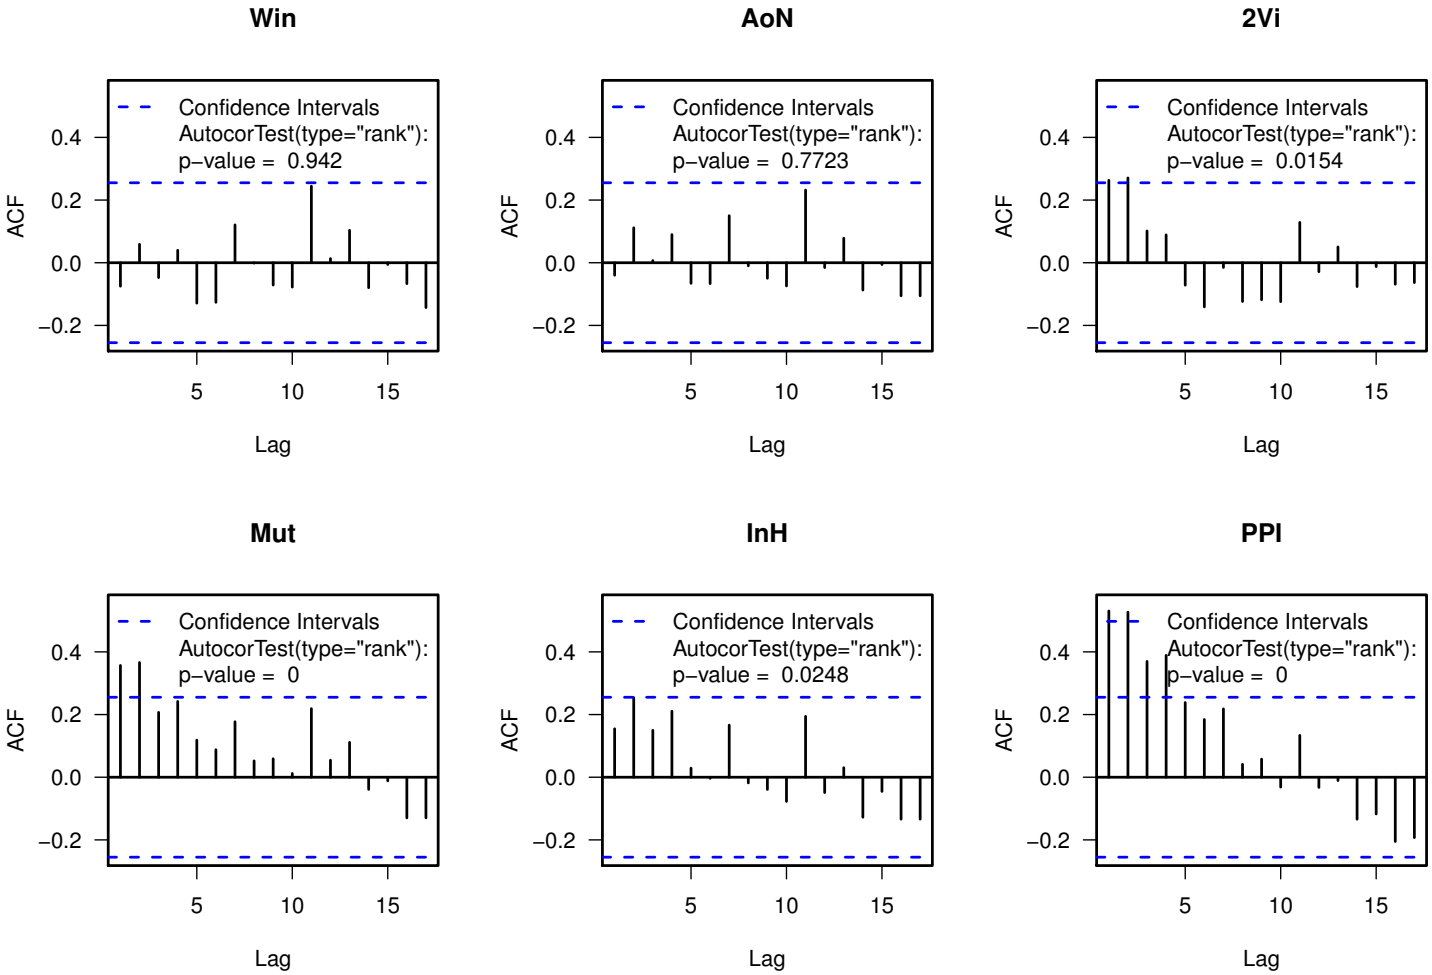

Figure S4: Plot of the autocorrelation function of the standardized residuals $\hat{\epsilon}_t$ for the prediction of each ML-fitted model (function Acf from the R [8] package FinTS).

Second, we quantified the differences between the ML-fitted model in reproducing specifically the characteristic second wave infection dynamics on Tristan da Cunha. A standard approach is to perform a goodness-of-fit test that checks model behaviour with respect to a test statistic T. Here, we will consider two test statistics jointly, T_w corresponding to the second wave and T_e corresponding to the extinction period after the second wave. T_w and T_e are simply defined as particular parts of the full time series $y_{1:T}$,

$$T_w(y_{1:T}) = y_{t_{21}:t_{36}}, \quad T_e(y_{1:T}) = y_{t_{52}:t_{59}}.$$

Furthermore, to enable a comparison of the predictive checks across models, we will synchronize the goodness-of-fit tests across all H_i, $i = 1, \ldots, N$.

Goodness-of-fit tests typically calculate a tail area probability under the hypothesized model H_i to quantify the extremeness of the observed value $T(y_{1:T})$. The common approach is to "repeat" the data observation process in a "thought experiment" under H_i and the fixed value of θ that produced $y_{1:T}$ to derive the reference distribution of T, and to locate $T(y_{1:T})$ on that reference distribution. In our case, it is not possible derive the reference distribution analytically. A simulation-based solution in spirit of frequentist goodness-of-fit tests is to sample replicate data $x_{1:T}$ from the predictive

distribution defined by

$$m\big(dx_{1:T} \,\big|\, H_i, y_{1:T}\big) = \quad \text{Prob}\big(dx_{1:T} \,\big|\theta_{\text{ML}}, H_i\big);$$

see also [15]. Here, we obtained samples from the reference distribution of the two test statistics T_w, T_e with the following simulation scheme:

1. Simulate a time series replicate $x_{1:T}$ under the fitted model H_i, $x_{1:T} \sim \text{Prob}\big(\,\cdot\,|\theta_{\text{ML}}, H_i\big)$.
2. Evaluate and record $T_w(x_{1:T})$, $T_e(x_{1:T})$.
3. Go to 1.

To compare parts of the simulated and observed time series with a single real number, we will use a default distance $d\colon T_w(y_{1:T}) \times T_w(x_{1:T}) \to \mathbb{R}$, and likewise for T_e. Since $d\big[T_w(y_{1:T}), T_w(x_{1:T})\big] = 0$ if and only if $T_w(y_{1:T}) = T_w(x_{1:T})$ and likewise for T_e, we are left to quantify the extremeness of the point $(0,0)$ with respect to our samples from the above reference distribution. Our default distance corresponds to a weighted L_2 norm,

$$d[x_{t_1:t_n}, y_{t_1:t_n}] = \text{sgn}\big(\text{median}(x_{t_i} - y_{t_i})\big) \sum_{i=1}^{n} \omega\, (x_{t_i} - y_{t_i})^2.$$

If the majority of the pointwise distances $x_{t_i} - y_{t_i}$ is positive (negative), d will be positive (negative). For example, d is negative if the simulated second wave lies "below" the observed second wave. The weight $\omega \in [0,1]$ accounts for the number of data points and the frequency f of the pointwise differences $x_{t_i} - y_{t_i}$ and is set to $\omega = 1/nf^2$. Here, f is computed as 1 plus the number of times the pointwise differences switch sign. To illustrate d, suppose we observed a time series $y_{1:7} = (1,3,5,7,5,3,1)$ and two replicate time series $x^1_{1:7} = (3,5,7,9,7,5,3)$, $x^2_{1:7} = (3,5,3,5,7,5,3)$; $x^1_{1:7}$ is consistently above $y_{1:7}$ and $x^2_{1:7}$ oscillates around $y_{1:7}$. While the squared differences are the same, we obtain $d(y_{1:7}, x^1_{1:7}) = 4$ and $d(y_{1:7}, x^2_{1:7}) = 1$.

To diagnose the behaviour of the fitted models in reproducing the second wave and extinction after the second outbreak, we plot the discrepancies $d[T_w(x_{1:T}), T_w(y_{1:T})]$ versus $d[T_e(x_{1:T}), T_e(y_{1:T})]$ based on 10^5 samples $x_{1:T}$ in Figure S5 and tabulate the proportion of points near $(0,0)$ in Table S8. Concentrating on these two aspects of the influenza outbreak on Tristan da Cunha, it is apparent that the *Win* model fits these two characteristics best: the scatterplot is densely centred on $(0,0)$. The scatterplot of the *AoN* model is also centred on $(0,0)$ but less so. The *2Vi* model is as dense as the *AoN* model near $(0,0)$ but shows considerable scatter; 39% of the points are outside the plot. These outliers correspond to predicted epidemics that fade out before the second epidemic wave. Similarly, the scatterplot for the *Mut* and *PPI* model reveals that the dynamics of these fitted models differ significantly from the observed second wave. Regarding the *InH* model, although all the 10^5 points are inside the plot, the density is very low near $(0,0)$ which confirms that this model is unable to closely reproduce the time series.

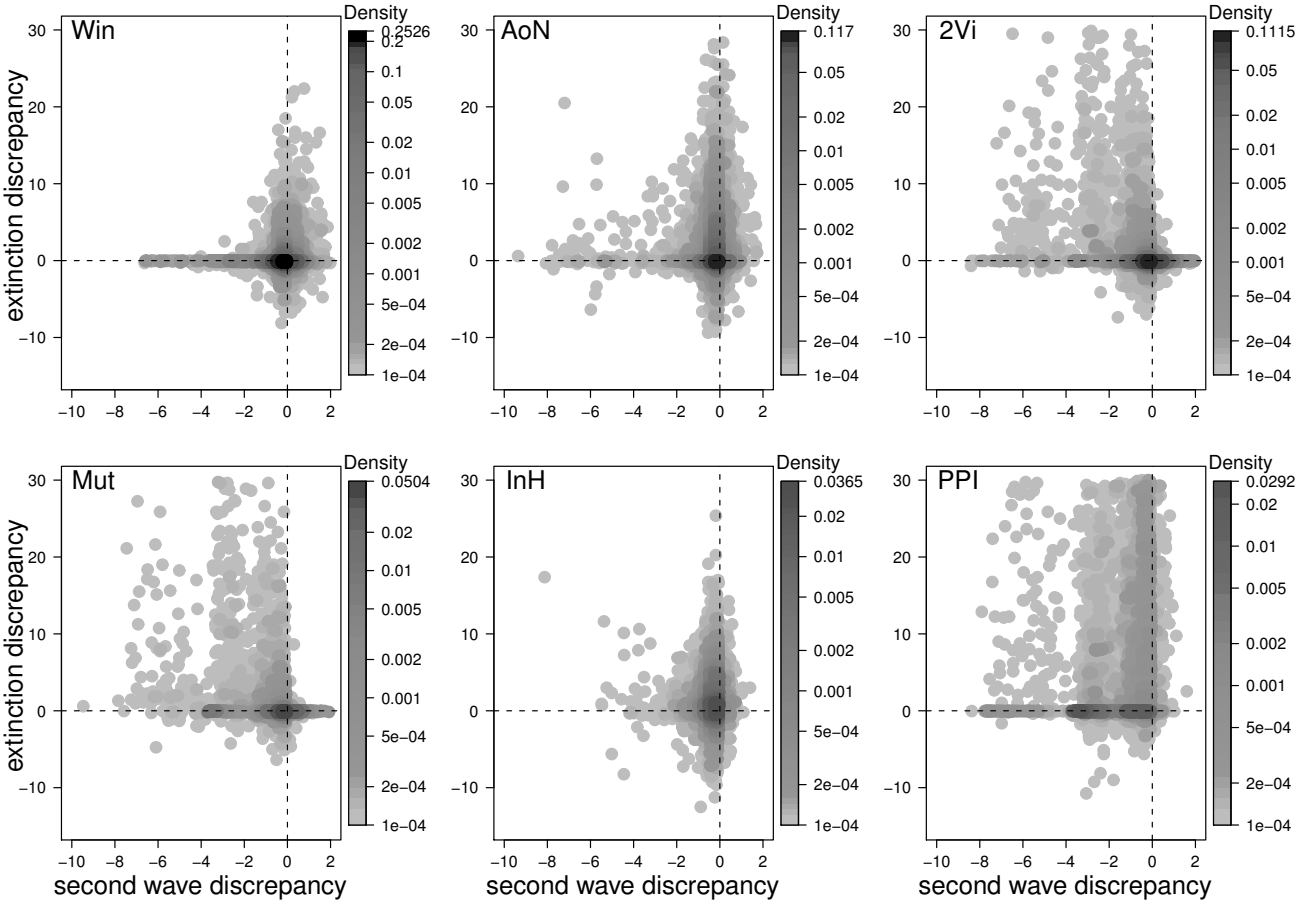

Figure S5: Dotplot of the discrepancies $d[T_w(x_{1:T}), T_w(y_{1:T})]$ and $d[T_e(x_{1:T}), T_e(y_{1:T})]$ for 10^5 replications of $x_{1:T}$ for each of the six candidate models. The range of the discrepancies has been constrained to improve visualization; points outside plot are *2Vi*: 39%, *Mut*: 65%, *PPI*: 40%. In order to compare the density of points across all models we have divided the plane in a grid of 0.25 by 0.25 and we have plotted each point with colour corresponding to the density of points within its 0.25×0.25 square (density are expressed as a fraction of the 10^5 points). The colour legend is represented on the right of each plot in a log-scale. The colour scale is the same for each plot but for ease of comparison between models we have limited each legend scale to the maximum of the corresponding plot (for example the density of points for the *InH* model is always below 3.65% of pts per square)

Table S8: Proportion of the 10^5 points within a radius R from $(0, 0)$.

Model	$R = 0.25$	$R = 0.5$	$R = 1$	$R = 2$
Win	0.26	0.5	0.7	0.81
AoN	0.12	0.26	0.43	0.60
2Vi	0.14	0.31	0.42	0.50
Mut	0.06	0.14	0.20	0.24
InH	0.02	0.14	0.40	0.65
PPI	0.01	0.05	0.11	0.18

The AIC_c differences correspond to diagnosable differences in model behaviour during the second wave and the extinction period. Aspects of this diagnostic plot can also be cast into the form of classical p-values [15]. For example, we can consider the one-sided p-values

$$\text{Prob}\Big(d[T_w(x_{1:T}), T_w(y_{1:T})] > 0 \wedge d[T_e(x_{1:T}), T_e(y_{1:T})] < 0 \;\Big|\; \theta_{\text{ML}}, H_i\Big),$$

which are *Win*: 0.112, *AoN*: 0.05, *2Vi*: 0.12, *Mut*: 0.05, *InH*: 0.01, *PPI*: 0.01. The large one-sided p-value for model *2Vi* can be explained by a substantial number of secondary outbreaks due to the competing virus that are too large in the absence of cross-immunity. Alternatively, we can also consider a two-sided p-value

$$\text{Prob}\Big(d[T_w(x_{1:T}), T_w(y_{1:T})] \times d[T_e(x_{1:T}), T_e(y_{1:T})] \notin \text{HD} \,\Big|\, \theta_{\text{ML}}, H_i\Big).$$

To facilitate model comparison, we will use the same high density interval HD across all models. We choose to calibrate HD such that the two-sided p-value of the *Win* model is 0.2, in which case HD $= [-2.56, 2, 42] \times [-2.7, 2.11]$. Then, the two-sided p-values (and the increase relative to the *Win* p-value) are *AoN*: 0.37 (1.84), *2Vi*: 0.50 (2.48), *Mut*: 0.74 (3.72), *InH*: 0.30 (1.52), *PPI*: 0.79 (3.97). That is, the probability that a time series replicate under the fitted models *AoN*, *2Vi*, *Mut*, *InH*, *PPI* differs from the observed second wave and extinction period by more than the HD bounds increases by at least 1.5-fold relative to the fitted *Win* model. Here, we turned attention away from the historic details of the influenza outbreak on Tristan da Cunha that are captured very well in a likelihood-based model comparison, and turned our focus towards the characteristic features of this outbreak. This analysis supports the view that the *Win* model explains the observed time series best. Moreover it also confirms that the *AoN* model is significantly the second best model. Indeed, among the remaining five models, the *AoN* model is the only one with uncorrelated residuals and it has the best trade-off between i) a high probability of occurrence for the second epidemic wave; and ii) a good accuracy in the reproduction of this second epidemic wave.

5 Text S5: Detailed analysis of the extinction probability

In stochastic models, during the early phase of an epidemic, when the number of infected hosts is small, the chain of disease transmission can be broken and the epidemic goes extinct before a significant proportion of the population has been infected. In our simple models, individuals remain infectious during a period that follows an Erlang distribution with shape $k = 2$. For this class of stochastic models, the probability that a disease fails to invade a population (also called the probability of minor outbreak) has an analytical approximation based on birth/death process theory. Indeed, the early stage of an outbreak can be approximated by a process in which individuals live for a random time (e.g. the infectious period), during which time each individual gives birth (e.g. transmits the disease) according to a Poisson process of constant rate β. Following classical results from the branching process theory [16, 17], one can easily obtain the probability P that the birth/death process goes extinct in a finite number of generations (the epidemic dies out quickly). P has the form q^{I_0} where I_0 is the initial number of infected individuals (see Table S12) and q is the unique root in $[0, 1[$ of the equation $s = f(s)$, with:

$$f(s) = \int_0^\infty e^{-\beta t(1-s)} g_I(t)dt \;\text{ for } |s| \leq 1$$

where g_I is the probability density function of the distribution of the infectious period. For the Erlang distribution with shape k and mean $1/\nu$:

$$g_I(t; k, \nu) = \frac{(k\nu)^k t^{k-1} e^{-k\nu t}}{(k-1)!} \tag{1}$$

One can easily integrate to obtain:

$$f(s) = \left(\frac{k\nu}{\beta(1-s) + k\nu} \right)^k$$

When $k = 2$, the cubic equation $s = f(s)$ has only one root q in $[0, 1[$:

$$q = \frac{1}{2} \left(1 + \frac{4}{R_0} - \sqrt{1 + \frac{8}{R_0}} \right)$$

so that the probability that the epidemic dies out before a major outbreak is given by:

$$P = \left[\frac{1}{2} \left(1 + \frac{4}{R_0} - \sqrt{1 + \frac{8}{R_0}} \right) \right]^{I_0} \quad \text{with } R_0 = \frac{\beta}{\nu} \tag{2}$$

This analytical formula can be compared to the results of Figure 3F obtained by stochastic simulations. For all six models we defined the extinction probability before the first epidemic wave (P^1_{ext}) as the proportion of trajectories (over 10^5 stochastic realizations) that have faded-out by day 10 after the introduction of the virus. Moreover, for the *2Vi* and *Mut* models, since the second virus, i.e. the poorly transmissible (*2Vi*, see Figure S10) or mutated (*Mut*) virus, is responsible for the second epidemic wave, we compared the analytical value of P for these viruses to the increase in the proportion of fade-out trajectories between day 10 and 31 (P^2_{ext}), i.e. the probability of stochastic extinction before the second epidemic wave conditioned by occurrence of the first epidemic wave. Note that for the *Mut* model, since infection by the first strain confers cross-protection against reinfection by the mutant, we have to slightly modify formula (2). We replaced I_0 by $I^2_1(T_{mut})$, the number of host initially infected by the mutated virus (both values equal to 1 in fact), and R_0 by the R_e, the effective reproduction number of the mutant:

$$R_e = R_0 \frac{S(T_{mut}) + \sigma L_1(T_{mut})}{N}$$

where $S(T_{mut})$ and $L_1(T_{mut})$ are the number of susceptible and cross-immunized hosts at the emergence time of the mutant ($T_{mut} = 12$), $\sigma = 0.18$ is the reducing factor of susceptibility for the cross-immunized hosts and $N = 284$ is the population size. We computed the mean values $\hat{S}(T_{mut}) = 10.46$ and $\hat{L}_1(T_{mut}) = 237.69$ over 10^5 stochastic realizations and we obtained $R_e = 1.72$.

Results are presented in Table S9 and show a very good agreement between the analytical formula (2) and the simulations regarding the value of the extinction probability before the first epidemic wave (P^1_{ext}). Despite a single infected host at the beginning of the outbreak, these probabilities are lower than 5% because of the high R_0 values. In contrast, for the *2Vi* and *Mut* models, the analytical approximation of the early extinction probability of the second virus is respectively 5.9 and 1.3 lower than the probability observed in the simulations. We suggest that this discrepancy is mainly attributable to the competition, due to ecological interference, between the 2 viruses during the first epidemic wave (*2Vi* model) or at the emergence time of the new variant (*Mut* model). Interestingly, despite having a high estimated number of initially infectious individuals (Table S12), the low transmissible virus keeps a high extinction probability during the first epidemic wave due to competition with the high transmissible virus.

Table S9: Comparison between the analytical approximation (2) and the empirical values (over 10^5 realizations) of the extinction probabilities before the first and second (*2Vi* and *Mut* models only) waves. The analytical probability * was computed as $P(1^{st}virus) \times P(2^{nd}virus)$: the probability that both viruses go extinct before the first epidemic wave. The analytical probability ** was computed by replacing R_0 by $R_e = 1.72$

	2Vi		*Mut*		*AoN*	*PPI*	*InH*	*Win*
Days	$0 \to 10$	$10 \to 31$	$0 \to 10$	$10 \to 31$	$0 \to 10$	$0 \to 10$	$0 \to 10$	$0 \to 10$
$P_{ext}^{1,2}$ (stochastic realizations)	9.1×10^{-4}	3.6×10^{-1}	3.5×10^{-2}	6.2×10^{-1}	2.5×10^{-2}	5.5×10^{-2}	5.4×10^{-3}	2.9×10^{-2}
P (analytical formula)	9.5×10^{-4} *	6.1×10^{-2}	3.4×10^{-2}	4.7×10^{-1} **	2.4×10^{-2}	5.5×10^{-2}	5.4×10^{-3}	2.7×10^{-2}

6 Text S6: Sensitivity analysis for the *Mut* model

Fig. 3 and Text S5 reveal that in the *Mut* model the newly emerging variant has a high probability to go extinct before initiating the second epidemic wave and that this early extinction probability is partly attributable to the single host initially infected with the new variant. Indeed, because second attacks occurred only three weeks after the beginning of the epidemic in a population of 284 individuals, we reasonably assume that if a new variant emerged during the first wave it had to happen within a single infected host. This assumption can be justified ad hoc since the estimated amount of antigenic escape, that is required to initiate the second epidemic wave, was similar to that of antigenic cluster transitions ($\sigma \approx 20\%$) [18, 19]. Here, we show that even allowing for more than one mutation with large antigenic effects, the *Mut* model remains far behind the two best models in the ability to explain the two-wave observed epidemic.

We performed a sensitivity analysis of the parameter estimates, the extinction probability and the AIC_c value, when the number of newly emerging variants increase. Moreover, in order to test the robustness of our argument we optimize the chance for a second wave to occur by assuming that the same new variant simultaneously emerged within M hosts infected by the mother strain. Results of the maximum likelihood inference are presented in Table S10 and show that parameters are weakly sensitive to M. In particular, the amount of antigenic escape remains high and on the order of 15%. Regarding the maximized log-likelihood, estimated values increase slightly with M thanks to the reduced extinction probability of the emerging variant (Fig S6). However, we note that the probability of inter-wave extinction remains high ($> 25\%$) even for $M = 3$ compared with that of the two best models ($\approx 5\%$). Finally, the still high ΔAIC_c values (> 6) indicate that none of theses *Mut* models can be retained has a likely alternative explanation to the best model.

Table S10: Comparison of the maximum likelihood parameter estimates and ΔAIC_c values for the *Mut* model assuming that M hosts are initially infected with the new variant at $t = T_{mut}$

Parameter	Description	$M = 1$	$M = 2$	$M = 3$
$R_0 = \beta/\nu$	basic reproduction number	9.16	11.19	11.37
$1/\epsilon$	mean latent period (days)	2.38	2.40	2.34
$1/\nu$	mean infective period (days)	1.01	1.35	1.57
$1/\gamma$	mean temporary removed period (days)	1.07	2.20	2.16
$1 - \sigma$	cross-immunity between the two strains	0.82	0.85	0.85
ρ	reporting rate for observation	0.69	0.68	0.68
T_{mut}	emergence time of the new variant	8/24	8/24	8/23
$\mathcal{L}(\theta_{\text{ML}})$	maximized log-likelihood	-115.24	-114.60	114.26
AIC_c	corrected Akaike information criterion	255.06	253.78	253.10
ΔAIC_c	$\Delta\text{AIC}_{c,i} = \text{AIC}_{c,i} - \text{AIC}_{c,Win}$	8.41	7.13	6.45

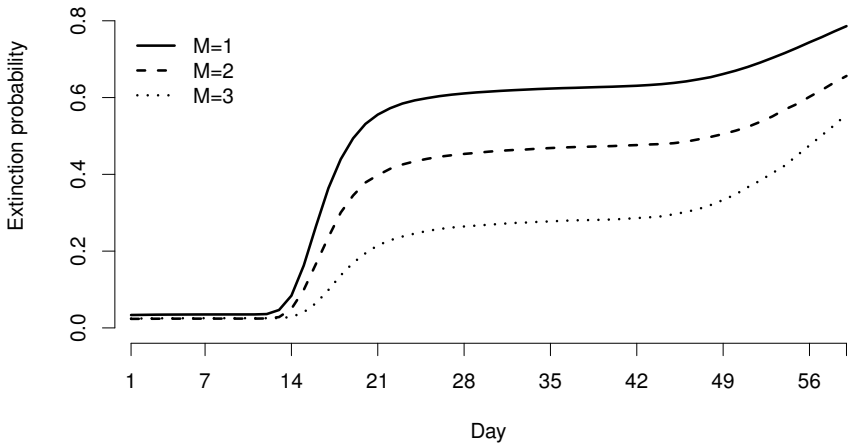

Figure S6: Evolution of the extinction probability of the *Mut* model for different values of M. Extinction probability is defined at each point of time as the proportion of faded-out trajectories $(E(t) = I(t) = 0)$ over 10^4 stochastic realizations

7 Text S7: Dynamics of the *Win* and *AoN* models in large populations

The ecological and epidemiological contexts of the 1971 epidemic on the island of Tristan da Cunha differ in several points from that normally observed in large and uninsulated populations. Indeed, realistic propagation models for these populations incorporate additional mechanisms such as seasonal forcing (i.e. school closure and climate forcing), external reintroduction (i.e. immigration), pre-existing immunity (i.e. past influenza incidents) or age-structure (i.e. heterogeneous mixing). However, in order i) to focus on the direct effects of our reinfection assumptions on the epidemic dynamics in large population and ii) to compare this dynamics with that observed in Tristan da Cunha, we chose to keep the same *AoN* and *Win* models as in the main paper with their maximum likelihood parameter estimates, allowing for only two modifications:

1. we used a population of 1 million individuals with 10 infectious at the beginning of the epidemic (in order to limit early stochastic extinctions), the remaining individuals being susceptible.

2. since in a large population the close-contact rate, and thus R_0, is much lower than among the islanders, we fixed $R_0 = 1.4$ in accordance with estimates for the 2009/A/H1N1 pandemic [20]

Moreover, since the *Win* and *AoN* models presented on the Figure 2 do not distinguish between first and second infections, we used "unfolded" versions of these two models. Obviously, from a mathematical point of view these "unfolded" models are equivalent to the original ones. We ran 1000 stochastic realizations of these two models and we computed the evolution of the mean total incidence, without observation measurement error, as well as the mean incidence due to reinfections (see Figure S7). By comparing the resulting epidemiological dynamics with that of Tristan da Cunha, we note that:

- for large population sizes and more classical R_0 values, the reinfection wave is absorbed in the first and single epidemic wave. We believe that this remains true whenever the time-scale of the temporary removed period ($\sim$week) is small compared to the time-scale of the epidemic duration ($\sim$month). As expected, the best period to observe reinfection cases is just after the epidemic peak because exposure is high and many hosts have already been infected once. Interestingly, this period can be sufficiently distant from the beginning of the epidemic to distinguish between the *AoN* and *Win* mechanisms by comparing the time elapsed between first and second infections: durations greater than 4-6 weeks should advocate for the *AoN* hypothesis. We suggest that experimental studies should focus on this period to validate our results.

- because of a reduced R_0 value, the proportion of hosts having experienced at least one attack in the large population is reduced by half compared to the Tristan population (63.8% vs 95.4% for the *AoN* model and 55.4% vs 95.1% for the *Win* model). However, the proportion of hosts having experienced two attacks is 4 (*AoN*) to 13 (*Win*) times less than in the island. This apparent protection against reinfection, in the large population, can be explained by the lower R_0 value and because most reinfections occurs during the downturn of the epidemic when the force of infection decreases.

- while both *AoN* and *Win* mechanisms lead to similar reinfection rates in Tristan da Cunha (42.6% and 39.4% respectively), simulations in large populations indicate that reinfections occur three times less with the *Win* mechanism (2.4%) than with the *AoN* mechanism (9.1%). This result suggests that reinfection should be rarer and thus more difficult to detect in practice with the *Win* mechanism.

- the proportion of individuals already infected but not immunized at the end of the epidemic is another important epidemiological difference between the *AoN* and *Win* mechanisms. While this proportion is nil for the *Win* mechanism the *AoN* dynamics leads to two times more re-susceptible hosts (21.1%) than re-infected hosts (9.1%). Obviously, these non-immunized hosts form a critical group as they can potentially break the population herd immunity and sustain a second epidemic wave if a fitter virus is reintroduced. However, this potential risk mainly depends on the value of the parameter α (the probability to develop a long-term immunity following infection) which could be exceptionally low among the Tristan da Cunha islanders due to their high susceptibility and/or their genetic background (see the immunological discussion of the main paper).

We stress that our immunological assumptions on reinfection must be empirically validated before investigating their epidemiological consequences on larger populations. Moreover, there is no doubt that the additional epidemiological mechanisms mentioned above, such like seasonal forcing, would be required to make quantitative predictions on a large time scale. However, we suggest that these qualitative preliminary results can nonetheless give some insights into the dynamics of the All-or-Nothing and Window-of-Reinfection mechanisms in larger population/time-scale but also help to identify the best time during the course of an epidemic to conduct experimental verifications.

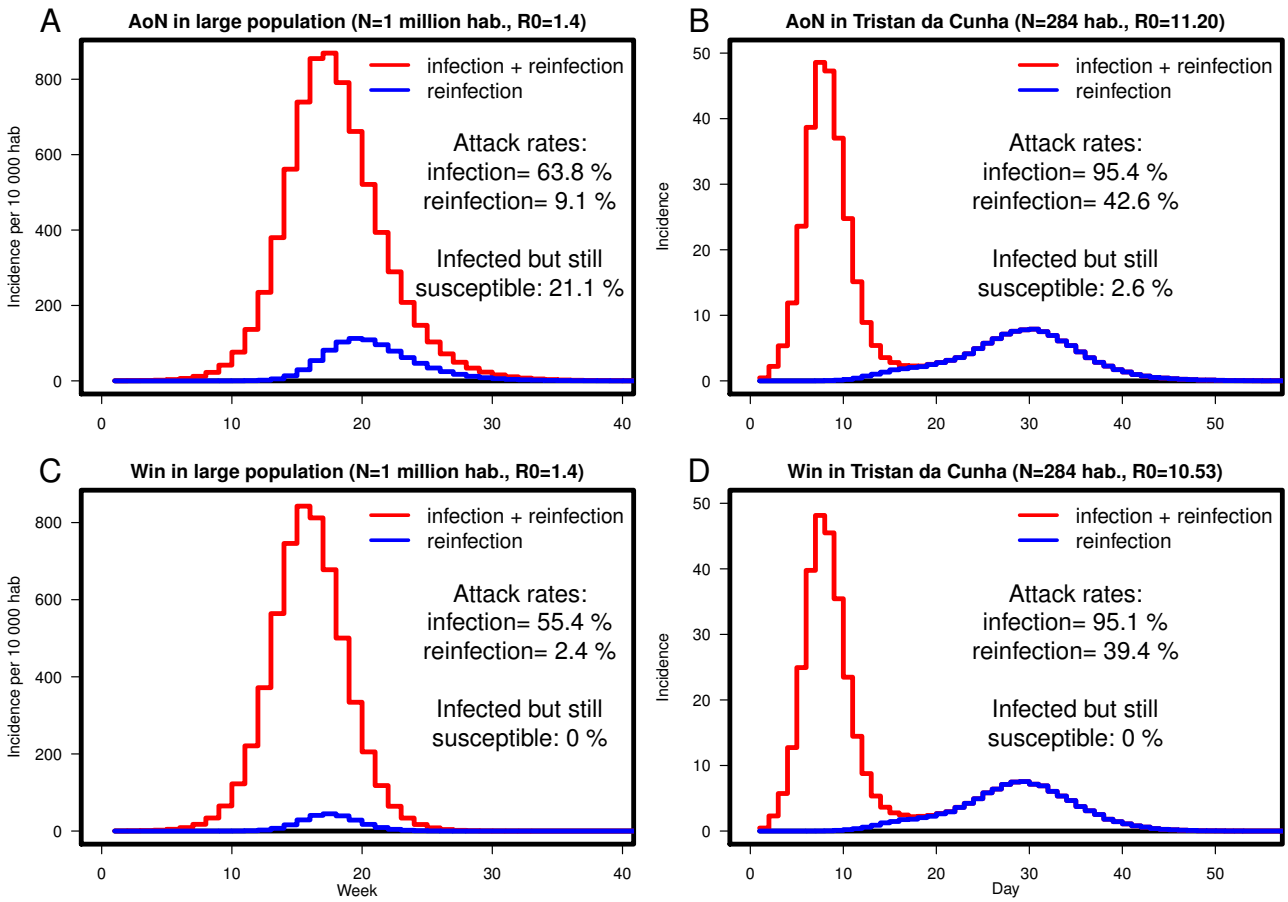

Figure S7: Dynamics of the *AoN* and *Win* model on a large population (A and C, weekly incidence) and comparison with the dynamics in Tristan da Cunha (B and D, daily incidence). We computed the evolution of the mean total incidence (infection plus reinfection, in red) over 10^3 realizations and we plotted the cases due to reinfection (in blue). Attack rates and proportions are expressed in percentage of the total population. A and C: the population represents 10^6 individuals with initially 10 infectious. The R_0 value is fixed to 1.4, and other parameters are those estimated by maximum likelihood (see Figure 4 in the main paper).

8 Text S8: Practical non-identifiability for the *InH* model

The ML estimates of R_0 and the infectious period $(1/\nu)$ are about twice higher under the *InH* model compared to other models. In particular, as shown in Table S11, the corresponding 95% confidence intervals (CI) for these two parameters is widely extended in the *InH* model, resulting in a practical non-identifiability issue [3]. This CI is computed via log-likelihood profiles which appears very flat. To investigate further the origin of the problem we used the results of the likelihood profile of R_0, as explained in reference [3]. Results are presented in Figure S8. As shown in Figure S8A, this profile is flat and the CI is thus widely extended. For each fixed value of R_0, the profile gives the maximum likelihood

estimates of the remaining parameters. As shown in Figure S8B, there is a strong positive correlation between R_0 and the duration of the infectious period ($1/\nu$). By fitting a linear regression over this correlation we found that the transmission coefficient of the virus ($\beta = R_0\nu$) remains approximately constant over the 95% CI and equal to 3.88 effective contacts per days. Accordingly, the evolution of the mean incidence is indistinguishable for all R_0 values included in the 95% CI (Figure S8C). In contrast, the evolution of the mean prevalence presents significant differences after the epidemic peak (Figure S8D) because hosts remain infected longer when $1/\nu$ increases. Together, these results indicate that for the *InH* model i) the observed incidence time series can be reproduced over a large range of R_0 values leading to a practical non-identifiability of this parameter; ii) R_0 is positively correlated with the infective period ($1/\nu$) in order to keep the transmission coefficient constant (ν is thus also poorly identifiable); iii) information on the prevalence after the epidemic peak (around day 15) would be valuable to fix this non-identifiability issue. Unfortunately, there is no data on the evolution of the prevalence in the original paper by Mantle and Tyrrell [1].

On another hand, one can ask why are R_0 and ν less identifiable in the *InH* model than in the other five models? Actually, the *InH* model is the only one where reinfection does not depend on a contact process: following recovery, unprotected hosts are doomed to suffer from an intra-host reinfection. Put it in another way, β (and thus R_0 and ν) is inferred only during the first epidemic wave, when infections are exclusively due to disease transmission, but not during the second epidemic wave, when infections only depend on the shape and the timing of the in-host reinfection via the parameters γ and k_R. In contrast, under the five other models, R_0 and ν are inferred over the whole incidence time series which is enough informative to prevent these non-identifiability issues.

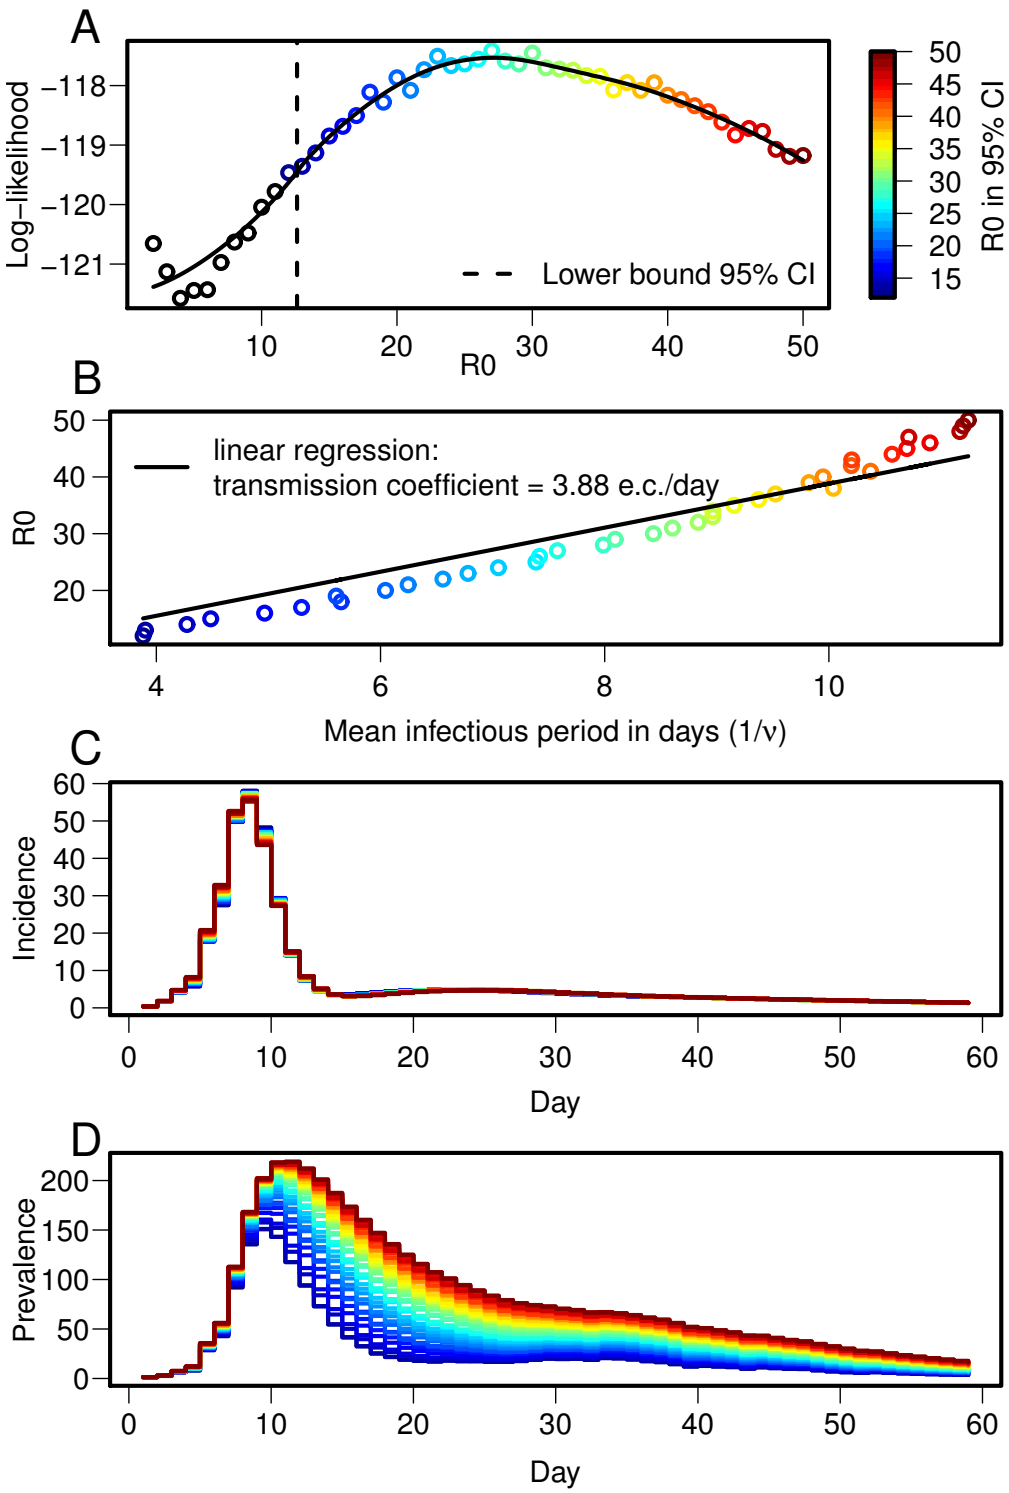

Figure S8: Analysis of the practical non-identifiability of R_0 and ν for the *InH* model. A: profile log-likelihood of R_0. Values in the 95% confidence interval (dashed lines) are colour coded. B: correlation between R_0 and $1/\nu$, for each fixed value of R_0 we plotted the maximum likelihood estimate for $1/\nu$ obtained by log-likelihood profile computation. We then fitted a linear regression and found a constant transmission coefficient ($\beta = R_0\nu \approxeq 3.88$ effective contacts per day) over the 95% confidence interval $R_0 \in [12, 50]$. C: for each value of R_0 we computed the evolution of the mean incidence. Curves are plotted with colour corresponding to the R_0 value and cannot be distinguished. D: as in C but for the evolution of the mean prevalence.

9 Text S9: Calculation and comparison of the effective reproduction ratio

Our maximum likelihood estimates of the basic reproduction number $R_0 = \beta/\nu$ are unusually high for influenza (around 10 under our two best models). We define $R'_e(t)$ as the effective reproduction number predicted by our mechanistic models at time t:

$$R'_e(t) = R_0 \times \hat{S}_t/N \tag{3}$$

where $N = 284$ is the population size and $\hat{S}_t$ is the number of individuals susceptible to be infected or reinfected at time t. We note that the expression of $\hat{S}_t$ depends on the reinfection hypotheses. Using the notations of Figure 2 (main text) we have $\hat{S}_t = S(t)$ for the *AoN* and *InH* models, $\hat{S}_t = S(t) + W(t)$ for the *Win* model, $\hat{S}_t = S(t) + \sigma L(t)$ for the *PPI* model, $\hat{S}^i_t = S(t) + L_j(t)$ for strain i in *2Vi* model and $\hat{S}^i_t = S(t) + \sigma L_j(t)$ for strain i in *Mut* model. Note that our estimates of the initial conditions (Table S12) together with the reported high attack rate (96%) indicate that at the beginning of the epidemic we have $R'_e(0) \approx R_0$.

Interestingly, since the dataset reports daily incidence counts, it is possible to reconstruct the time course of the effective reproduction number (R_e) using the algorithm described in reference [21] and thus to compare and test our results.

The method proposed by Wallinga and Teunis [21] is a likelihood-based estimation procedure that infers "who infected whom" from the observed dates of symptom onset as provided by the epidemic curve. The key assumption that allows the authors to write down a likelihood function requires the independence between transmission events from case j to case i and from case j to any other case k. Assuming this pair approximation, their algorithm takes as inputs the incidence time series and the distribution of the generation time (τ) which is the time from symptom onset in a primary case to symptom onset in a secondary case. Given these two inputs the likelihood that case i as been infected by case j is:

$$p_{ij} = \frac{f(t_i - t_j)}{\sum_{i \neq k} f(t_i - t_k)} \tag{4}$$

where t_i and t_j are the dates of symptom onsets and f is the probability density function of the generation time (note that $f(\tau) = 0$ for $\tau < 0$ so that if p_{ij}=0 if $t_i < t_j$). The effective reproduction number for case j is then:

$$R_j = \sum_i p_{ij} \tag{5}$$

and the average daily reproduction number $R_e(t)$ is calculated as the arithmetic mean over R_j for all of those cases j who reported symptoms of illness on day t. We refer to the original paper of Wallinga and Teunis [21] for a step by step derivation of equations (4) and (5).

In order to compare $R_e(t)$ with $R'_e(t)$ for each model, we must derive the analytical expression for the probability density function of the generation time associated with our mechanistic models. Its general form is equal to the following convolution [22]:

$$\begin{aligned} f(\tau) &= g_E * h(\tau) \\ &= \int_0^\tau g_E(t) h(\tau - t) dt \end{aligned} \tag{6}$$

where

$$h(t) = \int_t^\infty \frac{g_I(s)}{s} ds$$

and g_E (resp. g_I) is the probability density function of the latent (resp. infectious) period. In this study, we have assumed that these periods follow an Erlang distribution with shape equal to 2 and we have estimated the mean latent ($1/\epsilon$) and infectious ($1/\nu$) periods for each model. Using the general formula (1) given in the Text S5 for the probability density function of the Erlang distribution, we obtain $g_E(t) = (2\epsilon)^2 t e^{-2\epsilon t}$ and $h(t) = 2\nu t e^{-2\nu t}$ for $t \in [0, \infty)$. Replacing in equation (6) we obtain for $\tau \in [0, \infty)$ (provided that $\epsilon \neq \nu$):

$$f(\tau) = \frac{4\epsilon^2 \nu}{(\nu - \epsilon)} \Big[t e^{-2\epsilon t} + \frac{1}{2(\nu - \epsilon)} (e^{-2\nu t} - e^{-2\epsilon t}) \Big] \tag{7}$$

For each model, we computed the time course of R_e by replacing ν and ϵ in equation (7) by their maximum likelihood estimates (reported in Table S11). We also computed R'_e using equation (3) and taking the mean value of $\hat{S}_t$ over 10^5 simulations, conditioning on the occurrence of the second epidemic wave. Comparison is presented in Figure S9 from which we can draw the following conclusions:

1. R_e is found to be high (around 8) at the beginning of the epidemic which is consistent with R'_e (around 10) and gives support to a high R_0 value during the 1971 influenza epidemic in TdC. In agreement with the discussion in [21], we note that the observed discrepancy with R'_e is likely to result from the fact that Wallinga and Teunis' algorithm ignores under-reporting in the data counts (around 15%). By contrast, our maximum likelihood approach accounts for under-reporting.

2. The fluctuations in R_e throughout the epidemic (black dots) are fairly well reproduced by our mechanistic models (red and blue curves) and in particular by our two best models (Win and AoN). Small fluctuations of R'_e are not captured graphically (those are averaged values over 10^5 simulated time series) but occur in individual simulations due to stochasticity.

3. Since our mechanistic approach assumes a constant transmission rate, this analysis shows that variations in the number of susceptibles can fairly well reproduce and explain the observed fluctuations in R_e when the possibility of reinfections is explicitly included into epidemiological models. Indeed, we argue that in this small community (N=284) with a high attack rate (96%), variation in R_e is expected to depend crucially on the available number of susceptibles. This is especially true during the second epidemic wave when fluctuations in R_e above 1 strongly depend on the rate at which hosts become available for reinfection. We note that fluctuations in the transmission rate could also have played a role in the variation of R_e. For instance, the initial welcome gatherings might have played a role in accelerating the initial growth rate of the epidemic. Although this question could be addressed by assuming a separate transmission rate at the beginning of the epidemic, it would have no impact on model selection because the differences in log-likelihoods (and thus in AIC) between our six models accumulate during the second epidemic wave rather than at the beginning of the first wave. Indeed, all six models behave similarly during the first wave (see Figure 3 A–E in the main text).

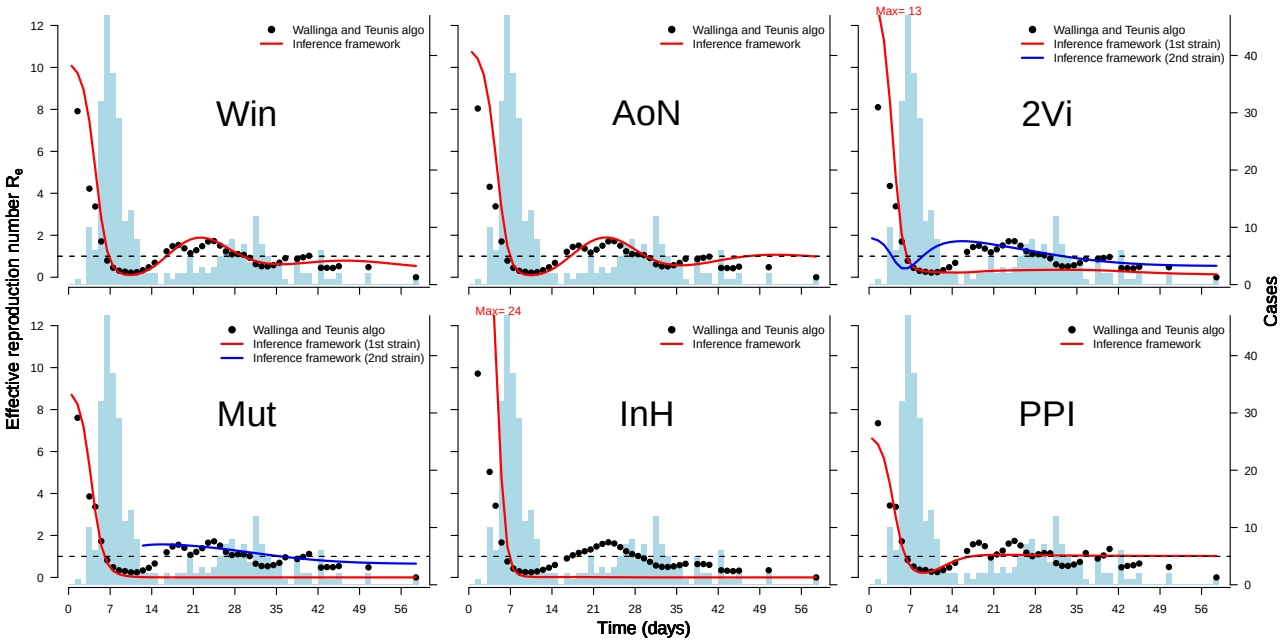

Figure S9: Comparison between R_e and R'_e throughout the epidemic. Note that for the InH model R'_e does not increase during the second epidemic wave because reinfection occurs within hosts.

10 Table S11: 95% confidence intervals for parameter estimates

Table S11: Parameters description, parameters bounds used for the procedure described in **Step 1** of Text S2, maximum likelihood estimate (MLE), Monte Carlo standard deviation of this estimate (σ^{MC}) for parameters and approximate 95% confidence intervals computed via log-likelihood profile. See Text 2 for computational details. Values with * indicate that the end of the confidence interval is beyond the parameter bound.

Parameter	Description	Bounds	MLE	σ^{MC}	95% CI
	Window-of-reinfection (*Win*)				
$R_0 = \beta/\nu$	basic reproduction number	$[1, 50]$	10.53	$< 10^{-2}$	6.14 – 17.06
$1/\epsilon$	mean latent period (days)	$[0.1, 7]$	2.16	$< 10^{-2}$	0.45 – 2.65
$1/\nu$	mean infective period (days)	$[0.1, 11]$	2.01	$< 10^{-2}$	0.1* – 4.20
$1/\gamma$	mean temporary resistance period (days)	$[0.1, 30]$	13.61	0.02	11.04 – 20.78
$1/\tau$	mean duration of the reinfection window (days)	$[0.1, 20]$	4.80	0.01	2.67 – 15.01
ρ	reporting rate for observation	$[0, 2]$	0.71	$< 10^{-2}$	0.59 – 0.79
	All-or-Nothing (*AoN*)				
$R_0 = \beta/\nu$	basic reproduction number	$[1, 50]$	11.20	0.03	5.50 – 19.91
$1/\epsilon$	mean latent period (days)	$[0.1, 7]$	2.07	$< 10^{-2}$	0.65 – 2.78
$1/\nu$	mean infective period (days)	$[0.1, 11]$	2.44	$< 10^{-2}$	0.1* – 4.49
$1/\gamma$	mean temporary resistance period (days)	$[0.1, 30]$	11.57	0.02	7.60 – 17.05
α	probability to develop long-term immunity	$[0, 1]$	0.53	$< 10^{-2}$	0.37 – 0.60
ρ	reporting rate for observation	$[0, 2]$	0.71	$< 10^{-2}$	0.59 – 0.80
	2 Virus initially introduced (*2Vi*)				
$R_0^1 = \beta_1/\nu$	basic reproduction number (1st virus)	$[1, 50]$	1.96	$< 10^{-2}$	1.50–2.73
$R_0^2 = \beta_2/\nu$	basic reproduction number (2nd virus)	$[1, 50]$	14.24	0.02	6.24–23.78
$1/\epsilon$	mean latent period (days)	$[0.1, 7]$	2.18	$< 10^{-2}$	1.23 – 2.98
$1/\nu$	mean infective period (days)	$[0.1, 11]$	2.32	$< 10^{-2}$	0.1* – 4.34
$1/\gamma$	mean temporary removed period (days)	$[0.1, 30]$	0.64	0.01	0.1*– 3.62
ρ	reporting rate for observation	$[0, 2]$	0.65	$< 10^{-2}$	0.56 – 0.76
	Mutation (*Mut*)				
$R_0 = \beta/\nu$	basic reproduction number	$[1, 50]$	9.16	0.02	5.05 – 16.39
$1/\epsilon$	mean latent period (days)	$[0.1, 7]$	2.38	$< 10^{-2}$	0.1* – 3.00
$1/\nu$	mean infective period (days)	$[0.1, 11]$	1.01	$< 10^{-2}$	0.12 – 2.86
$1/\gamma$	mean temporary removed period (days)	$[0.1, 30]$	1.07	0.02	0.1* – 5.61
$1 - \sigma$	cross-immunity between the two strains	$[0, 1]$	0.82	$< 10^{-2}$	0.64 – 0.90
ρ	reporting rate for observation	$[0, 2]$	0.69	$< 10^{-2}$	0.59 – 0.79
T_{mut}	emergence time of the new variant	$[8/21, 9/8]$	8/24	–	8/21* – 8/27
	Intra-Host recrudescence of infection (*InH*)				
$R_0 = \beta/\nu$	basic reproduction number	$[1, 50]$	25.43	0.08	12.62 – 50*
$1/\epsilon$	mean latent period (days)	$[0.1, 7]$	2.45	$< 10^{-2}$	1.21 – 2.66
$1/\nu$	mean infective period (days)	$[0.1, 11]$	4.97	0.01	2.68 – 11*
$1/\gamma$	mean time before in-host reinfection (days)	$[0.1, 30]$	15.90	0.02	10.13 – 20.12
α	probability to develop long-term immunity	$[0, 1]$	0.62	$< 10^{-2}$	0.52 – 0.68
ρ	reporting rate for observation	$[0, 2]$	0.75	$< 10^{-2}$	0.63 – 0.84
	Partially-Protective-Immunity (*PPI*)				
$R_0 = \beta/\nu$	basic reproduction number	$[1, 50]$	6.92	0.02	3.99 – 15.12
$1/\epsilon$	mean latent period (days)	$[0.1, 7]$	1.84	$< 10^{-2}$	0.1* – 2.55
$1/\nu$	mean infective period (days)	$[0.1, 11]$	1.04	$< 10^{-2}$	0.1* – 3.30
$1/\gamma$	mean temporary resistance period (days)	$[0.1, 30]$	5.72	0.01	2.74 – 9.38
$1 - \sigma$	partial protection induced by immunity	$[0, 1]$	0.83	$< 10^{-2}$	0.67 – 0.92
ρ	reporting rate for observation	$[0, 2]$	0.68	$< 10^{-2}$	0.55 – 0.75

11 Table S12: Initial conditions

Table S12: Initial conditions were also inferred via the MIF algorithm as explained in reference [9]. We divided the initial population ($N = 284$) into three classes: S, I and L. Individuals in the L class are supposed to be either totally immunized against the virus or not available during the time-course of the epidemic. We inferred $S(0)$, $I(0)$ independently and we fixed $L(0) = N - S(0) - I(0)$. The upper bound for $I(0)$ corresponds to the initial number of passengers landed on TdC. *For the *2Vi* model we inferred the initial number of infected individuals for both viruses (low/high transmissible virus).

Model	$S(0)$	$I(0)$	$L(0)$
Bounds	$[0, 284]$	$[0, 5]$	–
2Vi	277	3/1*	3
Mut	278	1	5
AoN	277	1	6
PPI	277	1	6
Win	277	1	6
InH	276	1	7

12 Figure S10: Dynamics of the *2Vi* model

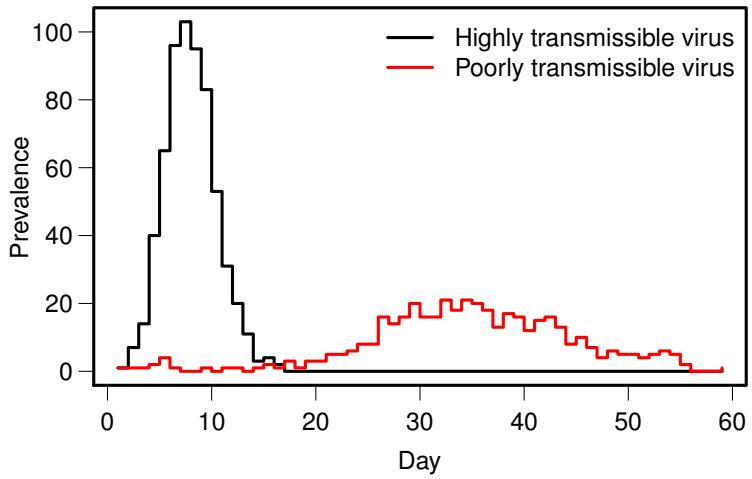

Figure S10: Evolution of the prevalence for the highly (resp. poorly) transmissible virus is plotted in black (resp. in red) for one typical realization of the *2Vi* model with maximum likelihood parameter estimates of Figure 4 (see main paper). The first epidemic wave is mainly due to the highly transmissible virus ($R_0^2 = 14.24$) while the second epidemic wave is exclusively due to the poorly transmissible virus ($R_0^1 = 1.96$). Because of the competition between the two viruses during the first epidemic wave, the poorly transmissible virus is just able to maintain a low prevalence, resulting in a 35% of chance of going extinct before initiating the second epidemic wave (see Figure 3 in the main paper).

13 Figure S11: Dynamics of the *AoN* with Mathews' parameter estimates

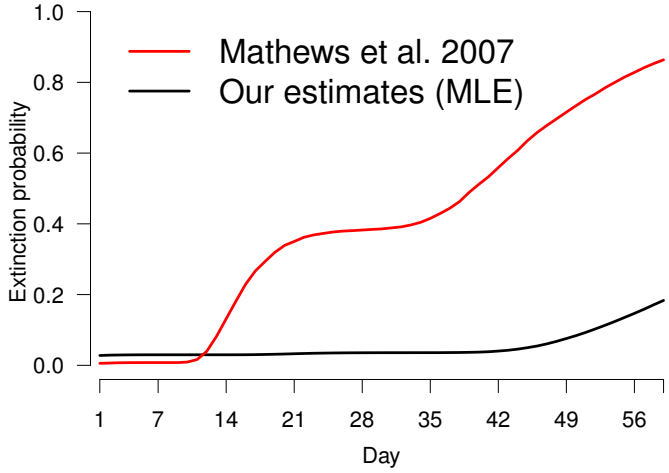

Figure S11: Effect of incorporating demographic stochasticity on parameter inference. We compare the evolution of the extinction probability of the stochastic *AoN* model when it is ran i) with the ML parameter estimates obtained via MIF and presented in Figure 4 of the main paper, ii) with the parameter set obtained by Mathews et al. [23] with a similar but deterministic model and using MCMC. Since our modelling approach takes demographic stochasticity into account, the maximum likelihood procedure implicitly try to minimize the inter-wave extinction probability. In contrast, the deterministic model of Mathews et al. [23] is insensitive to such stochastic extinctions and leads to parameter estimates that are far from being optimal regarding the chance for a second epidemic wave to occur. Parameter set of Mathews et al.: $R_0 = 6.44$, $1/\epsilon = 1.36$ days, $1/\nu = 0.98$ days, $1/\gamma = 12$ days and $\alpha = 0.49$ with initial conditions: $S(0) = 239$, $I(0) = 2$ and $R(0) = 45$ individuals.

References

[1] Mantle J, Tyrrell DAJ (1973) An epidemic of influenza on Tristan da Cunha. *J Hyg* 71:89–95.

[2] Rios-Doria D, Chowell G (2009) Qualitative analysis of the level of cross-protection between epidemic waves of the 1918-1919 influenza pandemic. *J Theor Biol* 261:584–592.

[3] Raue A, et al. (2009) Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics* 25:1923–9.

[4] Andreasen V, Lin J, Levin SA (1997) The dynamics of co-circulating influenza strains conferring partial cross-immunity. *J Math Biol* 35:825–842.

[5] Gog JR, Swinton J (2002) A status-based approach to multiple strain dynamics. *J Math Biol* 44:169–184.

[6] Ballesteros S, Vergu E, Cazelles B (2009) Influenza A gradual and epochal evolution: insights from simple models. *PLoS One* 4:e7426.

[7] Gillespie DT (1977) Exact stochastic simulation of coupled chemical reactions. *J Phys Chem* 81:2340–2361.

[8] R Development Core Team (2009) R: A Language and Environment for Statistical Computing.

[9] Ionides EL, Bretó CM, King AA (2006) Inference for nonlinear dynamical systems. *Proc Natl Acad Sci USA* 103:18438–43.

[10] Bretó C, He D, Ionides EL, King AA (2009) Time series analysis via mechanistic models. *Ann Appl Stat* 3:319–348.

[11] Ionides EL, Bretó C, King AA (2007) in *Statistical Advances in the Biomedical Sciences*, (John Wiley & Sons, New York), pp 123–140.

[12] King AA, Ionides EL, Bretó CM, Ellner S, Kendall B (2009) pomp: Statistical inference for partially observed Markov processes (R package).

[13] Akaike H (1974) A new look at the statistical model identification. *IEEE Trans Autom Control* 19:716–723.

[14] Burnham K, Anderson DR (2002) *Model selection and multimodel inference: a practical information-theoretic approach*. (Springer-Verlag, New York), Second edition, p 496.

[15] Gelman A (1996) Bayesian model-building by pure thought: some principles and examples. *Statistica Sinica* 6:6–215.

[16] Becker N (1977) Estimation for discrete time branching processes with application to epidemics. *Biometrics* 33:515–522.

[17] Diekmann O, Heesterbeek J (2000) *Mathematical Epidemiology Of Infectious Diseases: Model Building, Analysis And Interpretation* (John Wiley & Sons, New York), First edition, p 320.

[18] Smith DJ, et al. (2004) Mapping the antigenic and genetic evolution of influenza virus. *Science* 305:371–6.

[19] Koelle K, Cobey S, Grenfell BT, Pascual M (2006) Epochal evolution shapes the phylodynamics of interpandemic influenza A (H3N2) in humans. *Science* 314:1898–903.

[20] Fraser C, et al. (2009) Pandemic potential of a strain of influenza A (H1N1): early findings. *Science* 324:1557–61.

[21] Wallinga, J. & Teunis, P. 2004 Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *Am J Epidemiol*, **160**, 509–516.

[22] Svensson, A. 2007 A note on generation times in epidemic models. *Math Biosci*, **208**, 300–11.

[23] Mathews JD, McCaw CT, McVernon J, McBryde ES, McCaw JM (2007) A biological model for influenza transmission: pandemic planning implications of asymptomatic infection and immunity. *PloS one* 2:e1220.